Information Stewardship in Cloud Ecosystems: Towards Models, Economics, and Delivery

Adrian Baldwin
HP Labs
Bristol BS34 8QZ
England, UK
Email: adrian.baldwin@hp.com

David Pym*
University of Aberdeen
King's College, Aberdeen AB24 3UE
Scotland, UK
Email: d.j.pym@abdn.ac.uk

Martin Sadler and Simon Shiu
HP Labs
Bristol BS34 8QZ
England, UK
Email: first.last@hp.com

***Abstract*—We discuss the concept of information stewardship in cloud-based business ecosystems. The constituent concepts of stewardship — which we believe will be crucial to the successful development of cloud-based business of all kinds — extend those of security to encompass concepts of objectives, ethics/values, sustainability, and resilience: all familiar from the stewardship of natural resources. Our view is based on rigorous approaches from mathematical systems modelling and economics, and is informed by concepts from natural resource management and information assurance.**

I. Introduction

The current state of cloud computing is, essentially, captured in the NIST definitions [20]. The core characteristics of cloud services are described, together with the different service and deployment models. NIST defines infrastructure-as-a-service (IaaS) as the provision of the basic services of storage, compute and networking, moving up the stack to platform-as-a-service (PaaS), in which middleware is added to basic compute services. In the discussions within this paper, we group these two elements together describing them as a cloud platform. NIST also considers software-as-a-service (SaaS); that is, software applications running on cloud infrastructure. This is the level at which we see the business-level services emerging. Accordingly, for a clean abstraction and presentation, we adopt a three-layer model of the cloud (see Figure 1) in which we have companies that consume cloud services, software companies that provide their software as services on the cloud, and cloud platform providers who provide and manage the basic infrastructure.

The essential characteristics for cloud also help define a view on the cloud. NIST considers about the need for services to be network accessible.[1] From a customer's perspective, NIST considers the need for on-demand, self-service provisioning and rapid elasticity, hence allowing consuming companies to acquire standardized services easily (and to scale them to their needs). They also consider resource pooling, sharing, and multi-tenanted systems — properties required by service providers to control costs and meet elasticity requirements. Metering, or measurement at least, is required for billing, to allow service providers to optimize their operations and services.

These essential characteristics can be contrasted with the more traditional IT services model of outsourcing, in which, typically, the customer will negotiate a contract with a service provider to run a part or all of its IT systems. This negotiation typically leads to a unique contract for each customer [4], with a transition phase, and sets of operational policies, processes, metrics, and service-level agreements (SLAs) agreed with the customer. This should be contrasted with the cloud where self-service provisioning means that the service provider will, typically, offer a fixed set (or menu) of terms and conditions (implying policies and operating procedures) along with (shared) metering and measurements.

Much of the current usage of cloud is associated with organizations' renting CPU cycles on platforms such as Amazon's EC2. We believe, however, that the cloud will mature into providing compositions of business services, based on shared infrastructure. For an individual company, we can illustrate this transition happening with the company starting by purchasing less critical services in the cloud and eventually moving to placing all its business-critical IT systems into the cloud. In our three-layer ecosystem (see Figure 1), each customer typically has a number of applications supporting its business and they will decide whether to run the application within their own data centre or use a cloud service via a service provider (which may be the same organization as the infrastructure provider).

The SaaS (or 'app') providers are likely to start off as software vendors (cf. Apple's App Store) with each needing to make a decision: do we sell shrink-wrapped software, do we run the software on a cloud platform, or both? The platform providers must make massive investments in data centres and infrastructure, and hence are in practice limited to a few large companies. As the cloud ecosystem grows, the platforms will provide an easy delivery channel for SaaS providers who can offer services with minimal start-up costs. This will likely stimulate innovation, and customers may see new services emerge beyond the traditional IT business applications. This

This work has been partially supported by the UK's Technology Strategy Board via the project 'Cloud Stewardship Economics: Securing the New Business Infrastructure'.

*Corresponding author.

[1] We see the Lloyd's insurance market as a manifestation of a cloud ecosystem in which insurance brokers and underwriters are like the software providers, running business services that companies consume, with the market place being a platform providing basic services and hosting the insurance service providers.

might be expected to encourage further migration to the cloud. As businesses adopt cloud services, regulators, such as the UK's information commissioner, will need to respond to ensure good stewardship for society.

In this paper, building on [27], [28], we consider the concept of information stewardship in cloud-based business ecosystems. The declarative and operational concerns of information stewardship include those of information security — confidentiality, integrity, availability (CIA) and the operational mechanisms used to achieve them — and those of privacy. Critically, and characteristically, information stewardship is concerned also with concepts such as the management and supervision of values, respect for ethics, duty of service, responsibility, and, in the context of stewardship of the ecosystem itself, the promotion of resilience and sustainability [23], [13], [24]. We find that certain concepts from the field of natural resource stewardship resonate strongly with our concerns.

The purpose of this paper is to introduce and discuss some concepts of, issues in, and possible approaches to understanding the structure, dynamics, and use of cloud-based business ecosystems. We do not claim to provide definitive solutions.

We begin, in Section II, with an introductory discussion of the concept of a cloud-based business ecosystem (which we abbreviate to cloud ecosystem and ecosystem where it is convenient to do so). In Section III we give a brief sketch of the mathematical and computational framework that provides suitable tools for modelling cloud-based business ecosystems. Whilst we do not present models in this paper, the modelling concepts we introduce permeate our discussion throughout. In Section IV we review the concept of information stewardship and then, in Section V, we sketch our economic approach to security, stakeholders, and stewardship. In Section VI, we discuss the dynamics of cloud-based business ecosystems, considering the implications of evolving business and threat environments for the key stewardship concepts resilience and sustainability. In Section VII, we consider briefly some societal implications deriving from the emergence of cloud-based business ecosystems.

Finally, a note on references: this short article is a wide-ranging discussion drawing upon ideas from several disciplines: a comprehensive treatment of the relevant literature is not practicable. Accordingly, we give just illustrative references, deferring a fuller treatment to another occasion.

II. Cloud Ecosystems

We describe the emergent cloud system as an ecosystem — in the sense, for example of [23] — and within this paper we draw on this analogy looking at how ecologists have studied ecosystems. From an ecological perspective, an ecosystem is the complex set of relationships between organisms and the habitats in which they live. They are studied as a whole as the complex relationships, with the consequence that changes to one part of the ecosystem may impact upon other parts.

Rather than organisms and habitats, the cloud ecosystem consists of a number of companies that are service consumers, service providers, and platform providers, as illustrated in Figure 1. Instead of an organism's biology determining how it acts within the ecosystem, each company will have sets of policies and processes as well as the people within the company that control its actions. Each company will also have its own set of incentives, with associated utilities, that influence the design of these policies and processes. In studying the cloud ecosystem, we are primarily concerned with the way each of these companies provisions and manages its IT requirements and, as such, we can view business needs as parts of the habitat or environment in which they needs are developed and influenced. In the same way that weather cycles affect habitats, organisms and the overall ecosystem, the economic cycle will be one of the drivers for cyclical behaviour within the cloud.

Each individual entity within the ecosystem will make decisions based on its perception of the state of the ecosystem. As well as looking at these entities we also need to consider the communication channels between the entities, how information spreads and hence affects decision making. This could be seen as an asymmetric information problem between the various entities; however, in looking at who knows what, we should also understand the trust or weight that entities place in information from different sources. The information they are using may not be entirely accurate, but this influences their reactions to the ecosystem within which they exist.

Ecosystems can be studied at a variety of scales and levels of abstraction depending on the problem being addressed, with those elements that may affect the system but which fall outside the ecosystem (at the level at which it is being studied) being considered exogenous variables.

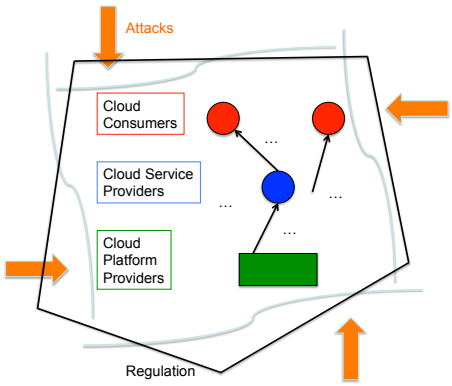

Fig. 1. A simple three-layer cloud ecosystem situated in its regulatory and threat environments

Throughout the remainder of this paper we will use a simple running example, described informally, but which could be formalized — following the modelling approach employed in, for example [29] — in order to build mathematical and economic models of the kind we describe below.

We consider a firm that is large enough to require a range of substantial operational departments, including finance, human resources, and IT, as well as customer-facing departments. We will illustrate our discussion using the firm's processes (i.e.,

both the business processes that deliver the firm's products and services to its customers and its internal management processes), its resources (i.e., its stock, funds, staff, computer systems components, information, etc.), and its structure (i.e., its physical and logical distribution and organization). We will also consider the rôle of the environment within which the firm exists, including the both the business environment and the security and stewardship threat environment.

III. Modelling Cloud Ecosystems

The first key notion in modelling cloud ecosystems is that of systems model. It is useful to argue — see, for example, [8] — that the key structural aspects of systems are the ones discussed below, a view that is consistent with the classical view of distributed systems, as described, for example, in [10]. There is also a close correspondence with the organization of natural resource ecosystems, as described in [13] and their analysis as complex systems (see several articles in [24], [13]).

Process. A synthetic system exists in order to deliver services, and services can be conveniently understood as processes that execute on the systems architecture. Typically, it will be necessary for many services to execute, concurrently and sequentially, in order for a service to be delivered. Similarly, natural processes execute relative to natural substrates. Our main focus here is on synthetic systems; in particular, large-scale information systems.

For example, suppose that our firm is a service provider (i.e., in the middle layer of Figure 1). There will be business processes that interact with customers, internal processes, and processes that interact with those of the infrastructure providers. These processes all execute concurrently and share and communicate resources. Typically, stakeholders in the ecosystem are, where necessary, represented as processes.

Mathematically, we use a model of processes that is based on Milner's synchronous calculus of communicating systems, known as SCCS [21], [22], developed to allow the co-evolution with processes of resources and locations, each of which is consider below. The process language we use admits, in addition to basic action sequencing, nondeterministic choice, concurrent composition, and a notion of resource-hiding, as well as recursion. The details are described in [7].

Resource. The infrastructure of a system, relative to which the systems processes execute, consists of a collection of resources that may be utilized by the processes in order to achieve their intended purposes.

In the example of our firm, the resources that may need to be modelled include the firm's stock, its available funds, its staff, its computer systems components, and the information of which it is steward on behalf its customers.

Mathematically, we model resources using ordered monoids, motivated by a conceptual analysis of the notion of resource that suggests that the key features of resource elements are that they can be *combined* (monoidal composition) and *compared* (ordering). This structure has proved a remarkably useful abstraction in practice [25], [7], [30] and, mathematically, fits well with the model of process mentioned above [7].

Location. In general, the architectures of systems are highly distributed, logically and/or physically. The systems resources are distributed around a collection of places, and these places have (directed) connections between them.

For example, our firm may have offices in may different cities, with physical (e.g., staff travel in cars to move between offices) and network (e.g., dedicated lines, public internet) connections between them. Within each office, the topology of the organization, and its associated distribution of resources, may also be significant for the firm's business processes. Location is also a logical notion, capturing, for example, firms, or even whole industries, as points in the ecosystem.

Mathematically, our treatment of location starts with the following basic requirements of a useful notion of location [5], [8], [6]: a collection of atomic locations — the basic places — which generate a structure of locations; a notion of (directed) connection between locations — describing the topology of the system; a notion of sublocation (which respects connections); a notion of substitution (of a location for a sublocation) that respects connections — substitution provides a basis for abstraction and refinement in our system models; product (again, monoidal) of locations (an inessential but useful technical property), suitably coherent with the other products [5]. Various classes of graphical and topological structures provide leading examples [5], [8], [6].

Environment. Systems exist within external environments, from which events are incident upon the systems boundaries. Typically, the environment is insufficiently understood and too complex to be represented in the same, explicit, form as the system itself. Instead, events that are incident upon a system's boundary are represented mathematically using stochastic methods; see [17], [8], [6].

For example, our firm may receive orders for its various goods and services at different rates for each different offering. Each of these might be represented using a negative exponential distribution. Similarly, the firm's IT systems may also be subject to attacks from within and outwith the ecosystem, with different types of attack occurring at different rates.

Combining our treatments of process, resource, and location, we obtain a calculus that describes how the state of asystem, expressed as a triple L, R, E of location, resource and process evolves to a new state when an action

$$L, R, E \xrightarrow{a} L', R', E'$$

These basic evolutions combine to describe, in terms of a structural operational semantics, the dynamics of the complex processes mentioned above.

Along with this process-theoretic set-up, we obtain a modal logic — combining ideas from Hennessy–Milner logic [22] and bunched logic [25] — that describes properties of the system states. This yields a logical judgement of the form

$$L, R, E \models \phi$$

which is read as 'relative to the available resources R at location L, the process E has property ϕ'. The modal logic of propositions that describes the properties is rather rich,

and includes a collection of 'separating' connectives of the kind found in Separation Logic [30]. These connectives admit compositional descriptions of system properties in terms of constituent subsystems.

A few points about the style of modelling employed using these mathematical tools are noteworthy. The mathematical tools used do not impose any choice of level of detail or level of abstraction. That choice remains with the modeller. Whilst the tools can, in principle, be used to describe systems at a very fine-grained level of detail, they have proved very useful — in commercial, industrial strength applications (see [29]) — for describing large-scale systems at relatively high levels of abstraction. The key to this lies in the modelling language Core Gnosis [9], [6], which implements the mathematical constructs mentioned above, including the stochastic representation of events incident upon the system, in a programming-language style, together with a well-developed modelling idiom [3], [8] that admits the representation of attributes of model components within the given semantic framework. Core Gnosis models are executable and provide simulations as a basis for Monte Carlo-style experiments.

But reasoning about system design requires not only models of the systems themselves, but also their designers' stewardship preferences, considered in their economic context. We now explore how these ideas and how they fit together with system models.

IV. Information Stewardship in Cloud Ecosystems

Information security is concerned with the confidentiality, integrity, and availability (CIA) of information — represented as stored data — in information processing systems, the objective information security operations being to protect these properties. Such protection is costly and is not absolute. Accordingly, the managers of information systems must determine not only their target levels of confidentiality, integrity, and availability, but also their target levels of investment or cost. In the analysis of information security architectures, with these concerns in mind, it has proved helpful to distinguish declarative and operational concepts [1]. This distinction sheds light on the inadequacy of many so-called refinements of the declarative concepts of confidentiality, integrity, and availability. Typically, such refinements confuse declarative and operational concepts and introduce concepts such as authentication, audit, non-repudiation, and even utility (see [26] for an extensive discussion along these lines). These are category errors: authentication (for example) should be seen as an operational mechanism by which aspects of the declarative objectives of confidentiality and availability can be delivered.

It is important to note, in the context of the cloud ecosystems in particular, that we are concerned not only with maintaining these declarative properties of static data, but also with protecting these properties of data during the execution of transactions. These observations lead us to consider what we might mean by information stewardship. In particular, we consider what it is for one agent (the steward) operating in a cloud ecosystem, to be the steward of the information belonging to another agent (the client) in order to obtain the provision of services to the client Information stewardship certainly includes the management of the core concepts of information security — confidentiality, integrity, and availability — as well as the management of privacy. But stewardship is also concerned with the management of the client's values and reputation as the client's information is manipulated by the ecosystem. We shall return to this concept after we have established some basic ideas about utility theory and its application to information ecosystems.

V. Utility and Modelling Security

Utility theory, a cornerstone of economics, provides a conceptual and mathematical set-up for modelling how declarative security properties, such as confidentiality, integrity, and availability trade off against one another and against cost [14], [16], [15]. It also allows us to understand how the magnitudes of these properties may deviate from their targets as a system interacts with its environment and evolves. Expressions of utility are thus, with respect to (declarative) security objectives representations of the system manager's policy.

Establishing such an understanding involves developing not only a utility-theoretic representation of a system's managers' preferences, but also a representation of the system itself and its evolution, together with a representation of their relationship. Roughly, the managers are concerned with the (expected) utility of a basket of quantities of interest — such as confidentiality, integrity, availability, and cost — where utility is expressed by, for example, a function of the form

$$U \quad \triangleq \quad \sum_{i=1}^{k} w_i f_i(Q_i - \bar{Q}_i) \tag{1}$$

in which we have the following:

- Each Q_i is one of the quantities of interest, such as (some or all of) CIA or proxies for these, or cost;
- Each $\bar{Q}_i$ is a target value for Q_i;
- Each f_i is a (possibly stochastic) function that expresses the shape of the managers' preference for the behaviour of each quantity relative to its target. A simple version of this set-up would take the f_is to be quadratic. Quadratics conveniently express diminishing marginal returns as the indicators approach target, but make utility symmetric around target. More realistically, Linex functions (e.g., [33], [31]), usually expressed in the form $g(z) = (exp(\alpha z) - \alpha z - 1)/\alpha^2$ are used to capture a degree of asymmetry that is parametrized by α;
- Where stochastic components are present, we are typically interested in expected utility, $E[U]$; and
- Each w_i is weighting, expressing the relative significance of each quantity within the managers' policy.

The variation of the Q_is over time is determined by the behaviour of the system as it evolves, as described by a model of the system's structure and dynamics. For example, we might — in the style of economic modelling — focus more-or-less

wholly on the dynamic behaviour, and represent the ecosystem by a system of (stochastic, simultaneous) equations[2]

$$Q_i \triangleq s_i(Q_1, \ldots, Q_k; c_i) \quad (2)$$

where each s_i is a function, possibly involving stochastic processes, that describes the behaviour of the ith quantity of interest, and each c_i is a control variable for the ith quantity.

Equation 2 provides the link between the dynamics of system models (as described in Section III and the dynamics of the system's utility [1], [6], [14], [16]).

A very familiar example (from macroeconomics) of such a set-up might be provided by the management of inflation and unemployment by a central bank [31]. A bank might be set targets for these quantities, which trade off against each other, by its government. The bank's control variable that is the interest rate, and the bank's task is to set a monthly sequence of rates so that inflation and unemployment stay on target.

In a highly complex situation, such as in a security architecture, it will typically not be possible to formulate system equations (in terms of functions s_i) in the way that is usually possible in, for example, macroeconomic modelling. Typically, though, the key control variables, such as system interconnectivity or investment in various aspects (people, process, and technology) of security operations, will be identifiable. Instead, however, an executable system model, using the key control variables, can used in order to simulate the dynamics of the utility function. Depending upon the requirements of the analysis, one might use a modelling language that incorporates a sophisticated mathematical analysis of system structure, such as described in [8], or employ a tool such as system dynamics [32], which emphasises more directly influence and feedback with perhaps less control of the detailed behaviour.

A key question here, explored in detail in [14], concerns the resilience of quantities of interest (with respect to maintaining target levels) when the system experiences shocks, such as a breach of confidentiality caused by a social engineering attack or the cracking of an encryption code, or the loss of a web-based service caused by a distributed denial of service attack.

A. From Security to Stewardship

This view of the economics of information security has proved to be valuable in advancing out conceptual understanding of the decision processes around the protection of information in situations in which the owner or manager of the information maintains an intimate relationship with the service-provider and the information being processed by the provider. As service-providers move to cloud ecosystems — that is, complex networks of interacting infrastructure providers, service providers and consumers — and as service-provision becomes more devolved and distributed within cloud ecosystems, this intimate relationship will be considerably weakened. Indeed, provided the information-owner's interests are properly protected, the opportunities provided by the cloud ecosystem may well be highly advantageous to the information-owner. To understand what should be meant by protecting the interests of the information-owner, we are led to the concept of information stewardship in cloud ecosystems. In this context, information stewardship would certainly encompass the security concerns that we have discussed, but would much more besides.

Informally, the notion of stewardship is understood to capture, in addition to the core concepts of information security, concepts such as the management and supervision of objectives, respect for ethics and values (e.g., duty of service, responsibility), and, in the context of stewardship of the ecosystem itself, the promotion of sustainability and resilience. The concepts and approaches that we have described above set out a collection of tools from economic and mathematical modelling that are of great utility in understanding the concept of information stewardship, it is useful to consider some notions of stewardship that have been found to be useful in other intellectual disciplines. One view (see, for example, various dictionaries) is that which has been developed in areas such political science, where the term is used to capture concepts such as the management and supervision of resources, adherence to principles, and the trusted prosecution of obligations. A steward, in this context, is one who is employed to carry out these functions on behalf of another or others. All of these notions might, at least in principle, be incorporated into the utility-theoretic and system modelling framework sketched above. However, we suggest that an alternative, and useful, point of departure is (as suggested in Section I) provided by the work of ecologists in understanding natural resource stewardship in natural social-ecological systems (for a comprehensive and thoughtful overview, see [13]). Here the key notion is that of a stakeholder in the ecosystem.

B. Stakeholders and Utility

A cloud-based business ecosystem includes many stakeholders. Each stakeholder has a perspective on the structure and function of the ecosystem which, together with its objectives, determines the formulation of the stakeholder's utility function for its engagement in the operation of the ecosystem.

As we have seen, examples of stakeholders include individual participants (e.g., consumers, service providers, platform providers), policy-makers (within and outwith the ecosystem) and regulators (e.g., politicians, government agencies), providers of professional services in support of ecosystem operations (e.g., auditors, lawyers), and equipment manufacturers (e.g., computer and network infrastructure manufacturers).

Each participant seeks to optimize, or at least satisfice, its own utility, according to the utility function that is appropriate for its perspective. The formulation of a stakeholder's utility function will, as usual, depend on a range of factors with, as has been argued in [14], [16], the perspective and techniques provided by macroeconomic and financial management being useful. The stakeholder will identify, as outlined above, a collection (sometimes called a basket) of quantities (the Q_is, as in Equation 1) that are of concern, together with target

[2]The system equations typically will include stochastic processes that represent the variability of the environment within which the system operates.

values and weightings, and will identify a functional form to describe the desired utility as the value of each quantity deviates from target.

The challenge for the regulators is to identify a utility function that adequately reflects the objectives of the policy-makers. The policy-makers determine what is socially optimal, and the regulators must seek to deliver appropriate behaviour by the ecosystem by formulating an appropriate utility function for the overall system. We conjecture that such a function will, at least in principle, be constructed from the utility functions of the different stakeholders; that is,

$$Overall\ Ecosystem\ Utility \triangleq R(U_1, \ldots, U_n) \qquad (3)$$

where R is the regulators' choice of (possibly stochastic) functional combination of the stakeholders' utilities ($U_1, \ldots U_n$). The regulators' task is complicated by the possibility of only partial knowledge of all of the U_is in Equation 3 — some stakeholders may be secretive. Several concepts from the treatment of commons within economic theory are relevant. For example, some aspects of the system will be like public goods (e.g., those parts of the shared infrastructure that are concerned with the ecosystem's information security architecture) and some will be more like club goods (e.g., payment systems).

Overall, the stakeholders in the ecosystem are faced with the need to make multiple, multi-objective decisions about highly complex systems ([19], [2] are excellent starting points among many for the relevant theory). For policy-makers, and hence for regulators, it is likely that important objectives will be appropriate levels of resilience of the ecosystem as it is subject to shocks — such as changes in the economic conditions within which the ecosystem operates and security attacks against the technology or business processes — and the sustainability of the ecosystem over its lifetime of operations. Note that we are concerned here with the sustainability and resilience of the concepts of stewardship (including sustainability and resilience themselves) and not merely of CIA.

VI. Ecosystem Dynamics

In the world of natural resource management (e.g., [13], for a range of pertinent articles and a wealth of references), resilience and sustainability are perhaps the key drivers for the ecosystem's stewards. For Chapin, Kofinas, and Folke (in [13]), the key concept is that of *resilience-based ecosystem stewardship*, which 'involves responding to and shaping change in social-ecological systems to sustain the supply and opportunities for use of ecosystem services by society'.

In this section, we discuss some of the key factors in the dynamics of cloud-based business ecosystems.

A. Influences Between Ecosystem Participants

We can describe cloud-based business ecosystems in terms of the various participants (companies, etc.) who are either service consumers, service providers, or cloud platform providers. Each will interact with the others within the ecosystem, as well reacting to exogenous controls. Participants will use their knowledge of the state of the ecosystem to assess how they should interact with it to maintain their desired utility. For example, a participant such as our firm may be considering moving its IT to the cloud. The decision to proceed will depend on many factors, including the state of the cloud market place, its understanding of the risks, and its understanding and assessment of its ability to provide its own IT effectively.

We can consider the state of the ecosystem as being represented by a number of variables, some of which may fluctuate and change quickly (fast variables) and some of which will change more slowly (slow variables). Within the cloud ecosystem, these variables will represent a range of factors that summarize both how all the entities interact and the effects of exogenous influences. There can be several different variables representing different aspects of the ecosystem:

- Some variables might represent the state of cloud adoption — this might include aspects such as the proportion of companies using cloud services, the proportions of companies mainly using cloud — other variables may represent the diversity of the service environment;
- The relative security position — here variables might represent the likelihood of a cloud service being hacked, or of internal IT being attacked, along with average incident costs;
- Some variables may represent the different cost structures — for example, average transaction costs for different types of service; and
- Some might represent birth/death rates of services.

As well as variables that represent the state of the ecosystem, we must also consider the resources that are required by the various members to perform their IT tasks. Two examples are staff (with IT skills) and capital to invest in new systems or services. We can think of each entity as pulling the resources it needs, but in doing so there is a cost (or resources may simply not be available). For example, a company with an internal IT department will need to attract staff to its own location, but, depending on the state of the ecosystem, suitably skilled people may be more attracted to work for cloud service providers. This happened during the internet bubble, where companies found it hard to compete with the promises offered by start-ups. Overall, external drivers may limit or expand the pool of resources. For example, government-sponsored training programmes or the availability of credit because of economic cycles. The amount of resources required may also be influenced by technology changes; for example, automation technologies may reduce staff costs and improved IDEs and specialized languages may reduce software development costs.

A company seeking to make decisions around the cloud will look at the state of the ecosystem, as well as the availability of resources within the ecosystem, and use this information to make decisions that aim to maximize (or at least satisfice) its utility. Considering the stewardship parts of cloud utility, a company's decisions may depend on the availability of skilled security staff resources to fulfil the CIA targets within its stewardship utility and therefore look to use cloud services to help. At the same time, the decision will be influenced by the company's view on the relative frequency of security incidents in and out of cloud, along with the associated incident costs.

Clearly the state of the ecosystem is of critical importance for the way in which participants react to one another and for how the cloud-based business ecosystem evolves, but perhaps of equal importance is how information flows around (a given state of) the system. Many participants may make decisions based on having the wrong perception of the state of the system. A key factor within the cloud ecosystem therefore becomes the availability of communication channels that allow information to flow between participants. Examples of communication channels include the media's reporting of certain events and conversations on the golf course. A company looking to switch its financial processes into the cloud is likely to be influenced by its perception of the extent to which other participants in the ecosystem are getting business value and cost savings from the cloud, as well as its perception of the risks. It might not have an accurate view on the state of the ecosystem and instead might base its decisions on reports from neighbouring participants and/or the media.

B. Feedback Loops

Ecologists consider ecosystems that vary overtime because of feedback loops. For example, a fast variable may be the population size of a particular animal. 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. They then talk of two different factors being responsible for these changes, the ecological factors and the societal factors (i.e., the effect of humans on the ecosystem).

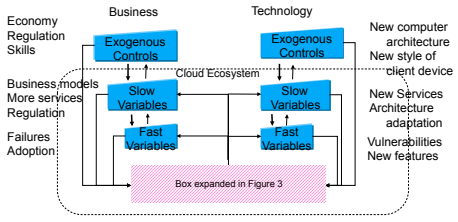

Fig. 2. Understanding the dynamic cycles within the cloud ecosystem

Within a cloud ecosystem, we can draw out similar feedback loops (see Figure 2). Instead of drawing out ecological and societal factors we draw out the business environment in which all the companies operated and the constant technological changes. This assumes the three-layer model introduced in Section II (see Figure 3).

As businesses transition key services (e.g., finance, HR, IT) to the cloud, utility satisficing may require quite rapid resource reallocations. This leads to an example (among many) of a reinforcing feedback loop affecting and affected by the fast variables of cloud adoption. Other examples of feedback loops will derive from compromised service standards in security and stewardship, which cause customers to withdraw from

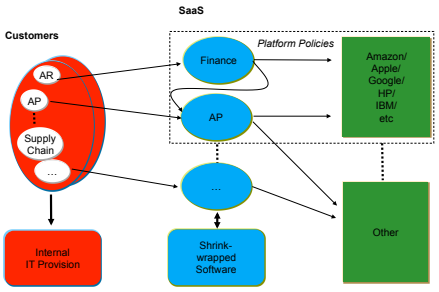

Fig. 3. Understanding the dynamic cycles within the cloud ecosystem (expansion of detail in Figure 2)

cloud service providers, so affecting overall rates of cloud adoption. A critical factor affecting the cycle time of feedback loops will be the 'lock-in' costs associated with withdrawing from the ecosystems. High lock-in costs may weaken the incentives of service providers to to act as good stewards.

Slow variables include things like the rate of addition of new services to the ecosystem, and corresponding changes to companies' business models, be they existing or potential ecosystem participants. Other examples include the adoption of technological developments, such as trusted infrastructure platforms, that may enhance stewardship models. Global economic cycles are also examples of slow variables.

C. Shocking the Ecosystem, Resilience, and Sustainability

As we have mentioned, for Chapin, Kofinas , and Folke (in [13]), the key concept is that of *resilience-based ecosystem stewardship*, which they say 'involves responding to and shaping change in social-ecological systems to sustain the supply and opportunities for use of ecosystem services by society'. This view is useful because it emphasises two aspects of stewardship that are of particular concern in cloud ecosystems:

Resilience: The capability of the system to recover from attacks that successfully compromise its declarative stewardship objectives (e.g., the confidentiality of a customer's PID is breached) or inhibit the effectiveness of its operational mechanisms to deliver its objectives (e.g., the loss of availability of authentication server); and

Sustainability/Adaptability: The capability of the ecosystem to adapt to changes in its composition (infrastructure providers, service providers, consumers), in its required functionality, in its regulatory environment, and its threat environment.

VII. Societal Impact

The emergence of cloud-based business ecosystems of the kinds sketched in this paper will raise a number of challenges for their host societies. We sketch a few immediate ones.

How can a sustainable, resilient marketplace of services be facilitated? eBay and Amazon have highly developed marketplaces for products, but services ecosystems will, we suggest, have inherently richer interdependencies and require correspondingly richer models of trust and assurance.

What are appropriate models of trust and assurance for businesses operating in cloud ecosystems? Accounting-based

models of stewardship (e.g., [12]) may help here, as may the vast literature in models of trust and reputation.

What models of (dynamic) resource allocation and congestion handling will be acceptable to host societies?

What are the possibilities for recombinant innovation (for products, this is about new products emerging from combinations of existing products) and combinatorial innovation (for products, this is about new uses of existing products) in service offerings? How can cloud-based ecosystems be configured to encourage the development of such innovation?

VIII. Conclusion

We are seeking to provide a new way to think about cloud that deals with the complexity of the ecosystem. This work is based on solid foundations (i.e., system modelling, economics, and ecology).

- Information stewardship is a concept that encompasses the basic concepts of information security, but includes also respect for objectives, for ethics/values, and emphasises concepts of sustainability and resilience.
- Cloud ecosystems will be subject to constant pressures that may tend to degrade them — we need to understand strategies to support sustainability.
- Cloud ecosystems are live and will be subject to shocks — it is important to understand when they will be resilient and to have strategies to improve resilience.

Successful information stewardship is fundamentally linked to the way the cloud emerges, and making the development of cloud to go in the right direction.

Future work is to build models of ecosystems as bases for simulations: with so many interacting entities, ecosystem behaviour cannot be expected to be analytically predictable. Models will be core for those using, providing, and regulating to make better decisions and ensure information is safe.

Other future work involves assurance models — possibly building on work in accountancy on stewardship in financial reporting (e.g., [11]) and stewardship-based economics (e.g., [18]) — for information stewardship in the cloud. Such a programme will help us identify the core concepts of information stewardship, extending those (CIA) of of information security.

Acknowledgments

We thank Anne McGeachin, Brian Monahan, Andrea Patacconi, Joe Swierzbinski, Julian Williams, and Chew Yean Yam for their advice.

References

[1] A. Beautement and D. Pym. Structured systems economics for security management. In Tyler Moore, editor, *Proceedings of the Ninth Workshop on the Economics of Information Security (WEIS 2010)*, 2010. http://weis2010.econinfosec.org/papers/session6/weis2010_beautement.pdf.

[2] Ken Binmore. *Rational Decisions*. Princeton University Press, 2008.

[3] G. Birtwistle. *Demos — discrete event modelling on Simula*. Macmillan, 1979.

[4] Yuanyuan Chen and Anandhi Bharadwaj. An Empirical Analysis of Contract Structures in IT Outsourcing. *Inf. Sys. Res.*, 20:484–506, 2009.

[5] M. Collinson, B. Monahan, and D. Pym. A logical and computational theory of located resource. *J. Log. Computat.*, 19(b):1207–1244, 2009. Advance Access 22 July, 2009. doi:10.1093/logcom/exp021.

[6] M. Collinson, B. Monahan, and D. Pym. A discipline of mathematical systems modelling. To appear: College Publications, London, 2011.

[7] M. Collinson and D. Pym. Algebra and logic for resource-based systems modelling. *Mathematical Structures in Computer Science*, 19:959–1027, 2009. doi:10.1017/S0960129509990077.

[8] Matthew Collinson, Brian Monahan, and David Pym. Semantics for structured systems modelling and simulation. In *Simutools 2010*. ACM Digital Library, EU Digital Library, ISBN: 78-963-9799-87-5, 2010.

[9] Core Gnosis. www.hpl.hp.com/research/systems_security/gnosis.html.

[10] George Coulouris, Jean Dollimore, and Tim Kindberg. *Distributed Systems: Concepts and Design*. Addison Wesley; 3rd edition, 2000.

[11] Accounting Standards Board (ASB) et al. Stewardship/Accountability as an Objective of Financial Reporting: A comment on the IASB/FASB Conceptual Framework Project. Technical report, ASB, Foreningen af Statsautoriserede Revisor, Deutsches Rechnungslegungs Standards Committee, Komitet Standardów Rachunkowoci and EFRAG, 2007. Available from www.efrag.org.

[12] F. Gjesdal. Accounting for Stewardship. *Journal of Accounting Research*, 19(1):208–231, 1981.

[13] Chapin III, F. S., G. P. Kofinas, and C. Folke (editors). *Principles of Ecosystem Stewardship: Resilience-Based Natural Resource Management in a Changing World*. Springer-Verlag, 2009.

[14] C. Ioannidis, D. Pym, and J. Williams. Investments and trade-offs in the economics of information security. In Roger Dingledine and Philippe Golle, editors, *Proc. of Financial Cryptography and Data Security '09*, volume 5628 of *LNCS*, pages 148–166. Springer, 2009. Preprint at http://www.cs.bath.ac.uk/~pym/IoannidisPymWilliams-FC09.pdf.

[15] C. Ioannidis, D. Pym, and J. Williams. Fixed costs, investment rigidities, and risk aversion in information security: A utility-theoretic approach. In Bruce Schneier, editor, *Proc. 10th Workshop on the Economics of Information Security*. Springer, 2011. In press.

[16] Christos Ioannidis, David Pym, and Julian Williams. Information Security Trade-offs and Optimal Patching Policies. *European Journal of Operational Research*, 216(2):434–444, January 2012. doi:10.1016/j.ejor.2011.05.050.

[17] R. Jain. *The Art of Computer Systems Performance Analysis*. Wiley, 1991.

[18] R.W.Y Kao. *Stewardship-based economics*. World Scientific, 2007.

[19] R.L. Keeney and H. Raiffa. *Decisions with multiple objectives: Preferences and value tradeoffs*. Wiley, 1976.

[20] Peter Mell and Timothy Grance. The NIST Definition of Cloud Computing (Draft). Technical report, National Institute of Standards and Technology, U.S. Department of Commerce, 2011. Special Publication 800-145 (Draft).

[21] R. Milner. Calculi for synchrony and asynchrony. *Theoretical Computer Science*, 25(3):267–310, 1983.

[22] R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.

[23] B. Nardi and V. O'Day. *Information ecologies*. MIT Press, 1999.

[24] Jon Norberg and Graeme S. Cumming, editors. *Complexity Theory for a Sustainable Future*. Columbia University Press, 2008.

[25] P.W. O'Hearn and D.J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, June 1999.

[26] Donn B. Parker. *Fighting Computer Crime: A New Framework for Protecting Information*. Wiley, 1998.

[27] David Pym and Martin Sadler. Information Stewardship in Cloud Computing. *International Journal of Service Service, Management, Engineering, and Technology*, 1(1):50–67, 2010.

[28] David Pym, Martin Sadler, Simon Shiu, and Marco Casassa Mont. Information Stewardship in the Cloud: A Model-based Approach. In *Proceedings of CloudComp 2010*, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering (LNICST). Springer, 2011. To appear.

[29] David Pym and Simon Shiu. Security Analytics. *IISP Pulse*, 4:12–13, 2010. HP's 'Security Analytics' service: http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA3-2046EEW.pdf.

[30] J.C. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. In *Proc. LICS '02*, pages 55–74. IEEE Comp. Soc., 2002.

[31] Francisco J. Ruge-Murcia. Inflation targeting under asymmetric preferences. *Journal of Money, Credit, and Banking*, 35(5), 2003.

[32] John D. Sterman. *Business Dynamics: Systems thinking and modeling for a complex world*. McGraw Hill, 2000.

[33] H. Varian. A bayesian approach to real estate management. In S.E. Feinberg and A. Zellner, editors, *Studies in Bayesian Economics in Honour of L.J. Savage*, pages 195–208. North Holland, 1974.