6 Enterprise information risk management:

Dealing with cloud computing

Adrian Baldwin

HP Labs, Bristol, UK

David Pym

University of Aberdeen, UK

Simon Shiu

HP Labs, Bristol, UK

Abstract Managing information risk is a complex task that must continually adapt to business and technology changes. We argue that cloud computing presents a more significant step change, and so implies a bigger change for the enterprise risk and security management lifecycle. Specifically, the economies of scale that large providers can achieve are creating an ecosystem of service providers in which the marketplace (rather than consuming enterprises) determines security standards and properties. Moreover, the ability to consume high-level services from different environments is changing the nature of one-size-fits-all security policies. At HP Labs, we are doing research on developing trusted infrastructure that will exploit and improve security management in the emerging cloud architectures. We are developing and using economic and mathematical modelling techniques to help cloud stakeholders make better risk decisions, and we are pulling these pulling these strands together to establish principles and mechanisms that will improve and enable federated assurance for the cloud.

6.1 Introduction

Managing IT risks remains a significant challenge for most companies, yet most companies are ever more reliant on IT. A typical company will have a vast number of activities, policies, and processes that help manage and mitigate digital risks. Ideally, these would be viewed as part of a coherent strategy wrapped around a security or risk management lifecycle. The complexity of the IT stack from network through to application and users means, however, that tasks tend to be carried out in isolation. The typical risk management lifecycle involves risk assessment, setting policies to mitigate these risks, implementing controls and running systems in accordance with these controls, and monitoring and audit to ensure risks are mitigated. The monitoring of risks can provide better information to understand the emerging risk situation. Moreover, having an integrated view of risk across the lifecycle and technology stack should improve risk management so that risk assessment can adequately assess technical controls and human behaviour, and the technical controls can be designed to be responsive to the changing risk landscape.

Cloud computing is not just another technology evolution to which this lifecycle must react. Rather, it brings a fundamental shift in how IT services are procured and provided. In this chapter, we argue that the use of cloud moves a company from a position in which it is largely in control of how it manages IT risks to one in which it is reliant on others as stewards of its information, charged not only with caring for its basic security information security — confidentiality, integrity, and availability (CIA) — issues, but also with respecting its objectives and ethics. We believe that cloud computing will become far more than a scalable computing platform and that an ecosystem of business-level cloud services will emerge [8,44,45]. As this ecosystem grows, it will enable companies to adapt their business models, based on innovation in the business services ecosystem. This will only be possible if companies are assured that cloud service providers will act as good stewards of their data and regulators ensure that the overall ecosystem is sustainable and resilient to shocks [8,44,45].

Many companies have outsourced parts of their IT operations, and even their IT and risk governance functions. This has the effect of breaking up the lifecycle, with each service provider taking responsibility for different aspects. Contractually, however, the company remains in control. Cloud is different. Each cloud service provider must scale its operations to run services for many companies. Indeed, this is how we expect to cost benefits to be gained, and has the implication that cloud services will be standardized, with terms and conditions defined by the service provider rather than negotiated between a company and an outsourcing service provider. From the perspective of engagement with a services ecosystem, there is a change in the procurement model. No longer do we see an IT stack that operates under a set of security policies; instead, we procure business IT services with appropriate terms and conditions. This enables smaller service providers to participate, running their services based on cloud platforms provided by the large IT service providers. This already creates a service supply chain which can become more complex as services are bundled. For example, a complete financial operations service could be offered by combining the offerings of smaller service providers running accounts payable, accounts receivable, and general ledger services. Complex supply chains complicate the stewardship concerns [8,44,45].

In this chapter, we consider how enterprises currently manage their IT risks and how this will need to change as cloud computing emerges and is adopted. Cloud adds complexity to the information risk lifecycle. Companies no longer control all the security activities, or even have visibility of them, and, when a company uses an array of cloud services, each will have its own sets of policies and procedures, and be based on different underlying technologies. We must therefore develop richer ways of assessing and managing information risk. We see three significant areas that must be addressed:

1. **Understanding risk:** Moving to cloud will remove control and flexibility from the users of services, meaning better risk planning must be achieved prior to contract negotiation and service initiation. From an enterprise perspective, we have been using ideas from economics, mathematical modelling, and simulation techniques to gain a better understanding of risk [46]. We provide a case study that uses these techniques showing how they help security decision-making and discuss the approach [11]. As the cloud develops businesses will rely on a complex set of interrelated services and as such when managing risk we need to understand the resilience of the overall ecosystem [25]. We will discuss how we are extending our current modelling approach to help these new kinds of risk analysis and decision [8,56].
2. **Need for monitoring and assurance:** Security may well be improved in the cloud, particularly for companies who lack a mature security management methodology, provided companies understand that their risks are being adequately managed. As IT functions are spread across the cloud, companies will need not only event monitoring systems that cross the cloud boundaries, but also assurance systems that demonstrate that each service provider is maintaining their required security policies and that the combination adequately manages risk. Here we may see a movement to automated audit and the sharing of audit information [3,4,5] rather than the expensive manual audit process.
3. **Better Infrastructure**: It is hard to get an accurate picture of what is happening through monitoring and assurance. An alternative approach is to have an infrastructure layer that enforces policies and provides attestation. Here we look at how developments in trusted infrastructure (TI) will change the rules for risk within the cloud. We will draw from previous work on TI see [43], where virtualization and TCG [47,54] are combined to provide mechanisms to attest to system properties, draw boundaries around services, and allow policies to control data flow. We have recently started a 'Trust domains' project[1] to build on these ideas, further exploring the required technologies and linking it to the previously mentioned modelling and simulation. These technologies could be used to create trust domains, with predictable expected behavior, that span multiple service providers and so help re-establish a company's control of its risk lifecycle.

Section 6.2 starts with a general overview of typical enterprise architecture and operations and uses this to frame and discuss the risk and security lifecycle. We then focus on risk analysis and decision making, covering standards and state of the art, and leading on to case studies we have done using techniques from economic and mathematical modelling. The main contribution of the chapter is in Section 6.3. We start with a discussion of how cloud changes the kinds of risk analyses and decisions that must be made. Section 6.3.2 provides a significant description of information stewardship and why we think it is an important expansion of information security. In Section 6.3.3 we provide an example where we have used a real options switching model to help enterprises frame the problem they have when they consider using cloud services. Much of the value of this

[1] 'Trust Domains' is a collaborative project funded by the UK's Technology Strategy Board and EPSRC. It is led by HP Labs and includes the Universities of Aberdeen, Birmingham, and Oxford, and Perpetuity Group.

model is that it frames the uncertainty that business and IT managers must handle — which we relate back to stewardship. In Section 6.3.4, we review our early work, looking at how to model the ecosystem and how this will facilitate decision making for all cloud stakeholders. Section 6.3.5 considers and shows the impact that trusted infrastructure will have on security and risk analyses. The new layers of relationship and technology will mean providing assurance will be both harder and more important. Section 6.3.6 relates both the stewardship and architecture research to challenges and solutions for assurance. Finally, 6.3.7 ties all these points back to the risk lifecycle we see for cloud computing. Section 6.4 describes our future directions, which largely tie in with our ongoing collaborations in the Cloud Stewardship Economics and Trust Domains projects, both funded by the UK Technology Strategy Board.

6.2 Background

Within this chapter we argue that cloud will fundamentally change the way in which enterprises consume IT, radically changing aspects of their current security lifecycle and security decision-making. We start by looking at current enterprise IT architecture and how this will change as business services emerge in the cloud. We then review the current best practice for enterprise security management along with research aimed at improving the security decision-making

6.2.1 Enterprise Architecture and Cloud

Before looking at risks and the way the security management lifecycle must change, it is useful to consider the current enterprise IT stack and the transformations of it that may happen with cloud. Most large companies will have built up a complex mixture of legacy data centers, infrastructure, and applications. Many will have gone through centralization and consolidation efforts, which will have produced rationalized and documented enterprise architectures [2,51]. Even if not, the IT layers and management controls described here are fairly typical and indicative of the architecture. Our point is that, more so than previous technology and service trends, cloud computing is radically changing this architecture.

A typical company will have a set of IT services supporting its business processes such as finance, supply chain and order management processes. These may be standard applications although they will often have undergone considerable customization to fit with the company's business processes. That is, new applications will be rolled out to support new or changing business processes, the number of interfaces between applications will grow, and different applications and components will often be administered by different IT teams.

Many of these enterprise applications sit on middleware platforms that provide a variety of services such as identity management, messaging, and databases. These services are often critical to the security of the enterprise and can be difficult to manage effectively. For example, frameworks such as COBIT [36] commonly identified access management as a risk, but since most enterprises have a complex mix of centralized provisioning and single-sign-on, with distributed application and component access control lists, it is often difficult to provide sufficient assurance [4].

Below the middleware sit the datacentres that provide compute power and storage. The datacentre administrators will often run the back-end operating systems, manage the physical hardware and its security, as well as run some pieces of middleware, such as databases. A large company will have a number of geographically dispersed datacentres to enhance resilience. Companies will typically have a separate IT teams running each of the network and the client systems (laptops), and another managing their operating systems.

There are standards that help companies manage their IT stack. ITIL [41] sets out a number of strategic planning and operational processes that should be followed to ensure the smooth running of an enterprise IT system. The security team will set policies across all these IT layers to ensure that risks are mitigated. They will often work with standards such as COBIT and ISO27000 [35] to help them design a comprehensive security management system. In addition to considering all these layers and teams, the setting of a security policy is a negotiation between the business, the operational staff, and the security team. Many enterprises have risk committees to review decisions and provide a way to informally discuss trade-offs between the different needs of a business.

We can see the structure of cloud computing emerging in a similar way; see Figure 6.1. Looking at the NIST definitions [42], we have Infrastructure-as-a-Service (IaaS) providing the basic datacentre capabilities; that is, the compute power and storage. As we move up the stack, we have Platform-as-a-Service (PaaS); that is, where

the service provider runs middleware on top of the infrastructure. Then we have the Software-as-a-Service (SaaS) layer providing the applications to run the business processes. Some would argue that we should talk of Process as a Service which allows a number of applications to be combined to support a business process or even of Humans-as-a-Service (HaaS) where services are augmented with people's skills.

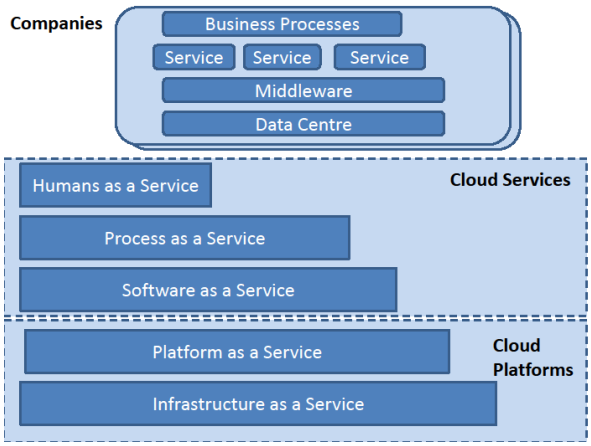

Figure 6.1 The structural components of enterprise IT and cloud computing

Traditionally, many companies have struggled with running their own IT systems — or do not see it as their core competence — and this has led many to follow IT outsourcing strategies. That is they hire a company to take on their IT systems and deliver them as a service. Typically, this will be involve a big deal and bespoke contracts will be drawn up setting out terms and conditions that meet the customer's needs [20]. When needs change, these terms and conditions can be changed, albeit at a cost.

Cloud provides a very different model of control. Cloud providers aim to provide scalable services at low cost — much lower than through outsourcing — and this can only be achieved by offering the same application and terms and conditions to many customers. Thus a cloud service provider may offer a small menu of choices rather than designing and running bespoke services. The result is a customer can no longer control the terms and conditions, including security policies, used in running a service. Control over this has switched from the customer to the service provider and customers must rely on service providers to be good stewards of their data.

Companies seeking to move their IT operations to the cloud could consider just moving their applications onto an infrastructure or platform as a service provider's systems. This would relieve them of the need to buy hardware and help them scale their compute and storage needs. One of the current models for enterprises using the cloud is to offload compute-intensive tasks. In these cases, security (i.e., confidentiality and integrity) and availability (cf. sensitivity and criticality) are not normally critical. An alternative usage model is where a company will seek to get some of its enterprise applications from the cloud; for example, looking to salesforce.com to provide a customer relationship management system.

Rather than going to the basic platform providers, a company might directly procure its business services. This would relieve companies of many of their IT functions although they would still need to perform network and client management. As this happens, the ability to provide robust messaging and identity services across a range of different business-level services becomes critical.

Considering how cloud systems have been emerging, and in the interests of conceptual clarity, we have chosen to work with a simplified three layer model as a basis for our discussion; see Figure 6.2). Here we combine the infrastructure- and platform-as-a-service into a cloud-platform-provider layer. We believe there will be a limited number of companies offering such services due to the massive investment costs needed to build datacentres. The presence of such cloud platforms will allow a large number of small software developers to offer their products as services in a cost effective manner. The third layer of our cloud model is companies that consume services[2] or platforms. We develop this model further both in [8] and later in this chapter.

[2] Throughout this chapter, we are concerned with how companies use the cloud and ignore consumer cloud services.

6.2.2 Pre-cloud Enterprise Risk Analysis and Decision-making

Even without the cloud, year on year enterprises see more failures and attacks, meaning they are compelled to spend more on information security. However since resources are limited, organizations need principles and guidance for how much and where to spend on security. The common answer is that it should be based on risk — that is, focus resources on highest impact and likelihood events.

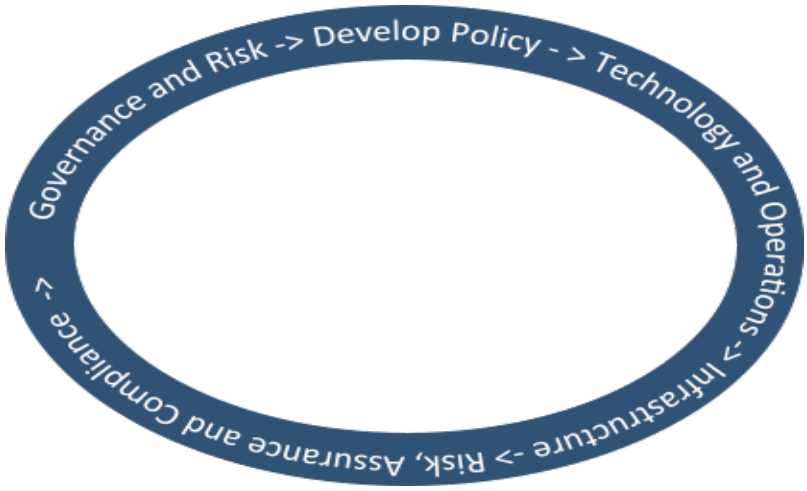

Figure 6.2 The typical enterprise risk management lifecycle

Ideally, an organization would drive all security investment and operations from a risk perspective, but a challenge to this is that there are many other stakeholders, incentives and processes with which risk choices have to live. That is, in theory (and to some extent practice) an organization will follow a lifecycle such as that in Figure 6.2, whereby business driven risk assessments set the context and priorities for security controls, policies and investments, which in turn guide infrastructure procurement and operations, and the monitoring and audits of the IT environments test whether the controls are effective and mitigating risks thus closing the loop. There are many standards and frameworks to help govern and apply best practice through the lifecycle, such as

- the ISO27000 [35] series of standards for information security that provide a framework referencing known security best practices, and how to organize them in a coherent governance framework, and
- COBIT [36] that sets out control objectives for IT from a business perspective and complements the slightly more technical focus of ISO27000.

Complementing these there are several methodologies [52] that aim to help organizations perform systematic evaluations of their IT risks. Typical steps include:

- A scoping phase to determine the nature of the risk; that is, assets, components, and boundaries. For digital information and services that depend on many layers of abstraction, and distributed inter-dependent systems this can be a difficult and subjective task;
- A threat analysis, where the attacks, motivations, opportunities, and vulnerabilities are considered;
- An analysis of the likelihood and impact of any of the threats occurring, which in turn guides prioritization and choices as to whether to accept, mitigate, or transfer risks.

The standards and frameworks are extremely useful and relevant. In practice, however, there are still many challenges applying the described principles in specific complex environments, primarily because

- the teams procuring and running infrastructure are under pressure to improve and maintain service levels which often work against the risk priorities,
- performing risk analyses that take account of all the trade-offs is inherently difficult,
- the enterprise architecture and processes are always changing, making assumptions made during risk analysis out of date, and
- many complex IT environments rely very heavily on human behaviour — and we lack rigorous ways to incorporate this into risk models.

There is a lot of current research looking at many of these problems. See [49] for examples of work seeking to integrate insights from human behaviour or psychology. [28,14] are specific examples of recent research on how better to integrate knowledge of human behaviour into the analysis of risk. There is also research on the

economics of information security, ranging from showing how to apply standard (business understood) return on investment cases for security investment [30] to the application of sophisticated utility functions [31].

HP Labs has developed the idea of combining economic and system models to help organizations with risk assessment, security analysis, and decision-making;[3] see [46]. Economic models, represent as utility functions, are used to help stakeholders think about and share their preferences and priorities for different business outcomes. We then use structural models to help stakeholders think about and share their assumptions for how different investment and policy choices will affect the outcome. Finally, we use a discrete process simulation tool [22,24] that allows stakeholders to explore and predict consequences of their different assumptions. Figure 6.3 provides a schematic of this methodology.

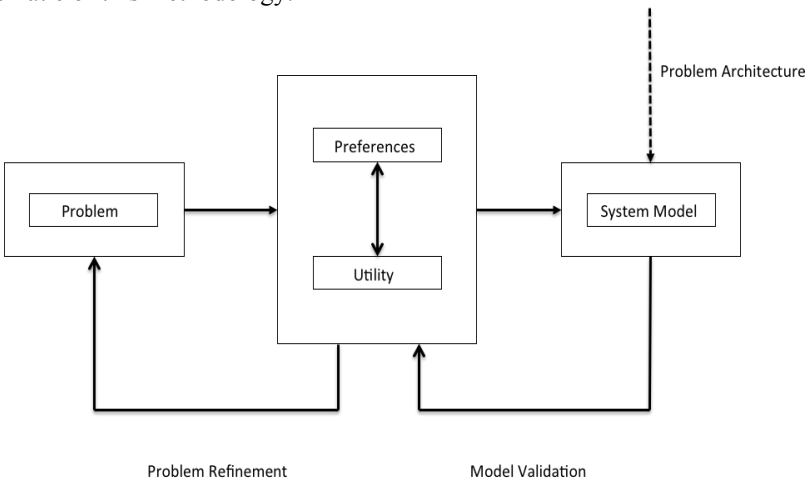

Figure 6.3 A framework for using economic and system models to support organizational decision-making

HP Labs has also conducted a series of customer case studies [7,48] to develop and refine this process. An early example was to help a large enterprise decide between a range of policies and investments to manage risks from software vulnerabilities, see [11].

Most large organizations have evolved a complex set of processes and technologies to deal with software vulnerabilities. These include testing and deploying patches over multiple environments and regions, deploying and updating anti-virus signatures, reliance on gateway and network boundaries and protections, intrusion prevention systems, processes to accelerate patching processes and so on. The risk question was whether to spend more resources to improve the effectiveness of all these controls and, if so, where. More specifically, should the organization invest in some new host-based intrusion-prevention technology, invest to improve the patching process, or spend the money on other (perhaps non-security) projects.

The first task is to find some components and abstractions that to help the stakeholders to think about the complex system. In the given study, it was decided that it would be useful to model how long (typically) it takes to mitigate risks from a known vulnerability. The team separated aspects that were under the control (or influence) of the organization (how fast they patch, when signatures would be deployed) from external factors (when vulnerabilities are publicly disclosed, when vendors release patches, when malware starts spreading). They then separated out concurrent processes that effect risk mitigation (i.e., testing, patching, signature updates) and to discuss how architecture and decisions (typically) affect their progress. The result was something like that which is represented in Figure 6.4.

[3] Much of this work was based on the UK Technology Strategy Board-funded Trust Economics project, with partners from University of Newcastle, University of Bath, University College London, Merrill Lynch and National Grid.

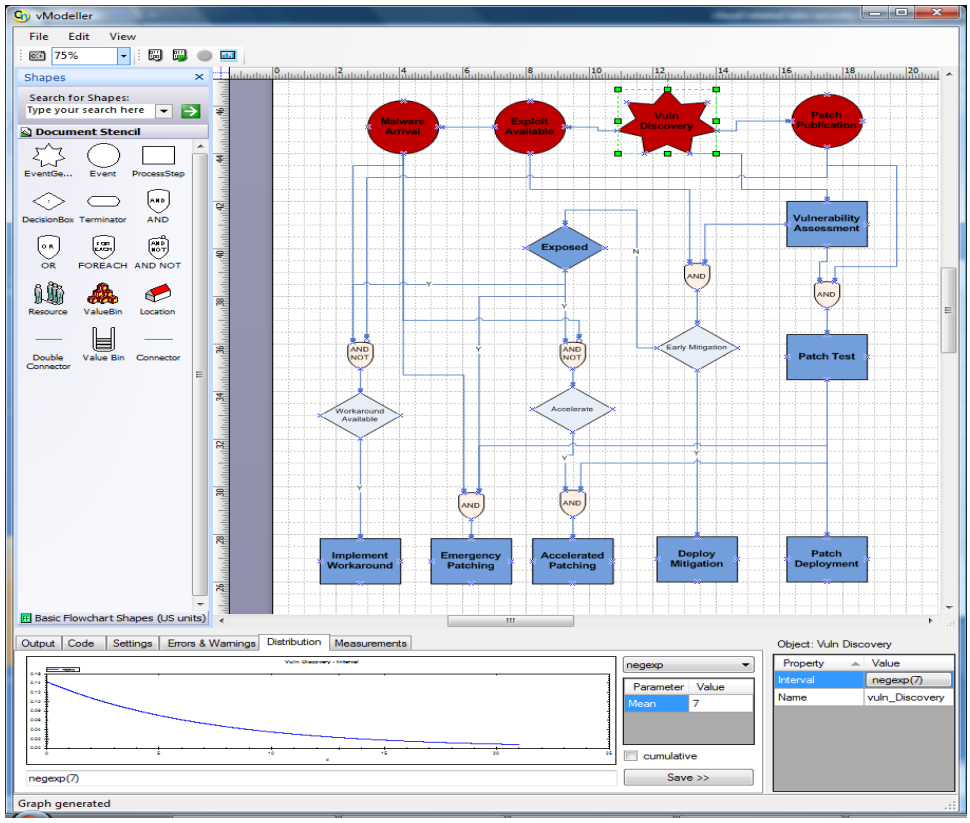

Figure 6.4 Structural components of a typical vulnerability management environment

The structural models already help ensure shared understanding between stakeholders, so they can discuss, say, whether scheduling is significantly delaying patch testing, or when and how often the assessment team accelerate patch processes. However, with such a complex system of inter-dependent concurrent processes, and actions it can be very difficult to see or reason about the cause and effects. To address this (using best available assumptions and empirical data) we use Gnosis [21,22,23,24] to create an executable mathematical model of the system.[4] Using Gnosis we can run Monte Carlo-style simulations to explore the interactions and their effect on time to mitigate risks. By varying parameters stakeholders can see the (model) predicted effect of different investment choices. Results are typically shared as histograms showing, say, the difference in time taken to mitigate risk for different investment choices. For example the results show quantitatively how an investment in HIPS should increase the number of early mitigations, whereas similar investment in patching will reduce the long tail of vulnerabilities that take a long time to mitigate. This would be a simplistic initial result, and further experiments can explore the effect based on different assumptions about the threat environment, or to differentiate on different types of mitigation.

So far, this example has only discussed the effect on risk mitigation. Most security decisions involve multiple trade-offs between mitigating different kinds of risks, maintaining services, and minimizing costs. To formulate this, we encourage stakeholders to define a utility function expressing their preferences between the multiple outcomes, building on the approaches to decisions with multiple objectives developed by Keeney and Raiffa [39]. We have developed some simple tools for preference elicitation and then use the system models to explore the effect of different security choices on these other outcomes, see [13].

Our experience is that focusing, via systems models, on the utility (of outcomes) provides a constructive way to engage multiple stakeholders (with different knowledge and incentives) in the complex process of risk assessment and choosing security investment and policy. Providing evidence for this is difficult as organizations, people and problems vary so much. We have done some preliminary studies that suggest our methodology affects the justifications security professionals might use, which fits with why it might be useful

[4] Gnosis is a discrete process modelling language that (partially) captures a discipline of mathematical systems modelling based on mathematical models of the concepts of location, resource, and process (all modelled using algebraic/logical tools) and environment (modelled stochastically) [21,22,23,24].

for multi-stakeholder decision making, see [50]. We are currently looking at further cognitive studies to generate more precise hypotheses for how and why economic and system modelling affect security decisions.

Part of the problem is that many security problems (like patching components and network security) are about mitigating risks on infrastructure that many applications rely on. Business stakeholders find it easier to reason about motivations and impacts of application failures, rather than on the complex dependency of many applications on shared infrastructure. One helpful outcome (from a risk perspective) of the shift to cloud computing is the ability to consume software-as-a-service from multiple environments. Although there will still be complex inter-dependencies between applications, it is helpful for a business to analyze discretely the risks for different applications or, conversely, to be able to look at the impact of infrastructure failures in the context of the only/few applications running on it. We shall return to this point when we discuss trusted infrastructure, which also allows us to de-couple risk analyses at different layers, so simplifying the whole problem.

This section has described some of the challenges enterprises face in aligning security policies and investments to business priorities. We have argued that, in some cases, the shift to cloud computing may disaggregate typically complex enterprise architectures (and so simplify risk analysis). Conversely, however, in the current lifecycle (even when outsourcing), the organization at least maintains control of all the components and services. The emerging problem with cloud computing is that organizations will lose this control, and will rely on a number of firms making different choices about how the applications and infrastructure holding their data and executing their transactions will be defended. It will be a bigger challenge to help businesses create effective risk and security strategies whilst tapping into the flexibility and cost structures of cloud computing.

6.3 Risk and Security Management in the Cloud

Moving to the cloud does nothing to reduce dependence on IT; instead, it means that a company is dependent on service providers to do the right thing and act as good information stewards. That is, enterprises must rely on others to manage their information and the processes that create, maintain, and use the information. The shift of control over policy, operational, architectural, and assurance options from the customer to the cloud service provider means that the customer must employ careful risk planning to choose the service that offers the best trade-off between the service provided, the level of risk, and the cost.

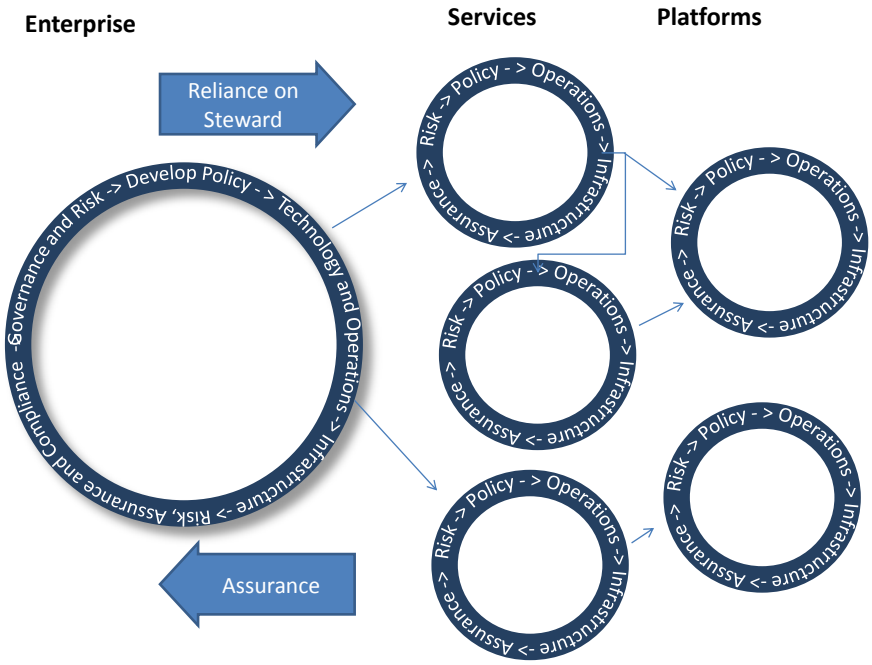

Figure 6.5 Security lifecycles where an enterprise must rely on other services to be stewards who must run a good security management lifecycle; these in turn may depend on a service supply chain; assurance information should flow back

Instead of running one security management lifecycle, an enterprise now must rely on a multitude of service providers, each running their own security lifecycle (see Figure 6.1). An enterprise is then faced with two issues: firstly, to decide on whom to trust for what service, and how much information they need about that company's lifecycle; and, secondly, how to gain a coherent overview of its risks. These issues are at the heart of the transformation of the security lifecycle as enterprises move to cloud. Here we argue that as the lifecycle breaks up and changes, companies need better risk planning methodologies (Sections 6.3.1 to 6.3.4), better architectural support (Section 6.3.5) and better assurance (Section 6.3.6). That is each of the individual elements of the security management lifecycle need improving in such a way that allows a company to understand the risks to which each service is subject and hence get this joined up risk view. Here we look at each of these three areas separately before bringing the discussion back to the overall lifecycle.

6.3.1 Decision-making about Risk in the Cloud

Although companies will often have some form of IT risk planning, security is often reactive to events. This may be inevitable, both due to the 'arms race nature' of security as well as the fast pace of technology innovation and change. A company's ability to adapt its security to circumstances relies on the ability to control security policies and infrastructure. As cloud services are used this ability to react is very much reduced. As one of many customers of a service it can be hard to influence service offerings and switching between services can be expensive (especially if there are no standardized data formats). That is, there is a danger that customers get locked into a particular provider or *de facto* standard service models — hence the need for good initial risk planning.

Part of this risk planning must take account of the opacity of the service and what assurance information is available to the customer to ensure that the promised levels of service are maintained.[5] In performing this risk planning, companies and their security officers need a much richer concept of stewardship. The must consider more factors than just the confidentiality, integrity, and availability of services that they would manage to internally. In particular, we would argue that they must consider about the objectives of the service provision, the ethics of the associated business decisions, and the sustainability and resilience of the services on which they depend [8].

Earlier in this chapter, we discussed the current enterprise IT stack with separate services all built on top of horizontal middleware, datacentre and network offerings. Security has become hard in this world (arguably, it always was [1]) since a security policy in a datacentre must be sufficiently strong to provide the sought level of protection, but not too disruptive of enterprise applications. Communicating the technology risks and the business needs up and down this stack has often proved challenging. This interconnection and dependency has made it hard to take an economic approach to security because of the difficulties in articulating changes in service levels and risk levels within one-size-fits-all policies.

Cloud offers a different way to procure IT services. Now each service can be thought of as a separate entity. Cloud does entail a degree of loss of control and ability to react, and hence the need for more up-front planning, so complicating the decision-making aspects of the lifecycle. However, the ability to focus on a single service and drop the one-size-fits-all infrastructure security policies typically simplifies each decision. Looking at each service separately means we can get a much clearer understanding of the trade-offs between risk, service level and cost. Thinking of services in this way provides a more modular way to think of security hence simplifying decision-making. However, much of the complexity is hidden in the way the service is delivered.

As businesses start to use the cloud for critical business functions, we believe that an eco-system of service and platform providers will start to emerge. A company may just worry about the way service providers they use manage risk but they should also worry about the overall resilience within the ecosystem. A simple example of this need arises in the construction of service supply chains, where one service provider uses others to provide parts services and platforms to deliver the underlying IT infrastructure. Risks may emanate from failures of any of these services, even though such failures may be unknown to the end-user. Other risks can occur because of interdependencies between service providers and the resources (e.g., investment capital and skilled staff) upon which they rely. As with enterprise IT, where risks occur because of the complexity of the infrastructures, we believe modelling decisions in the cloud ecosystem can help us gain an understanding of these risks.

The emergence of a cloud ecosystem will represent a major shift in the way in which enterprises purchase IT provision. This means that the surrounding environment will change and, in doing so, it will cause shocks to the

[5] This is particularly important for security processes, where failures may not be obvious or not until a serious incident has occurred.

ecosystem. Criminals will start to see concentrated value in certain services and may invest considerable resources in attacks. Regulators will be concerned about the stability of companies and their ability to deliver services and, so, rules, regulations, and laws will eventually catch up with the emergence of cloud. In turn, enterprises thinking about the risks of procuring from the cloud will also be taking into account implicit assumptions about the sustainability and resilience of this ecosystem.

In the next section, we explore each of these themes and consider the use of models to help understand assumptions, risks, and decisions.

6.3.2 Information Stewardship

Typically, when thinking about security, people think of confidentiality, integrity, and availability (CIA) as a distinct and complete set of declarative properties [15] that, to varying degrees, system managers will seek to maintain. Pym *et al.* [32,33,34] introduce the idea that these elements can be combined along with cost to form a utility function for security decision-making that characterizes the relative extent and form of the managers' preferences between the various properties. When we start to think about how risks are managed around cloud, it is still useful to think about the CIA concepts but we need a wider conceptual framework to understand the implications of stewardship. Information stewardship for a cloud service needs to include additional concepts around management and duty of service (i.e., the appropriate achievement of the agreed objectives), the supervision of values and respect for ethics, along with the long-term sustainability of services and their resilience to rare-event shocks.

In choosing a service, a company must think about whether the service provider will act as a good information steward or, more accurately, a good-enough steward for the required service and at the given cost. In doing this, the risks associated with the service must be considered. A customer may place few stewardship requirements on a service storing encrypted data whereas the requirements for a service running his financial processes will be much greater. Some aspects such as the ethics of the service provider may have a wider consequence than just the service being used. A manufacturing company's reputation could be hugely damaged by the use of child labour for even a small part of one of its products. In the same way, the use of a service provider who appears unethical could cause damage to wider the ecosystem rather than just to service in question. As a first stage in looking at stewardship of the service, the service provider's identity must be checked. Are they whom we expect them to be, is the company financially solid, are they owned by our competitors? In establishing their identity, we must also establish the jurisdiction under which they operate. Some of these basic checks must be done before looking more deeply into the stewardship concepts.

The information steward is responsible for maintaining the usual CIA properties, along with maintaining data privacy, and as such we would expect that he would run the normal information security processes. For example, running vulnerability management processes, ensuring access to information is controlled and ensuring software and hardware is of sufficient quality. The effort put into each security process must be traded off against the cost, the service provided, the value of the data, and risks when things go wrong. As well as trying to maintain these basic security properties, the steward must communicate how much effort he will make in achieving this trade-off. Models such as the one described in Section 6.2.2 could help form the basis of this communication.

Establishing the boundaries of stewardship is also important. For what, exactly, is the steward taking responsibility? For an enterprise service being moved into the cloud, the service provider will be acting as an information steward both for the basic information and for the way in which the transactions are handled, but the enterprise will probably still remain responsible for managing the identities and rights of its staff access when they access the service. Care should be taken that the stewardship boundaries and hence responsibilities are clear.

Even with the best technologies, management and risk planning there will still be security incidents. The information steward will have responsibilities around managing such incidents and keeping the service users informed. However, the use of a steward will not reduce the accountability that the enterprise has for ensuring their data (e.g., PII that they hold) is secured. In looking at stewardship we need to look at what happens when things go wrong and how incidents can be jointly managed to reduce loss and damage to reputations. Incidents that happen to one customer using a service provider may cause reputational damage to others using the same service — so it is the duty of the information steward to maintain their reputation. Disclosure may become an important aspect of stewardship forcing those acting as information stewards to disclose incidents publically. This has become the case in the USA [55] for personally identifiable information where disclosure policies support a loose regulatory environment.

Businesses need services to run constantly or continuously over time and to remain fit for purpose. This means that they can be relying on services for long periods of time. In looking at the information stewardship offered by a service a customer should think about these long-term requirements:

- Will the service respond appropriately to changes in the threat environment?
- Will the service respond appropriately to regulatory changes?
- Will the service itself change as needs change? For example, we would expect an accountancy service to change its rules as accounting standards evolve;
- Are there good (cheap, efficient) exit routes if things go wrong or other changes are needed?
- What measures can be taken when things go wrong (arbitration, law)?
- What happens if the service provider is taken over or spun out as a separate company?

In addition to thinking about stewardship at the level of an individual service, we also need to consider the overall stewardship of the cloud ecosystem. Here we must consider the properties of sustainability and resilience [19] provided by the overall system. The influence needed for this level of stewardship will be beyond the individual participants, except perhaps the large platform providers. However, regulators, or clubs of service providers or users, may form in order to set rules and ensure appropriate overall stewardship of parts of the ecosystem.

6.3.3 Migrating to Cloud

As business-level cloud services start to emerge a business will be faced with a decision: should it keep using its current IT systems or should it switch to using cloud services and, assuming so, which cloud services to use. The decision may be to move all IT to the cloud but a more likely and recurring decision will be whether to keep an internal application running or to replace it with a cloud service. This decision will often be triggered by a business need to upgrade an application. Clearly information stewardship should be part of that decision but a company needs a framework for thinking through the decision.

In Section 6.3.2, we discussed the need to frame security decisions with a utility function [32,33,34] that allows stakeholders to trade off different aspects of security, such as confidentiality and availability, as well as considering cost. In Section 6.2, we discussed how security policies tend to apply across the board, so that one decision will affect many parts of the business. If cloud decisions relate to single services, then security decisions can become more modular, and so simpler. That is, when a security decision affects many different aspects of a business it can be hard to estimate the cost and productivity effects. Isolating these factors to a single service can help focus the decision. Conversely, moving a single service to the cloud may still have wider impacts on the wider enterprise IT. For example, moving many applications out from a company's datacentre will reduce the economies of scale, increasing costs for those that remain. In other ways, the cloud decision is more complex in that the decision to move encompasses factors representing business decisions, as well as cost and information stewardship factors.

Formalizing preferences in a utility function provides a framework for comparing the different values that a company would get between different cloud services options or the internal IT option. It allows an enterprise to express how it wishes to value a gain in one factor, such as business value, against a loss in a stewardship factor. It can then use this function to look at the different service options' terms and conditions.

This still leaves the question as to the optimal time to switch to a cloud service, assuming it offers better utility. That is, even if the cloud service offers better immediate utility, there may be greater longer-term utility in using the option to wait (until company finances are better, or until uncertainty relating to the value is reduced). A cloud switching model, using real option theory [53], has been developed in [56] to explore this question. The problem of whether to use a cloud service is expressed as a choice between staying with internal IT, switching to a cloud service, or opting to wait (and monitor).

An advantage of this framing is that much of the financial economics relating to discount rates and the time value of money can easily be reused. That is, most real option models take account of the fact that decision-makers will have different cash flow, levels of capital, risk appetite, and patience for a return on investments made. This makes it natural to show, for example, how a large oil and gas firm with huge reserves, used to long cycles of investment, will feel quite differently about utility than, say, a start-up having little capital and worried about whether it will be in business in 12 months' time.

Clearly, there is uncertainty over security and stewardship, and the switching model [56] allows these issues to be framed as affecting uncertainty over the value the business can achieve with each decision option. For example, an enterprise will have some (but not complete) confidence in internal IT's ability to keep up with regulations and the threat environment. They will likely be more uncertain about whether using the cloud

service will enable them to maintain security and compliance. Moreover they will implicitly be concerned about all the information stewardship points discussed in Section 6.3.2.

Finally, there is clearly potential value in having an option to wait to decide. As time passes it will become clearer whether early adopters are having success, whether security incidents are more common and whether it seems resilient and sustainable. For some firms the value will be high enough to warrant early adoption, for others it will be a clear no, and for others a wait and see approach will be appropriate.

It is not that the model provides accurate predictions about the outcomes, but rather that it frames the assumptions, so that stakeholders can debate the options and trade-offs appropriately. Moreover, it is also not that decision-makers are unaware of all these aspects to their decision options, but, as discussed in [50], there is significant value in bringing all the issues together in an appropriate way.

In discussing this model with cloud stakeholders, many raised the question of being able to explore the issues of lock-in (being tied to a particular cloud provider) and the ability (ease) of switching back to previous states. These can be treated as uncertainties within the current framework, but it is also natural to consider the economic models that explore precisely this situation, see [38]. As discussed in the 'Future Directions' section, below, our current work in this area is focused on testing and refining this type of model stakeholders, and as part of this we are looking at these iterative migration models.

As cloud services start to emerge, a company will need a framework for thinking about when it should start to use a cloud service and which service is the most suitable for it. Taking a utility theory approach forces the company to think about the different outcomes of the decision and how they may trade off against each other. The framing provided by the 'cloud switching' model [56] provides a way to think about the costs of moving service as well as uncertainties associated with the utility. Much of this uncertainty may come from the loss of control and the need to rely on others to be good information stewards.

6.3.4 The Cloud Ecosystem

A company can try to choose a cloud provider that best meets its needs, including one that it believes will be resilient to failures. Each service, however, sits within an ecosystem of other services, service customers, and cloud platforms, and their success and resilience may well be affected by the success and resilience of other services. Within the ecosystem there will be limited sets of available resources, for which each service provider must compete if it is to be successful, and each service provider will work hard to develop its own reputation. However, news of incidents — and the overall reputation of the cloud ecosystem — may swamp their branding. Finally, there will often be a service supply chain within which a service provider relies on other services and cloud platforms to allow it to deliver the contracted service. These factors mean that as a rich cloud ecosystem emerges, the success, sustainability, and resilience of an individual service will be dependent on those properties of others in the ecosystem.

In our analysis of this ecosystem [8], we draw quite significantly on research carried out on ecological ecosystems [19]. The ecological ecosystem consists of various organisms that exist in a habitat or a series of linked habitats. The ecosystem will be affected by the way the organisms interact (due to their biology) as well as due to external influences such as the weather, fires or pollution. In studying an eco-system and its dynamic behaviors we can start to see how resilient the ecosystem is to different shocks and hence start to manage it in a sustainable way. Analyzing from an ecosystem perspective helps us develop good stewardship properties.

Instead of organisms in various habitats we have an ecosystem consisting of customers consuming cloud services, cloud service providers offering the services and cloud platforms providing the basic infrastructure for these services (see Figure 6.1). Instead of the interaction between these entities being driven from their biology, it is driven by their need to maximize (or at least satisfice) their utility, so influencing their policies and decisions. This utility will usually be very implicit in each company's decision making but will drive customer's choice of services as well as the terms and conditions offered by the service and platform providers.

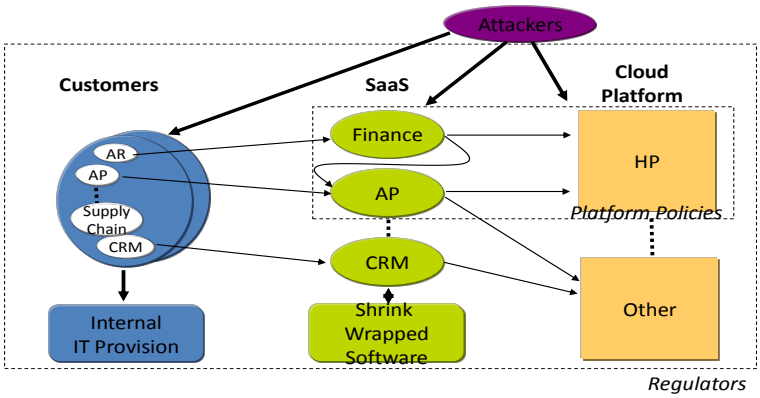

Figure 6.6 A three-layered cloud ecosystem

In modelling an ecosystem, one of the key questions is which entities should be included within the model and which can be treated as exogenous events. In an ecological system, weather events and the actions of the human population will often be treated as being exogenous. In our treatment of cloud ecosystems, we are keeping the effects of the overall economy, attackers, regulators, and technology changes as external to the overall ecosystem. We can consider how each of these external factors will affect entities within the ecosystem and how different economic, threat, and regulatory environments affect the sustainability and resilience of the ecosystem. In understanding the dynamic behaviour, we can start to think about how ecosystem stewardship helps maintain both sustainability in the course of normal operations and resilience in the face of shocks.

Ecologists consider how ecosystems vary overtime because of feedback loops. For example, a fast variable may be the population size of a particular animal, such as deer. This variable will determine how much biomass is eaten, which in turn determines the available food and reflects back into the population size. Slower variables may be things like changes in the capacity of soil or sediments to supply water or nutrients or changes in types of plants and animals in the ecosystem. Exogenous controls may be changes in the regional climate. Ecologists then consider two different factors as being responsible for these changes: the ecological factors and the societal factors (i.e., mankind's effect on the ecosystem). Within our view of cloud ecosystems, we can draw out similar feedback loops — see Figure 6.3 — drawing out the business or economic environment and technological environment, rather than, though analogously to, the societal and ecological factors described by ecologists.

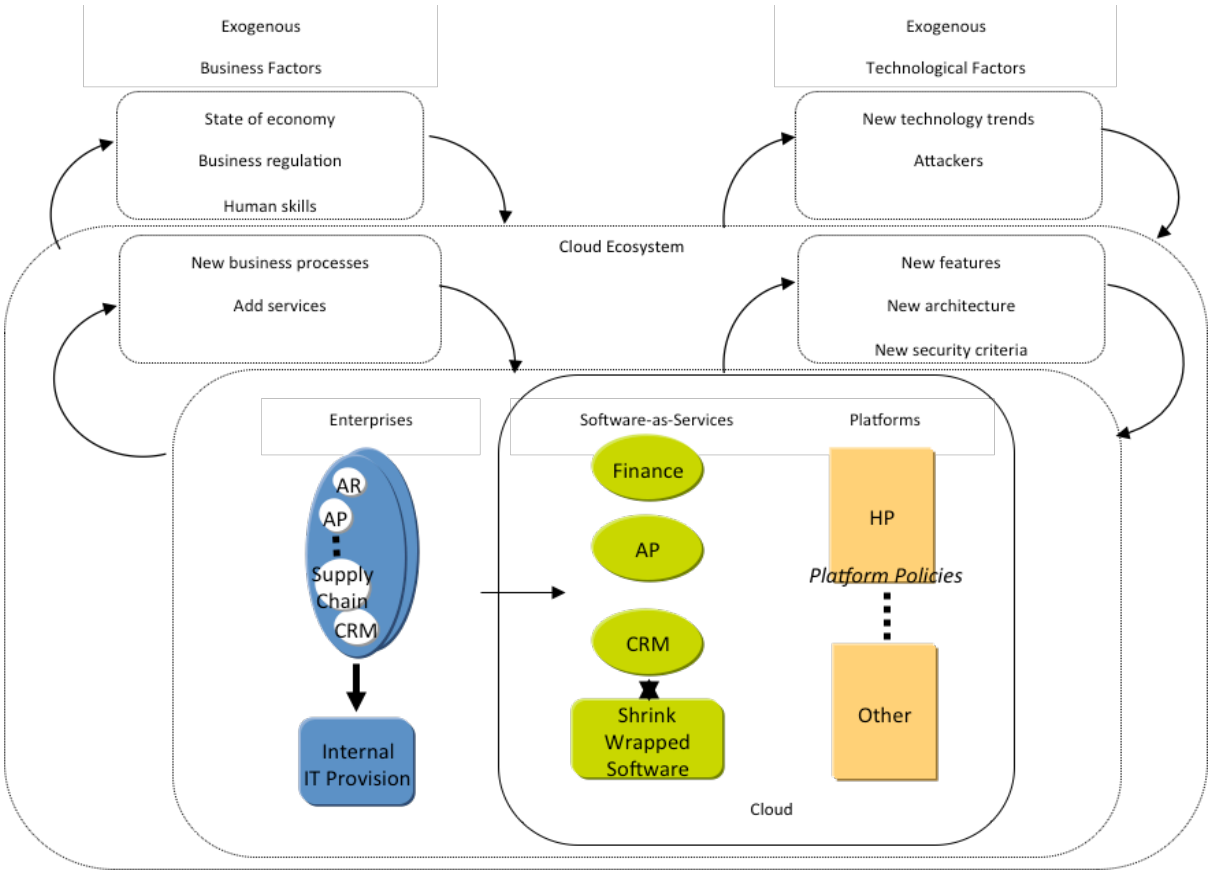

Figure 6.7 Dynamics of the cloud ecosystem

Here, in Figure 6.3, we can see the fast variables being linked with the regular IT decisions that enterprises make around their IT needs. Here they may make decisions to continue to run their own IT services or to move to the cloud. In making their decisions, they will seek to maximize (but, in practice, typically satisfice) their utility as discussed above. These decisions will be based on their needs as well as the state of the ecosystem. For example, their decisions will be based on the different costs, the the different services available in the cloud, the terms and conditions offered, and their beliefs in how good an information steward a service provider will be. Their decisions will of course affect the state of the ecosystem. For example, if a company decides to use a cloud service rather than its internal IT, it will release, or not require, resources such as IT staff and investment capital. On the other hand, additional network bandwidth may well be required.

In an equilibrium state, the resources moved by each of these decisions will cancel each other out. However, we can get reinforcing feedback loops that will help move the ecosystem to a different equilibrium state. For example, as a service provider starts to get more business it will be able to scale better and hence offer better or cheaper services. This will attract more customers as well as making it easier for them to get resources, such as skilled IT staff and investment capital. This reinforcement feedback will lead to a movement of IT from company's internal provision and into the cloud.

There are, of course, other feedback loops at this fast-variable scale. Some cloud services will fail either because they get insufficient business or because of stewardship failures (e.g., security incidents, failures to maintain services, failures in meeting the reputational needs of customers). Such incidents may lead to customers pulling services back from the cloud into their own datacentres. Other feedbacks will be caused by technology changes — for example, as the availability of new software features encourages companies to upgrade their systems. Where companies invest in creating new features in the cloud, or in shrink-wrapped software, this will help determine the cloud adoption rates.

One critical factor associated with how fast feedback loops work within the ecosystem will be the costs associated with changing provider, or moving back to internal IT. When a company chooses a particular service provider, it may be hard or costly to get its data out of that service and into the correct form for a different service; or there may be costs associated with integrating a new service into its business processes. High movement costs may mean that service providers have less incentive to update features and act as good stewards. It will also lead to a slowdown in the speed of these feedback loops.

Slower variables will be things like the addition of new services into the cloud ecosystem and the corresponding changes to companies' business models — an example would be how the first internet wave encouraged the creation of e-commerce websites and changes in the way in which many companies sell their products. These changes will occur as resources such as IT staff and investment capital become easily available to the cloud service providers, so encouraging innovation. Major technology changes will also lead to slower changes in the ecosystem. For example, were cloud service providers to use 'trusted infrastructure' (described late in this chapter) this would help service providers in being good information stewards, and hence assist the development of cloud ecosystems. Emerging standards and regulation changes will also change the way companies view and use cloud.

At a more global level, there are many influences that will affect the overall business environment, and hence the cloud ecosystem. For example, the overall state of the economy will determine many of the business needs to which each company responds, as well as determining how much investment capital is available. Governments may set up training programmes to ensure sufficient skills are available and support research and development programmes to help ensure that appropriate technology evolution occurs. Regulations around how businesses operate or around global trade may also change and reflect back into the ecosystem dynamics.

In thinking about risk in the cloud, we must consider the sustainability and resilience of the overall ecosystem and the effects that the normal evolution of the ecosystem and rare-event shocks may have on a given enterprise. A clear conceptual model of cloud ecosystems and their dynamics is a necessary prerequisite for allowing us to think through these effects. We can start to extend the system modelling approach [22] used to help us understand security decisions in the enterprise to understand the dynamics of the overall ecosystem, the effects of shocks and different ecosystem stewardship approaches. Pym *et al.* [8] describe such an approach to modelling based on a location, resource and process calculus [21,23] that has previously been used for systems modelling and security decision-making.

6.3.5 Trusted Infrastructure and Cloud

In much of this chapter, we have concentrated on the risk-side of the security management lifecycle and how this is affected as business operations move into cloud. Having good system architecture helps to reduce risk and can make reporting and assurance easier. Following our previous argument, as an enterprise moves from running its own applications to using cloud services, it not only loses control of the people, policy, and process parts of some aspects of its operations, but also of the technical architecture and the application code-base.

We have hypothesized a world in which there will be a relatively few cloud platform providers that provide the basic compute and storage platforms, along with much of the middleware, and service providers who will write software, run it on the platforms, and use it to offer business-level services. These services would be provided to many customers using common management processes and contracts. There are two natural technical architectures to support such multi-tenanted services: the first is to write the software to support multiple customers at the same time (application virtualization); the second is to run many instances of the software each within its own container (infrastructure virtualization).

The two different styles of virtualization that can be used to produce a multi-tenanted service carry very different risk profiles. Although in both cases many of the risks will be due to two factors:

1. The complexity of the trusted computing base — that is, the parts of the software stack that must be trusted to maintain separation and information security;
2. The complexities of managing the service and infrastructure — ideally, we will have services that can be designed around the principle of least privilege and ensure there are separations of duty between different administrative roles, along with strong audit. Traditionally, in the enterprise, separations between application, database, and system administrators have been viewed as very important, as have separations between developers and those running production systems. These are seen as risk-reduction measures aimed at limiting what each individual can do and know along with making it hard for a rogue employee to cover his tracks.

If we move to a world in which each application provider must code its application to support multiple customers, this has two effects on the trusted computing base. Firstly, the code complexity will increase as multiple customers need supporting and, secondly, much of this code will be bespoke (this may be improved with supporting libraries and coding patterns). With this style of service provision, we are thus very dependent on the skills of a given service provider and, since code will be proprietary, it will be very hard to validate the trustworthiness of the application. Third party code reviews and code verification techniques may help in producing certifications to enhance trust.

The administrative model with application virtualization is also somewhat unpredictable as it will be a consequence of the software design. A rogue administrator or a hacker who can gain administrative privileges will easily be able to gain access to the details of many customers. This concentration of data may encourage attackers, as the value may make the investment of time and effort in sophisticated attacks worthwhile. This was perhaps demonstrated in attacks on email service providers reported early in 2011, where sophisticated attacks were carried out on e-mail advertising distribution services, allowing spammers to get mailing lists for multiple customers as well as use the services to send out spam [40].

The alternative to creating multi-tenanted applications is to run each customer's instance of an application within a separate container. This has the advantage of keeping the application code simple, so reducing bugs and vulnerabilities, reducing the impact of a breach and, where trust can be gained in containment technologies, enhancing trust in the service. Many of the business applications offered in the cloud will be complex and need to scale, supporting many transactions and requiring each instance of the application to run on multiple servers connected by networks. This makes it hard to use simple sandboxing technologies, such as those available in Java.

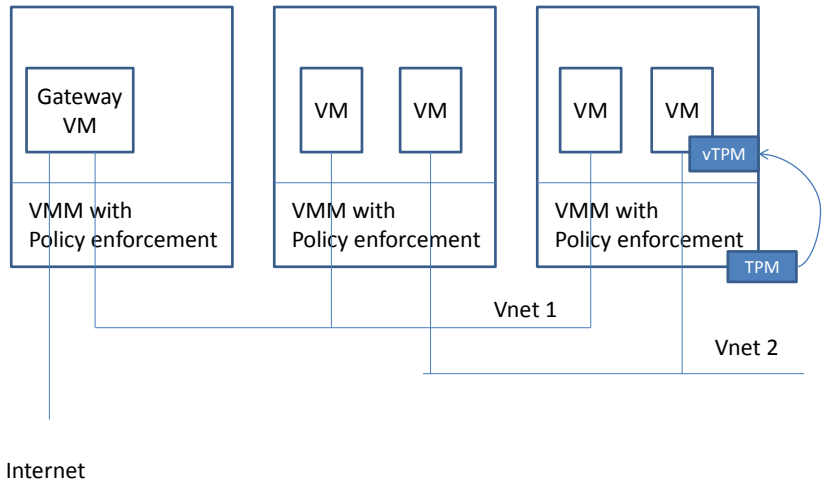

Figure 6.8 Virtual infrastructure can be created with a mixture of virtual machines, virtual networks and virtual storage with attestation provided by TPMs

Virtualization of servers [9], networks [37] and storage can be combined to allow us to build complex boundaries around applications. Each part or service within the application can be run within a different virtual machine, each with its own virtual storage, and the whole lot connected with one or more virtual networks (Figure 6.4), along with a gateway connecting them to the internet to receive/process transactions and management requests. This whole system can effectively replicate what would run within an individual company's data. This architecture can support flexing in that more virtual machine instances can be added into the network to match the required transaction rates. This way of running services could have an additional management burden in setting up the application for each customer and hence relies heavily on having automated configuration systems [29] that support deployment and flexing [27].

Trusted computing technologies, as defined by the Trusted Computing Group [47,54], can help in building trust in to some of our basic infrastructure, and these ideas have been extended to support virtualization. The trusted computing module provides a hardware root-of-trust upon which other trust functions can be built. In particular, three principle functions are provided:

1. A cryptographic identity for a device;
2. Attestation of the software stack that is running and that has been used to boot a system;
3. Safe storage for cryptographic keys that can be linked to attestation.

These three functions allow us to know that we are connecting and relying on the computer systems we expect to, and that they are in the form that we expect. In terms of basic virtualization, we can use TCG-based mechanisms to identify servers and check that the base system software is the virtualization layer. As we move up the stack, we can measure and attest to the various management components, such as the management domain and any separate driver domains. The idea of a virtual TPM [10] has also been introduced so that the integrity of a particular virtual machine can be tested and reported. As we build up the attestation in our virtualized infrastructure, we can also use the TCG mechanisms to help manage cryptographic keys necessary to secure network and storage virtualization [6].

We are concerned that cloud services be deployed in a secure and trustable manner, so having a simple containment strategy is not sufficient unless the container has the correct security and trust properties. Here we look to a number of principles for improving infrastructure that can be built using virtualization and trusted computing.

- *Reduced trusted computing base*. Complex code will always be subject to bugs that lead to vulnerabilities and hence, as a principle, it is important to reduce the code-base that is being trusted. For trusted infrastructure, this means ensuring that the software maintaining separations and managing critical system components and policies are kept to a minimum. For example, if we look at the Xen virtualization layer, not only must we keep the code within the hypervisor to a minimum, but we must also remove all the supporting functions from the management domain (dom0). This means we must work out a minimal set of services that are needed to support virtual machine management and remove the rest of the management stack into other virtual machines that are not part of the trusted computing base. As we construct larger systems, we still need to keep mindful of minimizing the trusted computing base. For example, critical application components can be kept small and run in separate virtual machines running a minimal operating system. Where we need protected storage and communications, the supporting keys can be linked to the attestations of this minimal image. In constructing cloud applications within containers, one particularly sensitive function is the gateways that expose the application to the internet; again, this can be built on a minimal code base.
- *Separate management components*. Defence in depth has long been a principle deployed in enterprise computing, so that a breach on the perimeter does not allow the hacker entry to all of the enterprise systems. The same principle must be applied to the infrastructure supporting cloud. Even as we minimize the trusted computing base, we can keep components separate, so that a break in, say, a network driver does not allow easy access to storage drivers or cryptographic keys.
- *Separation of policy enforcement from application space*. As we build an infrastructure using trusted virtualization, we create containers within which applications can run. We can also control the nature of the container by setting policies within the infrastructure controlling the containment. This means that the containers can be created with properties that are not under the control of those running the application software or anyone who has subverted the application. These policies can also be communicated as part of the TCG attestation measurements, thus giving others confidence that certain policies are enforced. For example, networking policies can be enforced so that messages to or from a container, or systems within a container, can only be sent to certain IP addresses.
- *Separate audit from application space*. As well as policy enforcement being separated from the application environment we can run audit functions outside of the normal application space. Keeping audit out of the application space means that audit records can be protected from tampering.
- *Attestation to communicate trust in the system*. As well as building more secure systems with a minimal trusted computing base, trusted virtualization allows the configuration and code-base to be communicated and attested to though TCG mechanisms. Those relying on systems can therefore gain confidence that the systems are in a trustworthy state.

As we develop infrastructures [16, 26, 37] that support each of these principles, we can build up trustable containment architectures. Considering the cloud ecosystem as described in Section 6.3.4, it seems that the use of trusted virtualization by cloud platform providers will aid each of the service providers in producing more secure services. Trusted virtualization provides a better architectural basis for systems — rather than increasingly more complex security processes and procedures, intended to compete in an arms race — suggesting that the security objectives as expressed by our utility functions may be achievable without big cost increases. From an ecosystem perspective, if such technologies become widely adopted, this should help increase the overall level of both trust and security, so improving the overall reputation of the cloud ecosystem and encouraging the move to cloud.

From an enterprise lifecycle perspective, the use of trusted virtualization has a number of advantages. A company can get information and assurance as to the properties of the infrastructure on which its cloud services are running. This can help both in the initial risk planning to ensure that appropriate levels of security are

achievable and in ensuring that both companies and service providers know that systems are being operated properly.

6.3.6 Assurance

Assurance is about providing confidence to stakeholders that the qualities of service and stewardship with which they are concerned are being managed and maintained appropriately. Cloud computing implies many stakeholders relying on many parties, so that efficient and effective assurance will be both complex and fundamental to a sustainable cloud ecosystem. There is considerable research and discussion on assurance (often driven by regulation) relating to cloud [3,17,18]. In this discussion, we start from some basic principles about assurance in federated environments. We discuss their implications, and the associated opportunities in the context of the stewardship and trusted-infrastructure research described in this chapter.

The principles of assurance are to decide what risks you are concerned about, understand how these risks are (in theory) mitigated, and then to seek evidence that these mitigations are effective. For example, when the Sarbanes-Oxley Act forced companies to demonstrate the integrity of their financial accounts, it was clear there were risks associated with how financial processes and reports depended on IT applications and infrastructure. This led to significant scrutiny of how people are able to change and access the IT infrastructure, which in turn meant many audits on identity management controls.

Analogously, information stewardship implies reliance (and obligations) on many parties to demonstrate how they are controlling risks. Access and identity management will be a part of this, and stakeholders will need different levels and types of assurance associated with all the controls in federated identity management. To expand, multiple parties are involved in registration of people and users, provisioning of credentials, revocation of credentials, creating authorization policy, authentication (of credentials), and enforcement of authorization policy. In turn, each stakeholder will care differently about each of these steps, and so will want different visibility into them. Moreover, we advocate the need for standard publicly reported metrics and data, and the ability for customers occasionally to be able to demand deeper views on specific data relating to their service. The nature of public versus private information for assurance and an expanded discussion of identity assurance are given in [3].

In general, cloud providers will not be able to offer cost-effective services if they must satisfy different audit and assurance requirements for each of their customers. Therefore, from an efficiency perspective, standardized approaches to assurance will be necessary. This is the same argument for standardized sets of terms and conditions — a cloud service provider cannot scale its business if it must accept auditors' checking different aspects from each of its many customers. Moreover, we expect similar analyses of risks and their federated mitigations will be needed to develop these standards.

In addition to assurance from individual providers, each stakeholder will (perhaps implicitly) be concerned about the stewardship of the whole ecosystem. For this reason, we expect there will be a role for metrics that hint at the sustainability and resilience of the whole ecosystem. Initially, they will likely emerge as requirements from clusters of stakeholders — for example, these could be vertical industries such as healthcare and financial services, but also disparate groups with common concerns and views on, say, privacy or law enforcement. It is too early to suggest what these metrics might be, but the conceptual and modelling work described here and in [8,44,45] are about exploring this question.

The work on trusted infrastructure also has direct links to assurance. From an efficiency perspective, it is hugely beneficial if application developers and providers exploit trusted separation in the infrastructure, as opposed to controlling and allowing sharing within the application. The former allows common assurance patterns to be established for when and how to trust infrastructure environments, and should significantly reduce the number of assurance patterns that need to be considered for applications. For example, if cloud platforms routinely run separate service instances in separate and contained trusted infrastructure domains, then service consumers need only seek assurance about a standard set of concerns how the infrastructure controls and maintains its boundaries.

Assurance in most enterprise environments is still a complex mix of automated monitoring and physical audits. Since the amount of assurance activity needs to rise, there will need to be much more reliance on automation. There are many immediate benefits from trusted infrastructure for this, including being able to trust the information, and attestation of components

6.3.7 The Lifecycle Under Cloud

Companies are already struggling with managing IT risks and maintaining an explicit security management lifecycle. Many rely on standards such as ISO27000 [35] as a way of maintaining discipline and, in the background section, we argue that better methodologies are needed to help understand security decisions. Cloud fundamentally changes the way that companies consume IT, as they give up control and rely on others to act as good information stewards. The use of cloud also means that companies have to think of IT in terms of the services rather than the technology components. This has long been the aim of system management through standards such as ITIL, but cloud forces this change in thinking. This means we need to reassess the way in which we think about the security lifecycle.

Each individual service has its own security management lifecycle with the service provider being responsible maintaining its smooth running. It may achieve this by running its own IT systems and operating a traditional security management lifecycle or may itself be a service consumer relying on the security lifecycles that others maintain.

A service consumer will need to maintain an overview on risk and hence needs an aggregated security management lifecycle. This still maintains the risk and governance aspects that are now associated with choosing the right stewardship characteristics for a service. The assurance elements then must be seen in this light: Are the chosen stewardship characteristics being maintained and is risk therefore managed appropriately? Cloud services will not remove all IT from a company and hence the company must still maintain the appropriate policy setting and operational control for the systems that it does run (e.g., clients for end customers or datacentres for the cloud platform providers).

Having a coherent lifecycle for each individual service becomes more important as organizations' ability to react to surprises is reduced. Hence, as we think about risks and how each potential steward will manage them, we must also consider what attestations and assurance metrics are necessary. As services are spread over multiple providers, or as business processes use multiple providers, we need to ensure a consistency between the cycles of each of the constituent services. In the past, this was achieved by having one set of policies, but now we need more careful planning between the various lifecycles; otherwise there will be weak points and potential threats.

We have argued that, for a single enterprise, modelling can help in gaining a better understanding of risk and the trade-offs associated with different policy options. Here, we argue that understanding risk, and the different stewardship options, is even more critical. Consequently, model-based risk methodologies will become increasingly important: system models can be used to explain risks and how they are mitigated.

Consider the VTM (Vulnerability and Threat Management) model discussed in Section 6.2. The model sets out assumptions about the threat environment as well as different controls within the enterprise. Having a model forces us to specify each mechanism coherently. Simply discussing these assumptions and controls can help in gaining common understanding between various stakeholders, but executing the model, to explore a range of system and policy design choices, allows the consequence of these assumptions controls to be explored, and provides evidence as to the likely outcome of employing different mechanisms. In this way, security management moves on from a world in which experts give their opinions to one in which assumptions and abstract mechanisms are clearly specified and their interactions can be explored. Further examinations of the model can help to explore how changes or failures in different processes and technologies can lead to different security failings, so providing a basis for deciding which assurance metrics are important [12].

As the move to the cloud-based services continues, it is not just security teams, or indeed the wider risk committee, who are trying to understand the different policy options and risks. Now both the service provider and the consumer, or those higher in the service supply chain, must adopt a view on risk and be satisfied that appropriate controls are being used. System models can help communicate the risks and mitigations, so allowing customers and service providers to explore different options. In this way, models may become a vital communication point in joining up the risk elements of the lifecycle. Modelling also helps in understanding which assurance metrics are important and hence linking all the pieces of the lifecycle.

No service is an island: in looking at risk, we cannot just look at the performance of the individual pieces on which we rely. We see cloud as an ecosystem in which resource movements or failures in unconnected pieces will affect our IT provision: these factors must be included the lifecycle. As well as selecting and monitoring the cloud services we use, we must consider other factors that may affect their function and, in looking at the security management lifecycle, we need to understand which ecosystem changes may cause changes in the risk profile. These changes may represent changes to exogenous variables within a system model of a particular service. As we reassess the way the lifecycle works, we must also consider the individual pieces of the lifecycle. In particular, we need to ensure that there are appropriate ways of thinking about the risks of handing over data

and what architectural and assurance controls will help mitigate these risks. Pre-cloud, our research agenda needed to be based around improving risk decision-making methodologies, better infrastructure, and assurance. As cloud emerges, these needs do not disappear, but methodologies must take account of this breakup of the lifecycle and be informed by the need for good information stewardship.

6.4 Future Research Directions

The work described here is a rich mix of

- empirical work with stakeholders and professional experts to understand the real dynamics and problems faced by risk and security teams,
- conceptual work with computer scientists, economists, logicians, mathematicians, and psychologists to develop rigorous and clear analyses and approaches, and
- design and engineering by technologists to develop alternative architectures to suit different lifecycles of security management.

Cloud and enterprise computing are continually changing and producing security challenges, and so we see the need to continue and integrate each of these activities. To that end, we will continue to use the partnerships in the TSB-funded Cloud Stewardship Economics and Trust Domains projects to help us to do this.

Thus far the Cloud Stewardship Economics project has involved both empirical and conceptual work. We have done workshops, surveys, and structured interviews with various stakeholders [25], built a series of economic models [32,33,56], and iteratively developed a conceptual framework for analyzing information stewardship [8]. These will all continue, but in addition we will begin to engage stakeholders in using the models to make better decisions. Our vision is to use models to enable structured war-gaming and scenario planning between stakeholders. This involves even tighter integration between the economists, security researchers, cloud/enterprise IT experts, and security professionals.

The Trust Domains [55] project is less mature, but has a larger ambition to integrate and affect architecture and technology. The focus is less governance and policy, and more about achieving operational assurance when multiple stakeholders must share infrastructure. These can be dynamic situations — such as a cross-border civil emergency, where multiple non-trusting groups with infrastructure and applications must suddenly share information and resources — or pre-planned situations, such as long-term (controlled) sharing of resources and information between non-trusting groups. What expectations do such stakeholders have, how explicitly can they describe sharing policy and requirements, and how would they be assured that their information and concerns are suitably managed?

The main focus at the moment is on empirical studies (structured interviews) with a range of potential stakeholders (typically enterprises). From these, we are developing and refining our view on expectations for how information should be managed and how assurance should be provided in shared environments. This in turn will drive both requirements for trusted infrastructure to realize trust domains (containers with relevant properties) and how to use models to more rigorously describe and communicate requirements and real-time operations.

Both these projects directly address the transformed lifecycle of enterprise risk and security management. Clearly, they will not solve all the problems. For example, even if we design strong and appropriate trust domains, and robust information stewardship strategies, there may well be many regulation and commercial drivers that strongly influence risk and security outcomes. Moreover, the nature of technology generally, and cloud specifically, are that they will bring many unforeseen changes. Nevertheless, by working closely with industry and customers, we expect to influence positively the context for cloud associated risk management.

6.5 Conclusions

The development of cloud computing may lead to significant changes in the way companies consume IT, moving from them managing large technology stacks to purchasing business-level services. As this happens, companies will lose control of the way in which their services are run, needing instead to choose between the terms and conditions offered by different service providers. This switch in control may increase the difficulties faced by companies responding to specific security events or failures. This observation emphasizes the need to have a better understanding of risk and how it is mitigated. To achieve this, we need better methodologies, based

on rigorous conceptual and mathematical modelling of systems, of human behavior, and of the wider environment.

As users employ cloud services, they rely on others not only to provide those services, but also to protect their information and appropriately control its interactions and evolution. To understand how this should work, we must widen our view of the declarative goals of information security from confidentiality, integrity, and availability (CIA) to include ideas of duty of service (to ensure that the desired objectives are addressed) and respect for values and ethics. Moreover, the sustainability and resilience of the ecosystem itself must be managed. We describe this broader concept as information stewardship.

This shift of perspective from security to stewardship implies a change in how security (now stewardship) is managed and, we contend, this will be best approach by considering the security management lifecycle, as already operated by many companies. Each IT service will have its own security management lifecycle, possibly dependent on the security management lifecycles of other services and platforms further down the service supply chain. We contend that these interdependent security management lifecycles must be viewed from the perspective of information stewardship. As the flexibility of management is reduced, so we can expect greater coherence between the different elements — such as risk analysis, policy-making, operations, and assurance — of the lifecycle. We must also draw together the various service lifecycles to give consistent pictures of risk, policy-making, operations, and assurance. We see mathematical modelling as playing a huge role in delivering the methodologies that must be developed to achieve all this in the form of practical tools.

Lastly, we consider the stewardship of the ecosystem itself. In the cloud, IT operations will be purchased from highly connected ecosystems of services, consumers and platform providers. Changes in one part of the ecosystem can affect many other parts, in complex ways that will, typically, be difficult to conceptualize. We contend that modelling, of the kind we have sketched in this chapter, can help decision-makers to understand these complex relationships and dependencies. From the perspective of managing the security lifecycle, managers must use this information to understand how different events in different components of the ecosystem may affect the systems for which they are responsible. From the wider perspective of the stewardship of the ecosystem itself, we must ensure that the ecosystem is managed to be sustainable and resilient. These features of the ecosystem are public goods, and we contend that there is a clear role for regulators in their stewardship.

Acknowledgments

This chapter draws on the work of and conversations with all of the security research team in HP Labs. Specifically, we thank Boris Balacheff and Chris Dalton for their advice all areas relating to trusted infrastructure; Yolanta Beres and Jonathan Griffin for their work on process modelling of vulnerability management; Chew Yean Yam and Christos Ioannidis (University of Bath) for work on the switching (real options) model; Matthew Collinson (University of Aberdeen) and Brian Monahan for work on foundations and process modelling across all the projects; Marco Casassa Mont for work on identity assurance; and Martin Sadler for overall vision. We would also like to thank and acknowledge all our partners in the Cloud Stewardship Economics and Trust Domains projects, and the UK Technology Strategy Board for its funding of these projects.

References

[1] Anderson R (2001) Why information security is hard: An economic perspective. In: Proc 17th Annual Computer Security Applications Conference (ACSAC) 358-365 IEEE Computer Society Press

[2] Armour FJ Kaisler SH Liu SI (1999) Building an Enterprise Architecture Step by Step. IT Professional 1(4): 31-39. doi:10.1109/6294.781623

[3] Baldwin A Beres Y Shiu S (2006) Using assurance models in IT audit engagements, HP Labs Technical Report HPL-2006-148

[4] Baldwin A Beres Y and Shiu, S (2007) Using assurance models to aid the risk and governance life cycle. *BT Technology Journal* 25:128-140. doi:10.1007/s10550-007-0015-7

[5] Baldwin A Casassa Mont M Beres Y Shiu S (2010) Assurance for federated identity management. *Journal of Computer Security* 18(4): 541-572

[6] Baldwin A Dalton CI Shiu S Kostienk K Rajpoot Q (2009) Providing secure services for a virtual infrastructure. *SIGOPS Operating Systems Review* 43(1):44-51. doi:10.1145/1496909.1496919

[7] Baldwin A Casassa Mont M Shiu S (2009) Using Modelling and Simulation for Policy Decision Support in Identity Management. POLICY 2009: 17-24
[8] Baldwin A Pym D Sadler M Shiu S (2011) Information stewardship in cloud ecosystems: towards models, economics and delivery. To appear: Proc of the 3rd IEEE International conference on Cloud Computing, Athens, 2011
[9] Barham P Dragovic B Fraser K Hand S Harris T Ho A Neugebauer R Pratt I Warfield A (2003) Xen and the art of virtualization. Proc Nineteenth ACM Symposium on Operating Systems Principles, October 19-22, 2003, Bolton Landing, NY, USA. doi:10.1145/945445.945462
[10] Berger S Cáceres R Goldman KA Perez R Sailer R van Doorn L (2006) vTPM: virtualizing the trusted platform module, Proc of the 15th conference on USENIX Security Symposium, July 31-August 04, 2006, Vancouver, BC, Canada
[11] Beres Y Griffin J Shiu S Heitman M, Markle D, Ventura P (2008) Analysing the performance of security solutions to reduce vulnerability exposure windows. Proc. Annual Computer Security Applications Conference (ACSAC), 33-42, CA IEEE 2008
[12] Beres Y Casassa Mont M Griffin J Shiu S (2009) Using security metrics coupled with predictive modeling and simulation to assess security processes. ESEM 2009: 564-573
[13] Beres Y Pym D Shiu S (2010) Decision Support for Systems Security Investment. Proc. Business-driven IT Management (BDIM) 2010. IEEE Xplore, 2010
[14] Beautement A Coles R Griffin J Ioannidis C Monahan B Pym D Sasse A Wonham M (2009) Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security.In *Managing Information Risk and the Economics of Security*, M Eric Johnson (editor), Springer 2009
[15] BeautementA and Pym D (2010) Structured Systems Economics for Security Management.. Proc. WEIS 2010, Harvard University. http://weis2010.econinfosec.org/papers/session6/weis2010_beautement.pdf
[16] Cabuk S, Dalton CI Eriksson K Kuhlmann D Ramasamy HV Ramunno G Sadeghi A Schunter M Stüble C (2010) Towards automated security policy enforcement in multi-tenant virtual data centers. *Journal of Computer Security* 18(1): 89-121 (2010)
[17] CAMM (Common Assurance Maturity Model Guiding Principles) (2010) http://common-assurance.com/resources/Common-Assurance-Maturity-Model-vision.pdf
[18] Catteddu D and Hogben G (2009) Cloud computing information assurance framework, ENISA Report (2009). http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-information-assurance-framework/ (accessed 01/01/2012)
[19] Chapin III FS Kofinas GP Folke C (editors) (2009) Principles of Ecosystem Stewardship: Resilience-Based Natural Resource Management in a Changing World. Springer-Verlag, 2009
[20] Chen Y Bharadwaj A (2009) An Empirical Analysis of Contract Structures in IT Outsourcing. *Information Systems Research* 20:484–506, 2009
[21] Collinson M Monahan B Pym D (2009) A logical and computational theory of located resource. *Journal of Logic and Computation* 19(b):1207–1244, 2009. doi:10.1093/logcom/exp021
[22] Collinson M Monahan B Pym D (2010) Semantics for Structured Systems Modelling and Simulation. Proc. Simutools 2010, ACM Digital Library and EU Digital Library. ISBN: 978-963-9799-87-5
[23] Collinson M, Monahan B Pym D (2012) A Discipline of Mathematical Systems Modelling. Forthcoming monograph, College Publications, 2012
[24] Core Gnosis (2012) http://www.hpl.hp.com/research/systems_security/gnosis.html (accessed 01/01/2012)
[25] Cloud Stewardship Economics (2012) http://www.hpl.hp.com/bristol/cloud_stewardship.htm
[26] Dalton C Plaquin D Weidner W Kuhlmann D,Balacheff B, Brown R (2009) Trusted virtual platforms: a key enabler for converged client devices. *SIGOPS Oper. Syst. Rev.* 43, 1 (January 2009), 36-43. doi:10.1145/1496909.1496918
[27] Degabriele JP Pym D (2007) Economic aspects or a utility computing service HP Labs technical report, HPL-2007-101, 2007
[28] Eskins D Sanders WH (2011) The multiple-asymmetric-utility system model: A framework for modeling cyber-human systems. Proc 8th International Conference on Quantitative Evaluation of Systems (QEST) 2011
[29] Goldsack P Guijarro J Loughran S Coles A Farrell A Lain A Murray P Toft P (2009) The SmartFrog configuration management framework. *SIGOPS Oper. Syst. Rev.* 43, 1 (January 2009), 16-25. doi:10.1145/1496909.1496915
[30] Gordon LA and Loeb MP (2006) Managing Cybersecurity Resources: A Cost-Benefit Analysis. McGra
[31] Ioannidis C Pym D Williams J Investments and Trade-offs in the Economics of Information Security in Proc. Financial Cryptography and Data Security 2009, LNCS 5628: 148-162, Springer, 2009
[32] Ioannidis C Pym D Williams J (2009) Investments and Trade-offs in the Economics of Information Security. In Proc. Financial Cryptography and Data Security 2009, LNCS 5628:148-162, Springer, 2009
[33] Ioannidis C Pym D Williams J (2011) Fixed Costs, Investment Rigidities, and Risk Aversion in

Information Security: A Utility-theoretic Approach. In: Schneier B (ed) Proc. Workshop on Economics of Information Security (WEIS) 2011, Springer. In press
[34] Ioannidis C Pym D Williams J (2012) Information Security Trade-offs and Optimal Patching Policies. *European Journal of Operational Research*, 216(2):434-444. doi:10.1016/j.ejor.2011.05.050
[35] ISO (2007) ISO 27000 Series of Standards (Supersedes ISO 17799), 2007. http://www.27000.org (accessed 01/01/2012)
[36] ITGI (2005) *Control Objectives for Information and Related Technologies (COBIT)*, 4th edition, 2005.
[37] Kallahalla M, Uysal, M, Swaminathan R, Lowell DE, Wray M, Christian T, Edwards N, Dalton CI, Gittler F, (2004) SoftUDC: A Software-Based Data Center for Utility Computing, Computer, v.37 n.11, p.38-46, November 2004. doi:10.1109/MC.2004.221
[38] Khwaja T (2002) Should I Stay or Should I Go? Migration Under Uncertainty: A Real Option Approach, Public Policy Discussion Papers 002-10, Economics and Finance Section, School of Social Sciences, Brunel University, 2002
[39] Keeney RL Raiffa H (1976) Decisions with Multiple Objectives: Preferences and Value Tradeoffs. Wiley, New York. Reprinted, Cambridge University Press, New York (1993)
[40] Krebs B (2011) Epsilon Breach Raises Specter of Spear Phishing http://krebsonsecurity.com/2011/04/epsilon-breach-raises-specter-of-spear-phishing/ (accessed 01/01/2012)
[41] Lloyd V (2011) Planning to implement service management (IT Infrastructure Library). The Stationery Office Books. http://www.itil.co.uk/ publications.htm (accessed 01/01/2012)
[42] Mell P Grance T (2011) The NIST Definition of Cloud Computing (Draft). Technical report, National Institute of Standards and Technology, US Department of Commerce, 2011. Special Publication 800-145 (Draft)
[43] Open Trusted Computing (2012) http://www.opentc.net/ (accessed 01/01/2012)
[44] Pym D and Sadler M (2010) Information Stewardship in Cloud Computing. *International Journal of Service Service, Management, Engineering, and Technology*, 1(1):50–67, 2010
[45] Pym D Sadler M Shiu S Casassa Mont M (2010) Information Stewardship in the Cloud: A Model-based Approach. In Proc CloudComp 2010, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering (LNICST). Springer. To appear
[46] Pym D Shiu S Coles R van Moorsel A Sasse MA Johnson H (2011) Trust Economics: A systematic approach to information security decision making. Final Report for the UK Technology Strategy Board 'Trust Economics' project. http://www.hpl.hp.com/news/2011/oct-dec/Final_Report_collated.pdf (accessed 01/01/2012)
[47] Pearson S Balacheff Chen L Plaquin D Proudler G (2003) Trusted computing platforms: TCPA in context. HP Books, Prentice Hall, 2003
[48] Squicciarini AC Rajasekaran SD Casassa Mont M Using Modeling and Simulation to Evaluate Enterprises' Risk Exposure to Social Networks. IEEE Computer 44(1): 66-73 (2011)
[49] Acquisti A Anderson R Schneier B (2011) 4th security and human behavior workshop, 2011, Carnegie Mellon University. http://www.heinz.cmu.edu/~acquisti/SHB/ (accessed 01/01/2012)
[50] Shiu S Baldwin A Beres Y Casassa Mont M Duggan G Johnson H Middup C (2011) Economic methods and decision making by security professionals. In Schneier B (editor) (2012) Proc. Workshop on Economics of Information Security (WEIS) 2011. Springer. In Press http://weis2011.econinfosec.org/papers/Economic%20methods%20and%20decision%20making%20by%20security%20profession.pdf (accessed 01/01/2012)
[51] Spewak SH Hill SC (1993) Enterprise architecture planning: developing a blueprint for data, applications and technology. QED Information Sciences, Inc., Wellesley, MA, 1993
[52] Stoneburner G Goguen A Feringa A (2002) Risk Management Guide for Information Technology Systems Technical Report, National Institute of Standards and Technology, U.S. Department of Commerce, NIST Special Publication 800-30. 2002 http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf
[53] Trigeorgis L (2001) Real Options: An Overview. In E. S. Schwartz and L. Trigeorgis. Real Options and Investment under Uncertainty: Classical Readings and Recent Contribution. MIT Press, 2001
[54] The Trusted Computing Group. http://www.trustedcomputinggroup.org/ (accessed 01/01/2012)[55] US Congress. S. 3742: Data Security and Breach Notification Act of 2010. http://www.govtrack.us/congress/bill.xpd?bill=s111-3742 (accessed 01/01/2012)
[56] Yam C-Y Baldwin A Ioannidis C Shiu S (2011) Migration to Cloud as Real Option: Investment ecision under uncertainty. In Proc IEEE TrustCom 2011 Symposiums & Workshops. In Press