DBSy in a Commercial Services Context♦

Brian Monahan, Frederic Gittler, William Horne, Simon Shiu, Adrian Baldwin, Chris I. Dalton, Patrick Goldsack, Richard Taylor, Chris Tofts, Mike Yearworth
Trusted Systems Laboratory
HP Laboratories Bristol
HPL-2005-141
August 4, 2005*

DBSy, domain based security, information security, information assurance, commercial ICT services management, ITIL, ITSM, COBIT, BS 15000

DBSy (Domain Based Security) is a set of notations and techniques developed by QinetiQ specifically for the UK MoD, a large distributed organisation. DBSy provides a way of describing and assessing business-driven information security requirements for network architectures. This focuses upon how the business requires information to be compartmentalised and how that might be achieved by strategic location of network-level security controls. In this paper we consider how DBSy-style modelling may be applied in a more commercial context of ICT (Information Communications Technology) services, defined and managed according to SLAs (Service Level Agreements). Although DBSy was not specifically designed to handle this situation, we discuss how ideas from DBSy can contribute to a broader security requirements and risk analysis approach that encompass the realm of ICT services and their management. We give a model of a commercial example in the style of DBSy and use that to illustrate our observations.

* Internal Accession Date Only
♦ 1st DBSy User Conference, QinetiQ Technology Center, Great Malvern, UK. 14 June 2005
Approved for External Publication

A version of this paper appears in the proceedings of the 1st Domain-Based Security User Conference, held at QinetiQ Technology Center, Great Malvern, UK, on 14th June 2005.

DBSy in a Commercial Services Context

Brian Monahan, Frederic Gittler, William Horne, Simon Shiu,
Adrian Baldwin, Chris I. Dalton, Patrick Goldsack, Richard Taylor, Chris Tofts, Mike Yearworth

Hewlett-Packard Laboratories

Bristol UK, Princeton USA, and Grenoble France

July 2005

Abstract: DBSy (Domain Based Security) is a set of notations and techniques developed by QinetiQ specifically for the UK MoD, a large distributed organisation. DBSy provides a way of describing and assessing business-driven information security requirements for network architectures. This focuses upon how the business requires information to be compartmentalised and how that might be achieved by strategic location of network-level security controls.

In this paper we consider how DBSy-style modelling may be applied in a more commercial context of ICT (Information Communications Technology) services, defined and managed according to SLAs (Service Level Agreements). Although DBSy was not specifically designed to handle this situation, we discuss how ideas from DBSy can contribute to a broader security requirements and risk analysis approach that encompass the realm of ICT services and their management. We give a model of a commercial example in the style of DBSy and use that to illustrate our observations.

Keywords: DBSy, Domain Based Security, Information Security, Information Assurance, Commercial ICT Services Management, ITIL, ITSM, COBIT, BS 15000.

> "Whether we entrust our decisions to machines of metal, or to machines of flesh and blood which are bureaus and vast laboratories and armies and corporations, we shall never receive the right answers to our questions unless we ask the right questions."
>
> - *The Human Use of Human Beings* (pub. 1950) by Norbert Weiner

Introduction

The concept of services is changing the way ICT is packaged and operated. These changes have considerable impact on the way risk and security analysis can and should be done. This paper considers how models in the style of DBSy could apply to this evolving commercial context.

The impact of the services approach to ICT provision is illustrated by the shift from monolithic applications running in privately managed data centres to one where infrastructure services are shared, and business value is delivered through a flexible mixture of application components. At the same time, there has been a broad realisation that various ways of modelling the business IT systems forms a vital ingredient in the management of networking and shared distributed infrastructure.

The 'state of the art' holistic approach to security is embodied in the ISO17799 standard which although capturing some structure and best practice, is really a checklist that requires significant security experience for its effective use. The DBSy approach uses structured forms of business and infrastructure models to express security requirements and facilitate the analysis of risk. Given the growing complexity of ICT services, this style of approach seems both relevant and highly appropriate.

However, the analysis and discussion in this paper shows that there are significant challenges to finding appropriate ways to model the services context. In particular, to take account of how reliant and how trustworthy the different infrastructure services are, the appropriate risk models will need to be much finer grained and more business-process focused. Similarly, much work is needed to

help commercial enterprises to understand and analyse the business priorities and risks in using ICT services.

Section 1 gives our overview of the DBSy methodology as we see it. Section 2 gives a simple commercial example in the style of DBSy, followed by some discussion of pertinent modelling issues. Section 3 provides a broader overview of ICT services, as well as describing key concepts such as service level agreements (SLAs) and the service lifecycle. We end that section with a discussion of emerging standards for ICT Service Management. This section attempts to show some of the complexities that can arise in the management and operation of an ICT service. Section 4 discusses the overall modelling challenges for any security requirement and risk analysis approach that takes account of the world of ICT services. Finally, we end the paper with a brief summary and draw conclusions. As an appendix, we also include a description of some of current HP Labs research into ICT services.

1. What is DBSy today?

In this section, we present our overview of DBSy and what it helps with today. Descriptive material concerning DBSy is available (see [DBSy1, DBSy2, DBSy3]) in addition to the material contained in these proceedings (see the conference CD).

Broadly, DBSy provides a way of capturing and assessing the network security requirements upon communications and more generally, network services, within a large distributed organisation (e.g. governments, military and public/private sector suppliers). The approach focuses upon how the business requires information and it's processing to be compartmentalised – that is, **network separations, services aggregation** and **compartmentalisation-in-the-large**. DBSy is strongly aligned with existing information security and assurance standards such as ISO/IEC 17799 (BS 7799).

Applying to a business situation DBSy involves developing the following models:

- **InfoSec Business Model**: Used to identify the operational business functions as **Domains**, and the required information flows between them, via business-level services such as e-mail, Web access and corporate data base access. The domains specify the maximum classification of information handled and the minimum clearance of the people working in the domain.
- **InfoSec Infrastructure Model**: Used to propose infrastructure **Islands** that provide networking and computing support for business functions. Causeways provide Security Enforcement Points that give protected, controlled access between Islands for each of the required network services.
- **InfoSec Security Architecture Model**: The two models above combined. Each domain must be contained within some Island and each network service is controlled via at least one Causeway. There is at most one Causeway between any pair of Islands.

An important and natural underlying assumption behind DBSy is that the infrastructure – that is, the Islands and Causeways – are typically operated, managed and ultimately owned by the organisation and its collaborative agencies. The management framework is therefore implicitly understood.

The Compromise Path Analysis technique gives a means to assess the strength of controls placed upon access paths needed to attain a given level of assurance. The placement and assessment of security controls in DBSy is consciously combined with other existing MoD procedures (e.g. JSP 440). In this way, DBSy provides a methodology assessing risks associated with sharing of information between business functions, consistent with MoD policy and doctrine. This assessment helps risk-owners understand the information leakage risks and the extent to which they can be controlled and mitigated.

The complexity of the security architecture captured by DBSy gives a broad indication of the operational risk level that it incurs – the more complex the security architecture becomes, the harder it is to operate assuredly and without serious mishap. It is imperative to *balance* the needs of

getting the business done in the first place, with the need to protect business information assets and resources.

The important characteristics of the DBSy InfoSec approach are that:

1. It is driven from business needs and function.
2. It provides a tractable way of concretely visualising a number of security issues in terms of compartmentalisation of network services.

DBSy makes it easier to see how to prioritise the defence of business components and network services, in accordance with corporate security policy. Notice that the prevailing assumption is one of *mandatory access control* – any service not explicitly permitted is, by default, forbidden and should not be provided.

This approach also provides a sound basis for actively testing the security framework for systematic leakage, using automated penetration testing, automated ICT health-check and automated traffic analysis. Such testing would involve measuring bandwidth adequacy of traffic flows for network services justified by the business – and conversely, active testing and inspection will seek out evidence of network services not supported by the business.

2. What does a DBSy for Commercial ICT Services look like?

DBSy was originally designed and formulated by QinetiQ for use within a large military organisation, the UK MoD. In particular, this means that the current DBSy will need to reflect the MoD's security policies and doctrine, and integrate with existing MoD security procedures and assessment techniques. In this section, we ask the question of how to reapply DBSy-like techniques to capture, analyze and assess large scale security and IT service requirements in a commercial business context.

Purely for the purposes of this paper, we have needed to take certain liberties with DBSy as it is practiced today within the UK MoD and Government. We have tried to keep to as much of the current style and spirit of DBSy as possible, to maximise the relevance of our discussion to the audience for this conference.

Our approach is to give a model of a commercial-world example in the style of DBSy, and then to discuss and highlight the modelling issues that arise.

2.1. Commercial example: Business Model

As promised, we give our DBSy-like model of the security and networking requirements for a (fictional) commercial enterprise, ACME Corp. The InfoSec Business model[1] is presented in Figure 1, with the **value system** used here is described in Figure 2.

In outline, the InfoSec business model identifies 5 company *business functions* (e.g. departments), two of which have direct customer-facing network services, with E-mail & messaging access for both Sales and Marketing, and Web access for the Marketing dept. There is also:

- A company wide internal email/messaging service.
- An internal corporate data base service, accessible by all business functions. (except Marketing).
- There is a shared file-store solely accessible by the Accounts and Legal & Contracts business functions.

For simplicity, we have ignored here the necessary *business IT governance* issues that surround, for example, business activities such as the preparation of accounts. Typically, there are externally appointed auditors and various regulatory oversight authorities for a vertical market segment (e.g.

[1] This model does not include Portals and other "environmental" aspects of DBSy. In our opinion, these should rightly be elaborated in yet another modeling "view" e.g. an InfoSec Environmental Model.

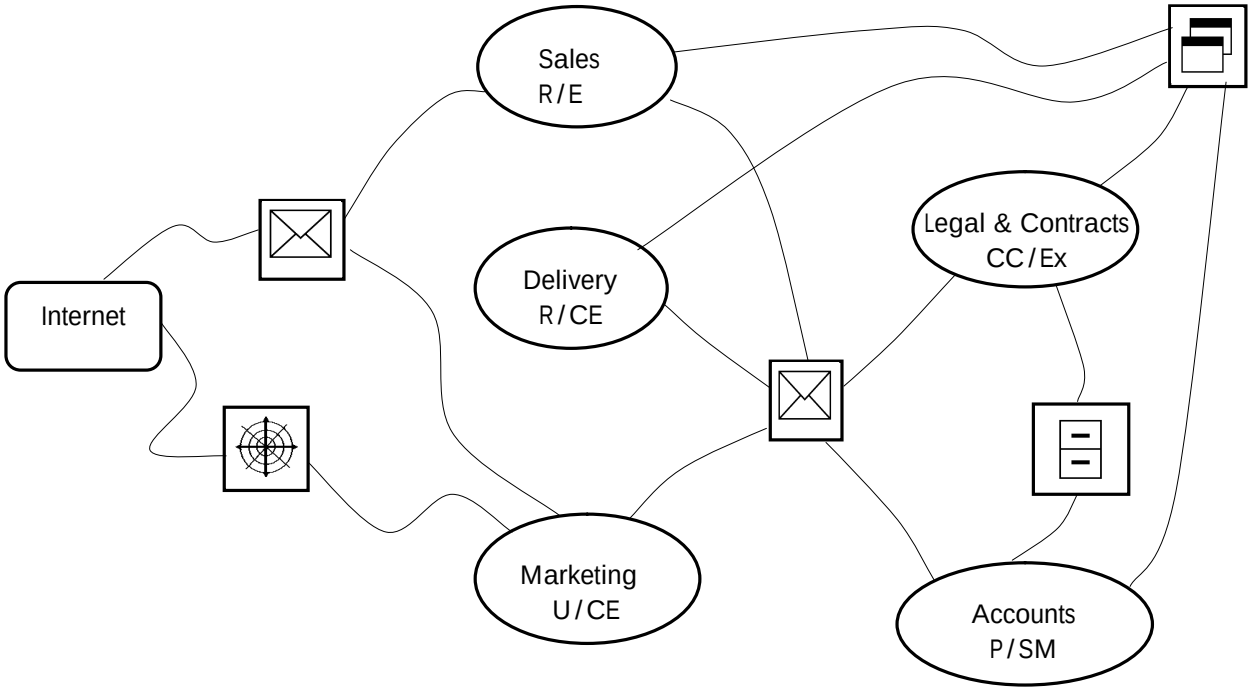

Figure 1: InfoSec Business model for ACME Corp.

ACME Corp. Protective Markings	
U	Unclassified
R	Restricted
CC	Company Confidential
P	Private

ACME Corp. Clearance Levels	
CE	Contractor or Employee
E	Employee and above
Ex	Executive and above
SM	Senior Management

Figure 2: Value System for ACME Corp. Protective Markings and Clearance Levels

the FSA in the UK and the SEC in the US) who would need to periodically review, and thus access, the accounts. More recently, the governance requirements on corporations to manage information appropriately, in accordance with the US based Sarbanes-Oxley and HIPAA legislations, is also of particular relevance here.

2.2. Commercial example: Architecture Model

An InfoSec Architecture Model that matches the business model given above is presented in Figure 3 below.

Here we have skipped formulating the Infrastructure model separately – and gone directly to the Architecture Model itself. In any event, the Infrastructure Model can be extracted, as required, from the overall Architecture Model.

In developing this model, we have had to "refactor" the Business model slightly to accommodate the discipline of having at most one Causeway between any pair of Islands. This implies that each Causeway has exactly two connecting links. Importantly however, the refactored business model still supports exactly the same information flows and service access capabilities as before – it is equivalent to the one given earlier. The changes were:

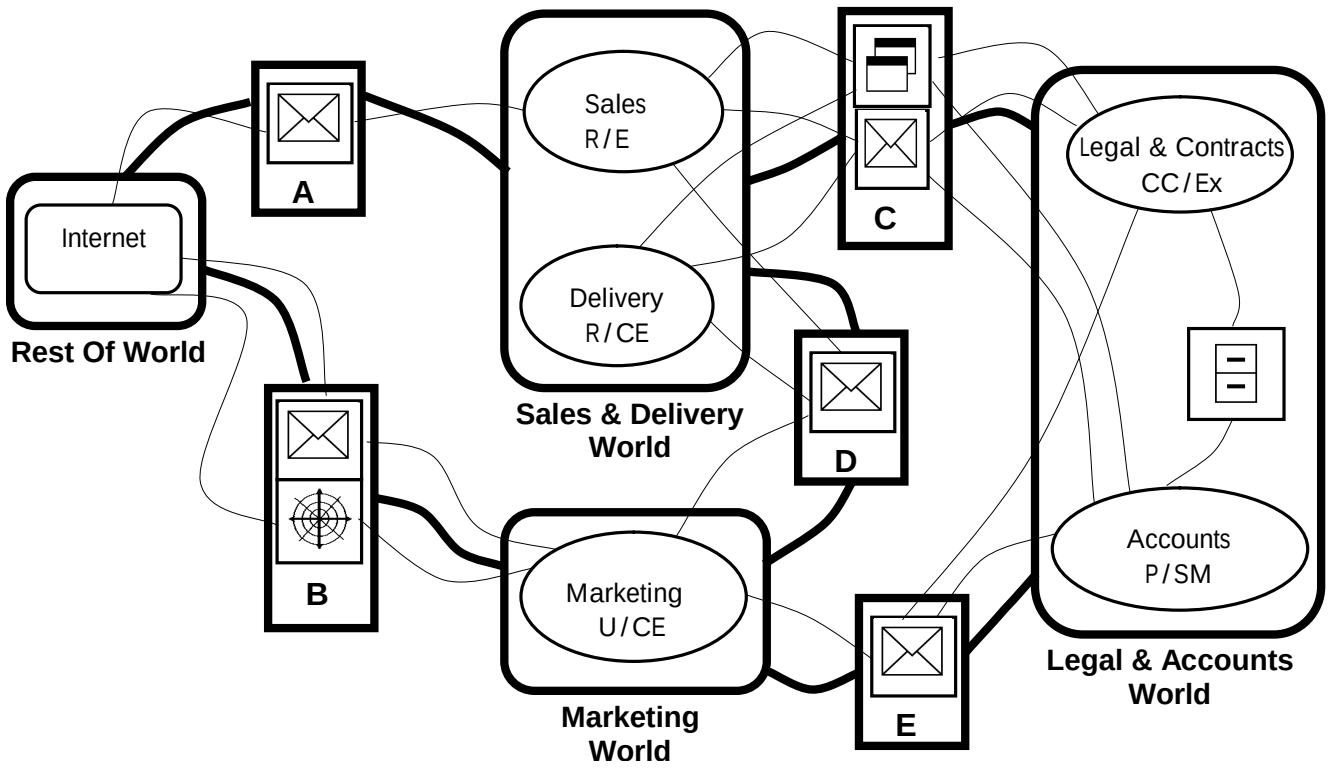

Figure 3: An InfoSec Architecture model for ACME Corp.

- The customer-facing e-mail/messaging network service from the Internet to the Sales and Marketing business functions has been split into two, supported by Causeways A and B.
- The corporate internal e-mail/messaging service is split into three, supported by Causeways C, D and E.

The model proposes five Causeways (two external facing) and three Islands, corresponding to strongly contained networks. The Islands are as follows:

- **Sales & Delivery World**: This contains the Sales and Delivery business functions – it contains information classified up to company Restricted and permitting access to general employees and contractors.
- **Legal & Accounts World**: This contains the Legal & Contracts, and Accounts business functions. This processes information classified up to company Private and is accessible by Executive grade personnel.
- **Marketing World**: This contains the Marketing business function. The information handled there is unclassified and is directly handled by general employees and contractors.

2.3. Modelling Issues Discussion

Having completed our models, we now discuss some of the modelling issues that arise.

Corporate-wide network services and the Internet:

We can immediately see that all of the companies business functions are interconnected via an e-mail/messaging service, quite *irrespective* of clearance and protective marking. Such a service will typically be regarded as a required business practice and enabler, permitting broad communication within the company. What the above business model does *not* capture are the various corporate mailing lists, mail aliases and project groupings that map into the organisational

structure. A finer grained analysis involving roles of personnel is necessary if any further constraint on information flow is required.

Some might have felt that the example above is too extreme in requiring universal e-mail/messaging for all employees. In actual fact, it is probably rather conservative in that some commercial organisations **require** customers to directly access their systems via the Internet – just consider any number of web-access portals (e.g. Amazon and Internet Banking). In such a case, the Internet would not just be a source of threats – it provides the means for customers (even new, previously unknown ones) to access corporate systems.

The true role of the Internet is neutral – it is neither fundamentally beneficial nor harmful, much like the telephone network is today. However, access via the Internet is now very much a business necessity in these days of e-commerce, web services and cheap, readily available networking. The real danger is one of network isolation at the beck and call of an adversary.

The "Value-system" and access-control:

The "value-system" defines the *maximum protective marking* (or classification) of the information handled and the *minimum clearance* of people that are permitted access. This imposes a natural *constraint* upon the access-control upon resources and assets in those domains. However, there may be other constraints needed to adequately specify access-control.

In particular, such analysis will involve characterising the processing systems, the roles of the people involved and the overall workflow for each type of information. For example, within an operating domain like "Accounts", there may actually need to be a number of different roles of differing seniority and responsibility (e.g. account clerks, account managers, account administrators, account executives). As an illustration using our example above, although everyone in the Acme Marketing department can be a general Contractor or Employee, this fails to say anything about who has *editing rights* to modify the content of web pages. There will necessarily be some form of business process for authoring and change control in which new draft content is submitted and authorised for publication.

Broadly, concepts familiar from Role Based Access Control will need to be combined with business process modelling to capture the necessary process-oriented constraints. An important concern here is that, in practice, people having roles with low clearance will, from time to time, have to handle information having *great sensitivity* and on the face of it, require higher clearance – for example, the accounts clerks preparing sales data that contributes to the quarterly revenue figures for a large publicly listed corporation. The point here is that the low-clearance people involved are not in a position to interpret the information overall – they may know that the information is of value – but not be able to fully appreciate its significance.

In commercial practice, an extension of Role Based Access Control known as "matrix management" is increasingly becoming used where people and the systems acting on their behalf have several roles, each with specific duties, responsibilities, permissions and capabilities, in usually multiple projects. In this approach, individuals are both empowered and expected to use their initiative to fulfil their obligations arising from all these roles. Inevitably, this creates conflicts which each individual is supposed to resolve as skilfully as possible. In particular, being a high-ranking executive may not confer you with immediate access rights to perform systems actions etc., just being a systems administrator does not confer executive capabilities. Clearly, the notion of "value system" as used in DBSy does not directly fit with this more complex, finer-grained approach.

User-modelling, accounts and strong identity:

Another related issue concerning access-control arises in connection with modelling interactions with data involving individuals. This situation arises when the organisation needs to maintain, as a core part of its business, personal profiles such as accounts relating to customers or individuals.

Examples where this arises in business settings abound, including law enforcement, healthcare, banking and financial services. Such data is generally tricky to deal with since the external

customer should be able to examine and manipulate details from their own account – but no-one else's. Similarly, corporate employees with specific roles may be able to view a range of accounts of a certain type, but not others. This kind of access-control and detailed business logic is not represented as a part of the DBSy InfoSec business model.

Broadly, for certain applications, DBSy needs to provide a more explicit account of services involving *strong identity* and *authorisation*, such as *single sign-on*, *ticket authorisation* and *data base account services*. In some sense, this has to involve a combining a network security view with the application services view of security.

Shared Networking Services and Causeways:

Each Island represents an isolated network in which gateways provide specific access to external network services. An important source of threats and vulnerability however, are the common network services such as DNS, DHCP, LDAP (for single sign on/corporate authentication) etc, that most be kept operational just to keep the networking alive. This is not represented at any level of detail. Granted, shared services may not need to be viewed at the higher levels – they can be assumed to be maintained, from a customer's point of view. However, from a providers point of view, the management and provision of these standard services will need to be explicitly taken into account (at least for planning/provisioning purposes).

In a similar vein, the Causeways represent a convenient abstraction at a higher level which it is understood needs to be translated into a more explicit networking requirement to provide gateway support at each end of the link. Thus, implementing a Causeway could possibly involve providing hardware such as routers and servers, plus appropriate network management and control software within each Island that the Causeway is linked to.

The main observation is that further drill-down is needed from a DBSy architecture model into the specific network architecture to see the implied requirement upon the networking infrastructure and its management.

Business processes, timeliness of information and dynamic security classifications:

In our example above, we glibly assigned the Acme Corp. Marketing department to have a Maximum Protective Marking of "Unclassified", with a Minimum Clearance of "Contractor or Employee". This perhaps reflected an assumption that the material handled would indeed be common knowledge (e.g. web pages cntaining published product information and price lists).

But what happens when Acme wants to release a new product or to conduct an advertising campaign that strategically promotes some suite of products? Presumably, the exact nature and timing of that information would be of extreme interest to a commercial competitor – and thus highly commercially sensitive, right up until the product or campaign is released. After that point, one hopes that everyone, including Acme's opposition, should be highly aware of what Acme has released! In practice, there will be people in Marketing who, according to their role and job function, will be aware of the new product plans prior to their release – and no doubt these individuals will have some continuing responsibility for the new product's marketing campaign following the launch. Thus, business function, role and process are all inherently intertwined, each having impact upon security and integrity concerns.

Similar concerns about timeliness of information classification arise in the military context – e.g. battle plans. Another common business example is the publication of quarterly revenue figures for any large, publicly listed corporation. These figures have to be compiled and prepared in a secure manner internally prior to their release – knowledge of these figures prior to release could be used to buy/sell shares very advantageously i.e. insider trading. Once published, their effect of these figures on shareholder value is extremely public.

This issue is to do with both the *upgrading* and *downgrading* of information – one can expect this to happen quite routinely, for all manner of reasons as information flows around the organisation. From the above examples, it is clear that the security classification and the significant nature of the information changes as a result of the actual business processes involving people and technology.

At the moment, DBSy provides too coarse grained a view of business process. What is needed is some way to extend the coarse grained views, to drill-down to expose actual information flows in greater detail.

Legacy networking in commercial organisations – automated network service discovery:

What is presented in the DBSy-style architecture model given above is something of a fantasy in today's commercial world. Unless the organisation already operates a defined network connection policy, commercial networks will not typically develop and evolve in this "compartmentalised" way.

Commercial networks have been traditionally characterised by a form of organic and ad-hoc growth in which authorisation requests from random parts of the corporation would be granted network access to other random parts of the business. Indeed, because of permissive default access policies, it may be technically difficult or very costly to *prevent* inappropriate connections from being established, without at the same time removing required business access. Without a corporate network connection policy that can be effectively applied, system administrators are unable to decide which requests to allow and which requests to deny – they would be compelled to accede to all access requests from management.

Common reasons/excuses for not having a network connection policy in organisations are:

- **Overhead costs of policy enforcement and administration**: An effective policy implies asset management and actively maintaining knowledge about the current extent of the networks. This inevitably involves having some kind of *configuration database* system for network components which will need to be actively kept up-to-date and well stocked with relevant information.
- **Perceived lack of business agility**: Policy enforcement requires those people wanting connections to actively seek business reasons or similar justifications for them. This process inevitably takes time and involves bureaucracy, whereas a permissive default access policy would always permit connections immediately without fuss and bother. The downside is that it rapidly becomes very hard to know at any given time what the extent of the local network actually consists of.
- **The organisation is never the right size – it is always either "too small" or "too large"**: Small to medium businesses may only have modest networking requirements – for which it seems like "overkill" to enforce a network connection policy. At the other extreme, large corporations that have rapidly developed and expanded their networking (perhaps by acquisition) without a network connection policy could find it very costly to impose one later on.

  The danger is that by the time the business need for a network policy has become obvious, the organisation may have grown past the point where that policy would be most effective.

Naturally, lack of policy would not be tolerated in any organisation where network security management is recognised early on as a necessary part of doing business. At the same time, there are legitimate requirements for flexibility of network service provision, in accordance with changing business circumstances.

The challenge for a commercial DBSy-style security requirements method is to be able to integrate with *network auditing* and *services discovery* mechanisms and technologies that can capture and map the presence of existing network services, along with their configuration topology, in a cost-effective manner. Armed with this information, it would then be possible to compare the discovered topology against those topologies that are required and permitted by security doctrine and policy. In this way, violations of network connection policy, accidental or otherwise, could be more routinely detected in an automated manner. An effective service discovery capability combined with a DBSy-style security requirements methodology would encourage more effective management of network connection policy.

Enforcement of corporate network service boundaries is getting harder – Virtual groupings:

Commercially available communications technology is moving ever further away from a simple-minded notion of network boundaries. Access to Instant Messaging technology such as Jabber (XML Messaging) and services like MSN, Sharepoint and Virtual Office are now a common business requirement. Such technology explicitly incorporates both e-mail with file sharing and distribution – which DBSy currently treats as separate functions. This technology allows easy creation, operation, and then disbanding of dynamic, ad-hoc virtual groupings for e-mail and project-based conferencing, in accordance with business demand.

It is becoming harder and harder to maintain any semblance of network boundary using just traditional static gateways and firewall technology at the pure network level. Security management and enforcement technologies need to co-evolve to meet emergent challenges resulting from continuing advances in communications and networking. In particular, implementing security requirements will increasingly have to exploit both network and application layer technologies in a coordinated manner.

Human behaviour vs. formalisation of business processes and security policies:

There is a healthy degree of scepticism about the value and the effectiveness of formalising business processes and security policies. If it were only a matter of regulating fully automated systems that only occasionally involved limited kinds of human interaction (interference?), then business processes could be safely subsumed within traditional software development practice.

The problem is that it is very hard to fully anticipate all the situations that a complex distributed system could encounter – especially so when that system is a large, distributed organisation involving people. Attempts to rigidly over-formalise the situation will tend to result in *inflexible* systems, often having the characteristics of being *brittle* and *unforgiving* – they tend to unexpectedly *jam*, to *abort* or act *erroneously*, typically in a highly inconvenient and damaging manner.

To be workable, processes involving people need some degree of "latitude". A positive benefit of this latitude (or shades-of-grey) is to encourage people to exercise their initiative and judgement in a responsible, accountable manner. This provides the opportunity for people to act positively and thus not condemn the project to predictable failure. There are many examples where human intervention has turned what could have been a fully automated disaster into a mere crisis.

More specifically, trying to impose *unnecessary* timing constraints and requiring *too many* guarantees on any kind of distributed system, including business level ones, classically has the effect of increasing the complexity of systems implementation. Unfortunately, unbridled complexity also has the effect of vastly increasing the number of ways that things can go wrong and misbehave. This effect is as true for software systems as for all others. Security vulnerabilities often arise at the interfaces between sub-systems – increase the complexity and ways of interfacing, and the potential for error and vulnerability increases.

The challenge is to construct our systems in such a way that they have the latitude and spare capacity to respond effectively to emerging, changing circumstances in both the threat profile and the business opportunity. This will only be possible by realising that no one can fully anticipate all future developments and that our systems need to be constructed *openly*, allowing the flexibility to adapt and change appropriately to the prevailing circumstances.

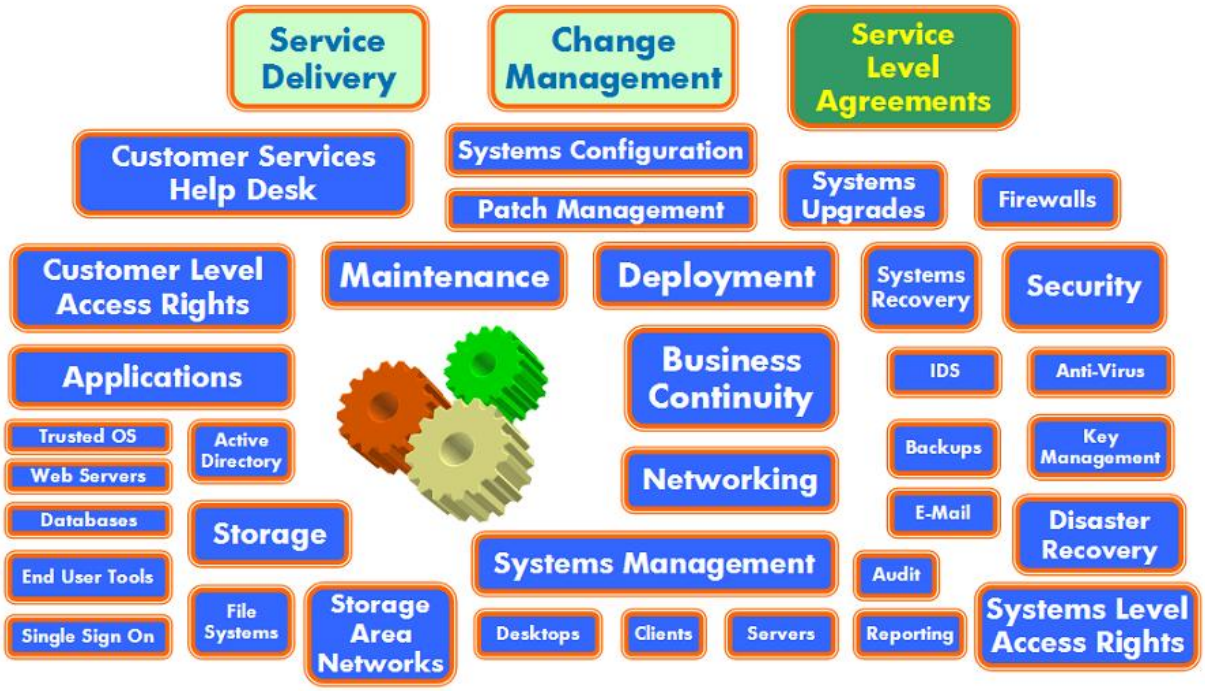

Figure 4: What are ICT Services?

3. Commercial ICT Services

The above DBSy models have neatly partitioned up the ICT infrastructure into separate, compartmentalised networks and services. The model identifies (in outline) the business functions to be supported by each network, together with business-level network services such as e-mail & messaging, which connect these networks to other Islands. Figure 4 above illustrates many of the ingredients that can go to make up typical ICT services.

3.1. Partitioning networks and services

Partitioning is a *necessary* step in deciding how to provision the underlying infrastructure and how this will be operated and managed. It provides key input to decision-making surrounding what services are to be operated by internal agencies (i.e. in-sourced) or operated by external agencies (i.e. out-sourced).

From the ACME Corp. example above, a natural contender for ICT infrastructure outsourcing might be the **Marketing World** Island since the business function (Marketing) handles unclassified material and the clearance of the corporate personnel accessing the systems is not high. In any event, some of the connecting network services, such as Web access (i.e. Web services) are unlikely to be strategic core competences for ACME Corp. and so running that Causeway service may be outsourced also. Finally, it may be decided that Marketing is not a core competency of ACME Corp., and thus the business function in its entirety might be outsourced.

We can immediately see that the InfoSec Architecture Model plays a key role in assessing and identifying the boundaries of different combinations of business function and/or ICT Infrastructure about which service provision decisions can be realistically made.

Simply by way of illustration, suppose that a business decision is taken to out-source the Marketing World as an ICT Service, whilst keeping the Marketing business function in-house. We tentatively suggest in Figure 5 below an amended DBSy infrastructure graphical notation, which shows that there is a "rehosted" **Marketing World Service** in a separate Island with a new Causeway linking to a vestigial **Marketing World** Island. Note that the new Island is now managed by another organisation, called *ICT Services Inc.*. However, where is the management interface? Who manages the intervening Causeway, X? This example suggests that management interfaces need to be explicit.

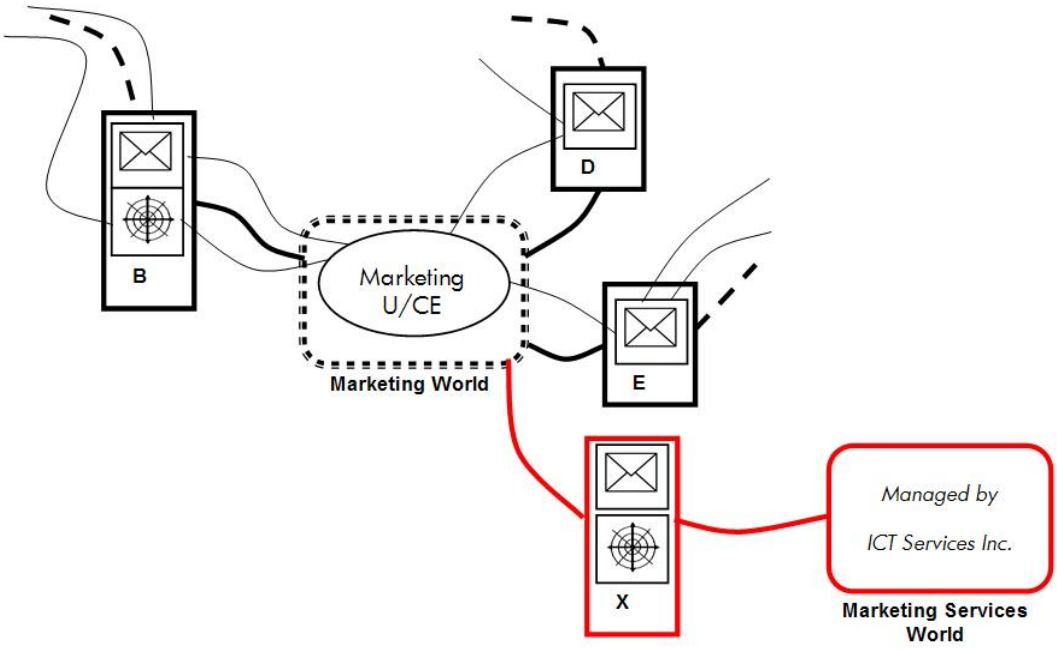

Figure 5: Rehosting the services associated with Marketing World

3.2. What does service provision involve?

It is important to carefully distinguish between two levels of service provision: the provision of some set of business functions in their entirety and the provision of the operations, management and provisioning of some portion of an ICT infrastructure serving a particular business function.

Both kinds involve a binding contractual agreement between the Service Customer and the Service Provider describing the services to be provided, usually in terms of SLAs (Service Level Agreements). This contract will also include a Service Credit Model which will describe the payment regime between the Service Customer and the Service Provider for the delivery of the service. The SLAs will be framed in terms of targets involving KPIs (**Key Performance Indicators**) that are numerical statistics gathered at an appropriate time resolution in a secure, trustable manner and used to assess service performance against the Service Credit Model and hence appropriate levels of payment for achieving or exceeding targets and penalties for failing.

Note that SLAs can and will be formulated however the ICT service is delivered, whether it be via an internal corporate department or via some external contractor. Thus, identifying and defining an ICT service in terms of appropriate SLAs is a necessary precursor to deciding how it will be provided in future. As we shall see below, there are many complex factors to be taken into account in taking such decisions.

External Service provision of Business Functions

In this situation, we assume that the business functions themselves are deemed by corporate management to not lie among the corporate core-competencies, and can thus be safely sub-contracted out to one or more specialist service providers. It will typically be the responsibility of the sub-contractor to then decide how the ICT infrastructure for this business function will be operated, managed and provided for.

However, there are complexities here caused by the needs of the surrounding context of corporate ICT infrastructure. Consider again our ACME Corp. example and suppose that the corporation decides, for purely economic reasons, to outsource the Delivery business function. Now, the corresponding ICT infrastructure for that function is hosted within the **Sales & Delivery World**

Island. Consideration of the DBSy Architectural model immediately raises some important questions:

– How would a new business sub-contractor securely interface with the corporate ICT in the **Sales & Delivery World**?

– What is the *impact* on the other business functions hosted within that Island e.g. the Sales domain? In particular, are there any security risks involving access to the *rest* of that Island that arise from that particular business function being outsourced? If so, how serious are they?

These issues might be resolved in any number of ways – for example, one way is to redesign the security architecture to more isolate the Delivery business function, perhaps into its own Island. Another way may be to consider increasing the extent of the outsourcing to the whole of the Island in question – in this case, consider the outsourcing of both the Sales and Delivery business functions to the same business-level Service Provider. This would enable the ICT networking infrastructure to be consistently provided for as well. Another course of action might be to decide not to outsource at all, because of the disruptive impact it will cause to the rest of the business.

We have suggested here that the outsourcing of business functions could well involve a consequential outsourcing of the supporting ICT Infrastructure. Even so, there is at least another option here, such as the existing ICT infrastructure provision being extended to include support for the new external business-level service provider. This could, for instance, involve an internal ICT services organisation offering their services externally, with all the attendant complications that this may involve for existing internal business engagements.

However, all this must be done so as to permit smooth integration with the remaining corporate ICT context. This is the topic of the next section.

ICT Services Infrastructure Operations, Management and Provisioning

The basic idea behind ICT service provision is straightforward: a specialist organisation operates, manages and provisions the ICT systems under a defined contract, providing responsible *stewardship* of some portion of the ICT requirements for their service customer. There are several aspects or dimensions that may be used classify and categorise such services:

1. **Internal vs. External Service Provision:**

   The first aspect deals with whether the service provider is external to the corporate customer or not:

   - **Internal service provider (in-sourcing)**: Corporate-owned specialist ICT services organisation provides and operates the required systems. In this scenario, the corporation has its own ICT services departments and organisations that can take on systems operations and management on behalf of other corporate business units and functions.
   - **External service provider (out-sourcing)**: External specialist ICT services organisation provides, manages and/or operates the required systems.

2. **Mode of operations – Local vs. Remote ICT operations:**

   This second aspect concerns mode of operations and refines the previous aspect: it concerns the need for Desktop/Workstation support in the corporate workplace and the need for server support supplied by local and/or remote data centres. These options are defined by the following matrix:

	SrvLocDC	SrvRemDC
NoD/W	No Desktop/Workstations Servers provisioned/managed in Local Data Centre	No Desktop/Workstations Servers provisioned/managed in Remote Data Centre
LocD/W	Local Desktops/Workstations Servers provisioned/managed in Local Data Centre	Local Desktops/Workstations Servers provisioned/managed in Remote Data Centre

Table 1: Local vs. Remote service provisioning and management

Local in this case means "on-site and/or only internal LAN required" and is directly embedded within the corporate workplace. Remote means "off-site and/or requiring external networking access" and service provision is located externally to the corporate workplace.

Significantly, it would be a matter of contract as to whether the service provider could further sub-contract and outsource (all or part of) the original service customer's requirements transitively to some other provider. The original customer may or may not need to have legal control of sub-contracting.

There are some further consequences of such a decision: use of a Remote Data Centre in this way implies an additional *cross-boundary* networking between the corporate workplace and the external service provider. The issue here is that all access to this service was previously contained entirely internally – but now it is being made external. Naturally, there will tend to be significant security issues and risks associated with creating and providing such access – and these risks may substantially offset the cost savings of service provision.

In any case, the particular DBSy Architecture Model describing access to the service should be revised to take the new access requirements into account. The model may then help to identify the risks involved and to help mitigate these security concerns. Because DBSy architectures assume that Islands are owned and operated internally by the corporation, the DBSy notation may need to be extended to clearly account for external IT service provision.

Interestingly, similar concerns will arise when going the other way around, when amalgamating IT service provision from potentially different levels of security regime. Bringing two or more IT services together where they had previously crossed organisational boundaries (e.g. corporate mergers) implies operational changes and potential for protective marking and clearance level mismatches that could need considerable effort to consolidate. Again, DBSy Architectural Models play a useful role here.

3. **Data and Software Management**:

The third aspect concerns data and software management – how and where is data/software to be handled? Where does data reside in the system? How are these locations protected? What about "temporary/transient" data storage? Is that really necessary and if it is, how is that managed? Clearly, if the entire network is embedded entirely within the corporate enterprise (i.e. no outsourcing), then all required database access can be adequately contained – in principle.

As before, this aspect refines the one before. The question is to do with the location of database systems and the kind of network access required for all the different kinds of processing over that data. In particular, there may to be several data base sources for processing whose results are delivered to yet other databases. It is thus conceivable that all the source and target databases are local with the processing being done remotely – or in some other combination, such as local processing and remote databases.

Software management involves maintaining the standard IT software stack to enable the customer's applications software and networking to perform in an effective manner. Sources of

change to the software environment are upgrades to hardware systems, upgrades and repair patches to systems software (e.g. security patches) and upgrades and patches of the customer's application software. Of course, any of these upgrade activities could have serious knock-on consequences for the other areas – for example, loading systems patches might require an upgrade of the application software – and vice versa.

All of this activity must be managed and scheduled to minimise disruptive impact. In particular, extreme care must be taken when upgrading business-critical applications and systems.

Another factor concerns the application software used for processing data. There may be business-derived constraints on the manner in which the software may be used to process this data. For example, the software may represent a significant investment by the corporation – it may embody some high performance and efficient processing capability, conferring significant competitive advantage to the corporation which must therefore be protected. Such constraints may arise to protect significant investments in R&D, due to IP or security concerns. As a result, the processing itself may have to be managed and protected in a way that is separate from the databases used to supply input and receive output.

None of the earlier discussion has yet touched on equipment ownership. One alternative is that all of the equipment is owned by the customer, so that the service provider is operating, administering and managing it on their behalf. This would imply that the service provider needs physical access to the corporate workplace (i.e. as contractors) on a routine basis. Another alternative might be that the equipment provisioning is entirely offsite and remotely accessed via networking. However, even in this scenario, some of the equipment will necessarily remain physically in the corporate workplace.

3.3. Lifecycle phases of an ICT Service

The purpose of ICT services is to operate, administer and manage ICT infrastructure on behalf of their service customers. Naturally, this is highly process-intensive, being principally concerned with the availability of service to their customers. Typically, revenue from customers is dependent upon service availability and, in general, how well the SLA between provider and customer has been met over a given period of time.

The lifecycle of an ICT Service (see Fugure 6) involves multiple phases which are shared between the customer and the supplier, as well as any subcontractors. To set up an ICT service for a customer, there are three main tasks that need to be accomplished: **Service Requirements Analysis**, **Service Definition**, and **Service Provision Resource Estimation**. Once these pre-sales negotiation phases have been completed to the satisfaction of both provider and customer, then the **Provisioning, Installation and Deployment** phases can commence, followed by **Acceptance and Commissioning/Handover.** At that stage, the service enters the **Operations** phase, during which **Back-ups, Update and Maintenance** activities will need to be prioritised, scheduled and preformed. At the end of its term, the service will undergo **Decommissioning/Handover**, where all data and materials are returned to their respective owners.

In outline, these phases are:

- **Service Requirements Analysis:** This phase involves the identification and characterisation of business objectives and their supporting business processes. Ultimately value must be placed upon both the functional (correct operation) and the dynamical properties of these processes. Candidate services can then be proposed and valued, and tests made against the relative merits (primarily cost and risk) of retaining such services in-house or outsourcing them.
- **Service Definition**: This phase involves describing and characterising the service required by the customer and the extent to which it can be provided. This is formalised in terms of a legal contract, consisting of a number of SLAs (Service Level Agreements).

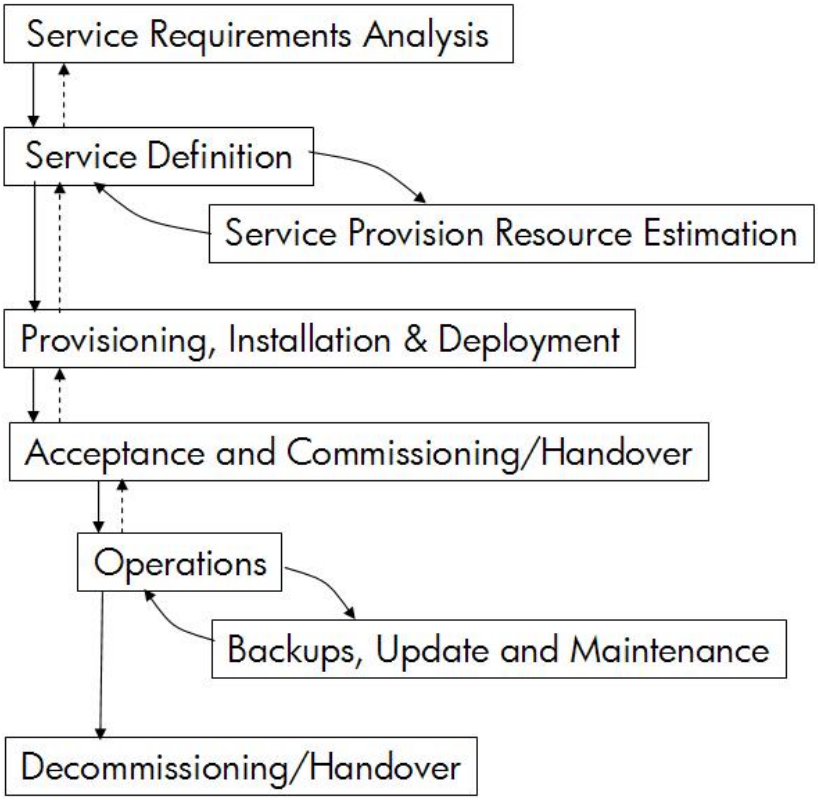

Figure 6: ICT Service Lifecycle

Typically, these SLAs are each expressed in terms of a number of SLOs (Service Level Objectives) defining the acceptable variation of measurable characteristics such as availability of network connections, availability of processor capacity, network bandwidth, and so on. Importantly, these characteristics must be phrased in terms of continuous random variables that can be statistically observed and effectively monitored in a secure, well-defined manner. Typically, the exact limits are defined and agreed in concert with the Resource Estimation activity described below.

Besides standard SLOs describing the normal operational characteristics of the service, the SLAs will also define and specify policy and procedures concerning Change control management, Business continuity and disaster recovery processes, and Systems patch management and software update policy. Many of these aspects are covered by standard operating procedures, and may be based upon emerging international standards for It Services Management, such as BS 15000.

Finally, the SLAs need to specify what non-compliance means for particular SLOs and the penalty charges that may be incurred as a consequence. Typically, non-compliance of a SLO means failure to achieve measurable performance targets for some minimum period of time. The penalty charges that are incurred will typically be proportionate to the measured degree of non-compliance and its duration. On the other hand, some SLOs may incur bonus payments based upon success in over-achieving base targets.

- **Service Provision Resource Estimation**: This phase captures and analyses the resource requirements for a proposed service and examines their consequences in terms of provisioning cost and economic viability. Clearly, this is a crucial task and will be interleaved with the Service Definition phase. The analysis will involve mathematical statistical modelling, plus some degree of systems simulation.

  The main source of complication here is in providing accurate estimates of the resources needed to provide degrees of *fault resilience/tolerance* against individual systems failure, for a given

level of service availability and cost. The kinds of failure envisaged can involve isolated faults in server hardware and networking devices – once the size of service provision required becomes appreciable, the Mean Time To Failure for the ensemble of physical systems involved will be severely reduced, making failure management a timely necessity.

For business-critical systems, a given level of fault resilience is only a basic requirement. Frequently, such systems must be made *disaster resilient*, involving multiple backup systems, having diverse construction and implementation. The overall systems design must then allow for a graceful switch over, back and forth, load balancing between these systems. One strategy for this is to massively over-provision the hardware systems – but there are several problems with this approach, not least of which is the economic cost involved. There is also a risk that merely over-provisioning will simply suffer from common mode failures, unless diverse solutions are also incorporated. An analytical systems modelling approach can therefore help to maximise the potential fault and disaster resilience for a given systems architecture and cost.

- **Provisioning, Installation and Deployment:** This phase involves the *construction* of the IT service itself and the general acquisition of resources. This may involve modifications to existing equipment and software or even purchasing of additional equipment and software. By the end of this phase, all the materials required have been obtained and the services constructed ready for operations. Naturally, DBSy-style models might be used at this stage in describing the necessary security requirements that permit customer interaction with the service and its management during operations.

  Traditionally, this is a risky and costly phase for service customers and providers alike. As a result, this is the focus of research investment into services automation to realise dynamic systems deployment and resource re/allocation, by taking a utility-based approach to computing provision.

- **Acceptance and Commissioning/Handover:** This phase involves trialling and testing the service, to achieve accreditation and (initial) acceptance of the service by the customer. This phase also involves ensuring a smooth transition into full services operation – either from a fresh start or as a handover from another existing service provider. This process will typically include the hand-over and transition from any existing legacy service. DBSy-style models may be helpful here in identifying the additional security protection that would need to be provided.

- **Operations:** This phase is, of course, the main activity that all the previous phases have been leading up to. This involves ensuring that the customer's applications software, processing and networking requirements are all *fulfilled* as much in accordance with the agreed SLA as possible.

- **Back-ups, Update and Maintenance:** All of these activities are interleaved in parallel with the Operations phase and concerns the ever-present need to perform back-ups, systems updates and maintenance activities for hardware and the systems and applications software components.

  Backups are particularly problematical for IT service providers and their customers, since they capture operational service data and materials owned by the corporation. This capture needs to be performed on a regular, routine basis in case this data is needed to restore the service operationally. Because of the clear security risks of maintaining long-term backups and so on, the service may have to be specifically designed to only handle and process transient data, having a well-defined expiry time, after which the data would be worthless to an adversary. Such a strategy essentially involves eliminating unnecessary state data that would need to be restored from backup on system failure. By this means, the need for back-ups of the service customer's data can be severely reduced and even eliminated.

  For isolated systems and services, the risk of disrupting the ongoing operation of essential, high value systems due to speculative updates will outweigh the need to apply patches to correct known flaws and to upgrade existing interfaces – at least, for a time. On the other hand, applying these patches and upgrades can only be put off for a while – the time will eventually

come when the upgrades become urgent and necessary, as requirements on the service mature and change. Patches and upgrades are not a matter of 'if', but 'when'.

For systems and services that have direct operational requirements for external network services, the need for maintaining systems updates and doing preventative maintenance may be far more acute and risky than updating the back office systems. This is because the external-facing systems form the front-line of defence against external attacks.

- **Decommissioning/Handover:** This phase happens at the end of life of a service. The data processed and results produced (including any backups made) will typically belong to the service customer, as might the applications software itself, particularly if developed specially for that service. All of these will have to be securely returned to the service customer. Other materials will typically be owned by the service provider and will need to be reclaimed for use by the provider's future customers.

  This process clearly presents a security issue for *both* the service customer and provider. The service customer needs to have *all* their data and software returned, without it being captured by a third party en route or retained in some accessible form by the service provider. However, the service provider wishes to rapidly sever their connections with the now defunct, non-revenue earning contract, and reuse the newly reclaimed resources as soon as possible. Furthermore, the service provider will wish to avoid any liabilities to future customers, due to pollution or contamination of resources by prior customers.

  Thus, a decommissioning process will typically include some kind of secure extraction and transfer of the customer's data and software, followed immediately by a "data scrubbing" process that securely removes all data related to the expired service from the provider's own systems.

  Typically, the SLA will specify an agreed level of "data hygiene procedures" to be applied at either one or both of the commissioning and decommissioning stages for the service. DBSy-style models might be useful to help identify any additional security enforcement required to maintain protection during these critical transitional phases.

  In reality, many systems are never actually decommissioned. At the end of a contract, the service is either retained by the existing provider, handed over to a new provider or brought back in house. Agreements for all three cases (often involving the transfer of physical as well as data resources, and even personnel) will have to be established.

  It is important to understand that the linear process chain outlined above is a simplification of what is in reality a far more complex set of interactions;

  - **Few battle plans survives intact first contact with the enemy**: specifications will change, both as business requirements change and technology refreshes become available;
  - **Processes merge**: there will be significant feedback throughout the stages as more becomes learnt about the actual as opposed to stated user requirements and their implications for design and cost.

3.4. Emerging standards for IT Services Management

As the industrial market for ICT services provision matures, a number of international standards are emerging that define a baseline of what can be reasonably expected of an ICT service provider by their service customers. This consolidation helps to create the market conditions, based upon industry-wide established best practice, for a commercially profitable international trade in ICT service provision.

There are certain benefits of standardisation for customers and providers alike:

- Rationalisation of definitions, classifications and ontology for the Terms and Conditions used in formalising SLAs will help to simplify the task of defining service offerings in legal terms. This helps to reduce potential for misunderstanding and ambiguity. Accordingly, customers can

have much greater confidence in their expectations about the service to be delivered, as their requirements can be communicated in a clearer and more effective way to the service provider. By the same token, the service provider can be clearer about the extent and scope of what is provided to the customer. An additional benefit will be in reducing procurements costs both for the Service Provider and the Service Customer by reducing procurement lead times and team sizes and also may lead towards reduced need for intermediaries. However, standardisation is likely to lead to commoditisation of service offerings with benefits to Service Customers but presenting Service Providers with considerable operational challenges to meet lower cost targets as margins erode.

- Common frameworks for describing ICT services enables service customers to make meaningful comparisons between offerings from different service providers. Importantly, this allows service customers to optimise the service provision they require from a market-place of service providers. For example, interoperability between different ICT providers allows a service customer to spread their requirements amongst several providers, rather than rely solely upon one single contractor. From the service provider's point of view, the creation of a standards based market offers scope for competition and the creation of competitive advantage from the pricing and bundling of value-added service product offerings.
- Industry-led best practice standards and accreditation of service provider promote a more confident and assured take up of services by IT service customers. Service providers also gain a degree of legal protection by adopting an accepted publicly defined standard.

Briefly, standardisation helps form a market of service providers for customers to choose from, for a given level and type of ICT service. It creates a defined baseline above which service providers can compete in the marketplace and develop their own value-added service product offerings.

We briefly outline some of the current approaches to standardisation of IT Service Management (ITSM) and security management:

ITIL – IT Infrastructure Library was created by the Office of Government Commerce (formerly known as the CCTA), with the explicit intention of providing guidelines and standards for IT service provision for the UK government. Since then, it is being rapidly adopted world-wide as a de facto, comprehensive best-practice standard for ITSM. Many leading consulting and educational bodies offer world-wide ITIL training and certification programmes for IT professionals.

BS 15000 is the first world-wide standard aimed specifically at IT Service Management. It describes a set of management processes for the delivery of IT services to the corporation and its customers. BS 15000 is aligned with and complementary to the process-based approach defined within ITIL.

COBIT is another de-facto standard for IT Services Management and stands for *Control Objectives for Information and related Technology*. This was originally released as an IT process and control framework associating IT with business requirements, aligned with vertical industry sectors. It was initially used mainly by the assurance community in conjunction with business and IT process owners. COBIT has begun to be used more and more as a framework for IT governance, providing additional management tools such as metrics and maturity models to complement the control framework. COBIT is maintained and developed by the IT Governance Institute.

BSI – IT Baseline Protection is a general service offering IT security advice and information to commercial businesses. The BSI IT Baseline Protection software tool (GSTOOL) supports commercial users in preparing, administrating and updating IT security concepts that meet the requirements of IT Baseline Protection. After gathering the information required, users have a comprehensive reporting system at their disposal that can carry out structure analyses on their compiled data and so generate reports in paper or electronic form. It is provided by the BSI, the Federal Office of Information Security which is the central IT security service provider to the German Government.

3.5. Issues and Challenges for IT Services Management

We briefly outline some emerging issues and challenges for IT Service Management. These are specifically concerned with how to maintain effective control in the presence of an ever-changing, complex business world.

Taming Organisational Complexity – Cutting through Silos

Both service customers and service providers tend to be large and complex organisations in their own right. Both organisations will have their own internal business procedures and priorities for action. The SLA defines what the service is, what is expected of it and what penalties may be incurred for non-compliance. The service provider will have a standard set of management processes for Customer Service that report incidents and problems, escalate them and hopefully, obtain solutions and take appropriate action accordingly.

Within these complex organisations, several independent departments may need to act collaboratively to provide the IT service for each of their respective customers. Given that these services are delivered according to SLAs, complex organisational structures will tend to get in the way and will make delivery harder to achieve. In particular, different parts of the organisation will have expectations of what the others are providing to them – and what they are expecting in return. It is therefore important to obtain a degree of consistency and shared understanding of these individual objectives. To this end, organisations can and do implement OLAs (**Operational Level Agreements**) that in effect provide internal SLAs between departments.

The main reason for this complexity is that simple economics dictates that service providers must consolidate their capabilities and share resources amongst their various service contracts as far as possible. It would simply be uneconomic to try and provide each set of services to the appropriate standard in a completely separate, independent set-up for each contract and customer. Thus resources will be aggregated and shared amongst service customers.

Because there may be onward connections from service customer to several service providers and back again to other service customers, and so on, this means that there will be indirect, unintended connections between competing customers and between competing providers, (see Figure 7). Future DBSy-style models could potentially help to identify where the organisational boundaries are and the necessary controls on a shared network architecture that serves and connects multiple service customers and their providers together. This implies that management of parts of the network is also shared with others and thus responsibilities should be identified, again in terms of appropriate SLAs within the consortium.

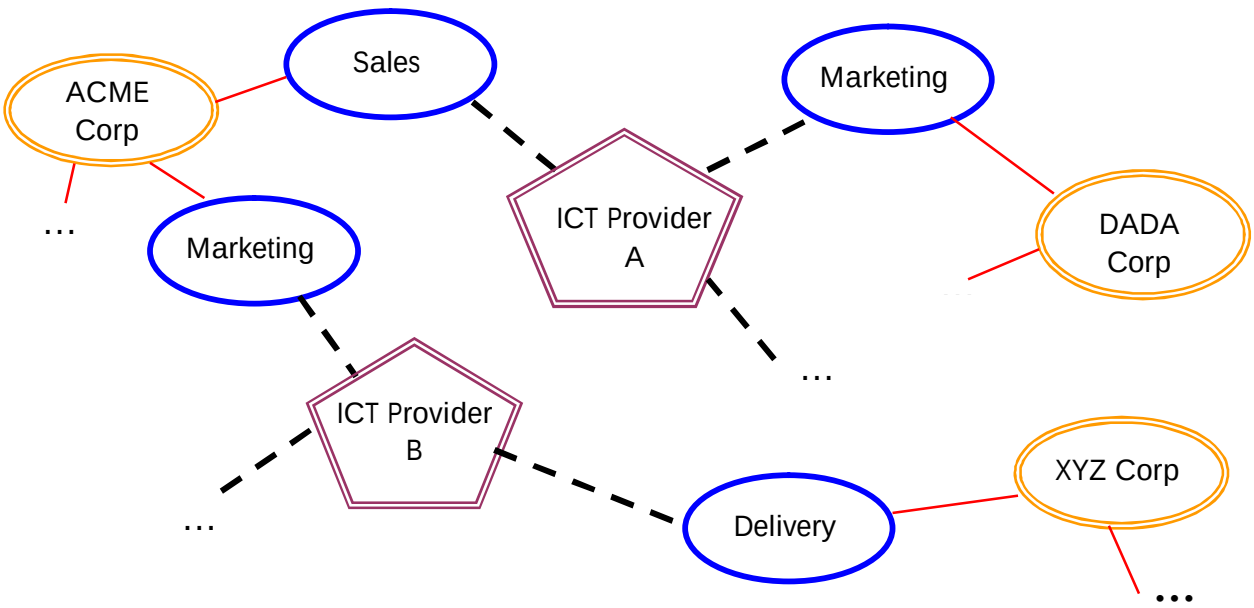

Figure 7: Indirect connections between customers and providers alike …

Corporate Accountability – Business Command & Control

With recent legislation (BASEL II, Sarbanes-Oxley, HIPAA) in the US and EU mandating greater attentiveness to IT governance, large organisations have an urgent need to provide a comprehensive command & control infrastructure that is tightly integrated with business objectives. These systems will need to keep each layer of management up-to-date and well informed about their portfolio of concerns.

However, irrespective of IT governance issues, there is a commercial benefit to having tighter command and control of the business. Future DBSy-style models may help provide ways to visualise, to map out and comprehend how, where and why information flows within the business. This may help reveal all kinds of ways of optimising business performance through more effective and secure communication.

How to measure realistic and effective SLOs for Information Security?

Given that ICT services are increasingly being driven by SLAs that try to capture business objectives and goals, how can realistic and effective SLOs be defined that are both measurable and encompass non-functional characteristics such as Information Security? The same question applies to any set of non-functional attributes related to the *quality* of what is delivered. It may be extremely hard to specify measurable attributes that effectively *parameterise* the intended content in a meaningful and scientifically credible manner.

As a brief example of an SLA purporting to be security-related, one might specify that there should be sufficient network bandwidth to download up-to-date anti-virus signature definitions in such-and-such a number of seconds to each bastion host on such-and-such a network. As an isolated SLA, this is certainly measurable and, at some level, worth having. However, it says nothing about when these definitions would be applied, or what they should be applied to. Are these definitions used in real-time to filter packet traffic and/or to scan file systems? What would happen if this particular SLA was violated? And so on.

In short, isolated performance-oriented SLAs like the one above can say little or nothing about the overall *effectiveness* of defensive measures.

As is widely appreciated, the effective measurement of security is problematic. In an ideal world, we would like to firstly have an objective assessment of all the threats we care about, followed by an effective measurement of their likely impact upon the business, and finally have a measure of their actual impact, all over a suitable time period. Using such data, we could then say how much security has achieved by comparing the actual impact encountered operationally with the likely impact of the threats identified. If the actual impact is low despite a high potential impact, then one can argue that security is effective in preventing disruption. If there is lots of actual impact that was either completely unanticipated or anticipated and not prevented, then one can say that security is not effective at preventing disruption. Notice that security is only meaningful once the sources and kinds of threat and potential attack have been identified. Unfortunately, this ideal approach is not particularly practical since it involves getting information that is fundamentally unknowable e.g. all potential threats of interest. Another, more pragmatic approach is therefore needed.

To improve on this situation, the approach suggested here is to identify several measurable aspects of the particular infrastructure and service provision that can be *correlated* with defensive effectiveness, *as an ensemble*. The idea is to establish a set of coordinated defensive controls and sub-systems (i.e. defence in depth) that are *designed* and *configured* to address the pertinent security concerns, whatever they may be. The next stage is then to identify an ensemble collection of SLAs that for example ensure that the defensive control is available and performing in a well-configured, desired manner. Although this cannot directly "measure" security as such, the statistical evidence gathered gives a *correlated indication* that the security systems are, at least, functioning as expected.

In this way, considerations of qualities like security can be effectively translated into a combination of *design integrity* concerns on the one hand and *active demonstrations* of measurable performance

on the other. To be sure, such an ensemble of SLAs will be strongly dependent upon architectural and structural security features specific to each ICT service. Thus, future DBSy-style models could be useful here in two ways:

1. For identifying and defining what the relevant defensive controls would need to be.
2. In helping to characterise meaningful, measurable attributes of these controls that can, as an ensemble, be effectively correlated with the level of security protection required.

4. Security Requirements and Risk Analysis for ICT Services

With each year that passes, ICT services and the way they are provisioned is becoming more and more specialised, interdependent and complex. This means that ICT services will increasingly be placed in the hands of organisations having the core competence and the specialist skills necessary to provide and operate them. As a result, we can expect increasing use of measurable SLAs to explicitly state and govern what providers deliver and what customers can expect for their money.

This trend is set to increase, irrespective of the prevailing fashion towards contracting in or contracting out. How each ICT service is actually provided – whether it is through an internal corporate department, via an external consortium of suppliers or some combination thereof – will be determined by the prevailing needs of the customer's business.

In this paper we have seen that operating and managing ICT services mostly involves creating and following business processes (e.g. change management). Security and ICT services are each concerned with people, the roles they adopt, the technology they use and the business processes that connects them. This means modelling the business processes, roles, actions, capabilities, effects, responsibilities and duties. The purpose of this section is to summarise and discuss various issues arising earlier in the paper that could help when formulating approaches to security requirements and risk analysis, applicable to ICT services as defined by and managed according to SLAs.

In summary, these issues are:

- **The unique ownership & management assumption** for the infrastructure components contained within Islands and Causeways is a serious barrier. Each ICT service, however it is provided, needs to be operated and managed in accordance with specific SLAs between provider and customer. The physical situation, the networking and the necessary separations involved will require a more subtle representation to accurately capture the various shared responsibilities involved. In particular, this sharing will be subject to relevant SLAs and thus the representation must also make reference to them.

  One may argue that the main question facing us here concerns how to map a number of different information views together in a coordinated manner to yield a synthetic view for the purpose of control. For example, aspects of physical location and physical access forms one such view, business functions and their network requirements forms another view and the logical interconnection that aggregates the various sub-networks and the security controls forms yet another view. Using these basic views, a number of compound, derived views could then be constructed from some selection of these elements to capture pertinent aspects of how the business is operating. For example, the use of ICT services challenges and breaks the conventional assumption that such services are necessarily *hosted* as a part of the organisation. Thus, one might imagine a "hosting" mapping that associates each service with its set of SLAs and the current provider. Consequently, it is this "hosting" mapping and its consequences that, fundamentally, the customer needs to maintain control over.

- There is a need for an openly available standard approach to systems security assessment that can be combined with Compromise Path Analysis. The security assessment approach used currently in conjunction with DBSy is linked strongly to internal MoD requirements, policies and procedures (i.e. JSP 440 and MPS).

  The Compromise Path Analysis technique involves digging down into the infrastructure and characterising the strength of security controls to meet risks on each of the access paths. The

process of risk assessment considers if these controls are acceptable according to the "risk appetite" of the risk owner. If not, then further refinements to the controls may be called for – or a change to the business processes involved will be necessary. To do this, the technology dependencies associated with security controls need to be made a lot more explicit.

- Management channels for business-level command & control of ICT service functions have to be made explicit, with a clear statement of scope and functional responsibility. When we need to break barriers and to cross corporate boundaries in providing IT systems, the management responsibilities and how they maintain operational effectiveness needs to be made explicit. Typically, these channels are also part of a shared, common infrastructure that requires coordination and hence management. Thus, the management command & control channels will themselves need to be highly protected.

  This shared coordination of managerial responsibility serves to encourage and enhance mutual stewardship of goals by both providers and customers. Making the management channels explicit can have the beneficial effect of increasing managerial transparency and accountability to all the participants involved. Actions taken by each participant can be recorded in an auditable and computationally hard-to-forge manner, providing a trustable record of activity. This can be used to provably show that certain actions were taken in a timely way [iTrust05].

  The overall effect of active, continual record keeping is to make urgent causes of disruption more visible and obvious, hopefully leading to more accurate and timely diagnosis of the set of most urgent issues needing to be resolved. Problems and issues will invariably arise during operations – the main question is how swiftly they can be resolved to reduce disruptive impact and consequential financial losses for all concerned.

- Because management of ICT services is very process-focused, a finer-grained mapping between the people roles, the systems technology and the business processes is necessary when analysing and managing ICT services. In particular, a deeper appreciation of what the relevent business protocols and processes are and how information flows between participants having particular roles is required. Approaches to BPM (**Business Process Management**) that help to identify, organise and define business processes and their activity have been developed – for example, see [BPMI, Ould05, BIS99].

- Developments in communications technology are forcing radical challenges to security technologies. Depending solely upon well-configured gateways and firewalls is no longer adequate for network and service defence. In particular, coordination of security measures at both the network and systems level on the one hand and the application level on the other is increasingly required.

  Virtual groupings utilising mobile phones, e-mail and instant messaging technologies are entirely commonplace within all manner of commercial organisations. Organisational structures and their management are being represented entirely in terms of these virtual groupings and other forms of dynamically generated access control list. To remain relevant, it is clear that modelling techniques need to be extended to include these application level concepts, in addition to characterising what they enable in terms of business process and associated technology.

  To provide joined-up management of these systems, drill-down to the more detailed technological aspects involving networking and systems must be provided. In practice, this may involve creating, identifying and/or integrating with explicit management interfaces to existing legacy network and systems management systems.

  Security depends critically upon end-to-end integrity and quality – all the way from organisational structure, business processes and protocols, services architecture, networking and security protocols, operating systems and their internal systems architectures, trusted device technology and all of the version and configuration control that this implies.

- Security classification of information and the roles of those entities that handle such information are dependent upon business process and what the information is used for. This means that the security classifications on bodies of information could depend on timings of real-world events that will occur or have already occurred, on what processing was done to yield or generate that information as output, and on what processing will be performed using that information as input. In reality, security classifications on information are therefore time-dependent and processing dependent – and will need adequate representation.

4.1 Commercial drivers and benefits

We now briefly turn to the question of the overall commercial drivers and benefits of a DBSy-style modelling. These include:

- Understanding the organisation better and how the ICT infrastructure is deployed and managed to the overall benefit of the business. Such information can help to optimise communications and identifying strengths and weaknesses in the communications infrastructure.
- Identifying the critical ICT services and to determine the consequences of decisions concerning how they should be provided – in particular, examining the security consequences of either contracting in or contracting out.
- As a service provider, DBSy-style models of relevant pieces of your ICT architecture could be exposed in a useful form to service customers, existing and potential. These models can reinforce the value proposition around security by showing where the controls and their management interfaces are. This helps assure customers that their information security risks are capable of mitigation on their behalf by the service provider.
- Compliance to IT governance regulatory requirements (e.g. Sarbanes-Oxley, HIPAA) involves presenting evidence that the business is well-run and efficiently managed. In particular, any events or situations that could adversely impact shareholder value have to be reported to appropriate audiences in a timely fashion. DBSy-style models could be a powerful source of evidence showing how the network infrastructure is well-managed and that information security risks can be appropriately mitigated.

Conclusions

DBSy provides an approach to security requirements and risk analysis for **network separation, services aggregation** and **compartmentalisation-in-the-large** that can highlight some of the responsibilities and risks of connecting different component networks together. Providing ICT systems and services is mostly about the managing of processes and the execution of workflows; ITIL/ITSM provides templates and criteria for workflow processes and represents current best practice for IT Systems Management. As seen in the rest of the paper, DBSy helps identify the boundaries and interfaces of ICT services, and the consequences of their separation. However, the unique ownership & management assumption for infrastructure does limit DBSy's application to those situations where this holds.

The difficulty to adequately model business processes and infrastructures is a major challenge for the effective use of modelling frameworks like DBSy in commercial environments. But unrelated changes may help improve the situation. An increasing number of corporations are investing now to streamline business processes and simplify infrastructures in order to reduce costs and be ready to leverage changes in business opportunities. Driven by competitive pressures, software vendors are continually changing the structure of their application packages, often to embrace the principles of the service-oriented architecture. Providers of management software introduce applications designed to manage business processes and infrastructure based on models. Corporations define and implement enterprise architectures, and specify their business processes more formally. These changes all combine to ease the task of building the base models used by DBSy-like modelling frameworks, which significantly improves the affordability and applicability of the methodology for commercial environments.

Acknowledgements

This paper would not be the paper it is without several direct contributions from HP Labs colleagues: Adrian Baldwin, Chris I Dalton, Frederic Gittler, Patrick Goldsack, William Horne, Simon Shiu, Richard Taylor, Chris Tofts and Mike Yearworth.

Brian is also very grateful to Abdel Boumakoul, Chris I Dalton, Frederic Gittler, Keith Harrison, William Horne, Simon Shiu, Mike Yearworth and Stefek Zaba for their helpful and constructive comments on various drafts of this paper.

References

[BPMI] Business Process Management Initiative, http://www.bpmi.org/

[BSI] BSI - IT Baseline Protection - http://www.bsi.bund.de/english/publications

[BIS99] Warboys, B., Kawalek, P., Robertson, I., Greenwood, M., *Business Information Systems – A process approach*, McGraw Hill, 1999.

[DBSy1] Domain Based Security White Paper, QinetiQ 2004 (http://www.qinetiq.com/dbsy)

[DBSy2] C.L. Robinson and K.J.Hughes, *Managing Infosec Risk in Complex Projects*, 4th Annual Systems Engineering for Defence Conference, RMCS Shrivenham, 15-16th February 2001

[DBSy3] C. L. Robinson, *Security Requirements Models to Support the Accreditation Process*, 2nd Annual Sunningdale Accreditor's Conference, 10th – 11th September 2001

[EDS] ELECTRONIC DATA SYSTEMS CORPORATION. SECURITIES AND EXCHANGE COMMISSION Washington, DC 20549. FORM 10-K ANNUAL REPORT PURSUANT TO SECTION 13 OR 15(d) OF THE SECURITIES EXCHANGE ACT OF 1934 For the fiscal year ended December 31, 2004 Commission File No. 01-11779.

[FT04] *Blackout puts spotlight back on IT failures*. Financial Times. MAIJA PESOLA and NICHOLAS TIMMINS, 27 November 2004

[iTrust05] Baldwin, A., Beres, Y., Plaquin, D., and Shiu, S *Trust Record: High-level assurance and compliance* in LNCS Vol 3477, Trust Management: Third International Conference, iTrust 2005, Paris, France, May 23-26, 2005. Proceedings Editors: Peter Herrmann, Valérie Issarny, Simon Shiu

[Mon03] Monahan, B, *From Security Protocols to Systems Security*, in Proc. of 11th International Cambridge Workshop on Security Protocols (2003), LNCS, Springer, 2003 (Extended version: http://www.hpl.hp.com/techreports/2003/HPL-2003-147.html)

[Mon05] Monahan, B., Infrastructure Security Modelling for Utility Computing, HP Labs Technical Report, 2005, http://www.hpl.hp.com/techreports/2005/HPL-2005-4.html

[Ould05] Ould, M. A., *Business Process Management – A rigorous approach*, British Computer Society, 2005.

[SoftUDC] Kallhalla, M., Uysal, M., Swaminathan, R., Wray, M., Christian, T., Edwards, N., Dalton, C.I., Gittler, F., *SoftUDC: A Software-Based Data Center for Utility Computing*, IEEE Computer, pp 38 – 46, November 2004

[TL01] Choo, T.H., Dalton, C.I., *An operating system approach to securing e-services*, CACM, Vol 44, Issue 2, 2001

Appendix: ICT Services related research at HP Labs

We take this opportunity to present an overview of some of the ICT services-related research and development that is underway within HP Labs.

Open Analytics: Understanding the value of Service Level Agreements

Richard Taylor, Chris Tofts & Mike Yearworth

Although most of the public never get to hear about the details of the bidding process for large IT outsourcing deals the press is full of IT system procurement failure stories post contract award [FT04]. When these failures arise the customer does not get the service required for their business with impact on operations and consequent financial loses whilst the vendor loses profitability on the deal due to penalties. For example, see the EDS failure to meet delivery performance requirements for the Navy Marine Corps Intranet (NMCI) system and the subsequent cash impairment charge of $375 million in Q3 of 2004 and operating losses of $487 million [EDS].

In a rational negotiation, customers cannot ask for an unrealistic performance specification and then expect the delivered and managed service to be inexpensive; yet stories of system failures show that this mismatch of expectation and delivery occurs frequently. Customer's business requirements are expressed to a vendor as a set of business objectives which are then refined into SLOs (**Service Level Objectives**). Such SLOs form an expression of measurable attributes of the service which have a direct bearing on the business objectives of the organisation and will be governed by the SLAs (**Service Level Agreements**) that form the contract. Contracts are thus intended to clearly state the obligations on the parties and the associated payments and penalties. Despite this apparently clear path from business objectives to information systems, in practice the pursuit process and contract negotiation often lead to poorly specified systems and services which do not match the real business requirements of the customer (headroom, availability, performance and agility, for example).

These problems have been caused by a failure of communication between customers and vendors. What every organisation is trying to do is value activities and their supporting processes. Value is ultimately monetary in most of these systems and value of both the functional and non functional (risk, security, transaction rates, response times) should then be used to assess appropriate, attainable and measurable SLOs and from there, SLAs with appropriate penalties - and importantly rewards[2]. However, business objectives are often not captured adequately by SLOs, and the result is that the procured information systems may not represent real value for the organisation. With the increasing need to demonstrate accountability for financial decisions (e.g. Sarbanes-Oxley) the capability to demonstrate that, for example, a performance/cost trade-off analysis was performed for all IT system procurement, could become essential.

HP is advocating the use of analytical techniques based on mathematical modelling of systems and equipping its services businesses with the capability to analyse customer's business goals and their link to SLOs. This can greatly help to communicate the business critical tradeoffs that need to be made between the performance, availability and cost of information systems. This consequentially helps avoid bidding for and/or building solutions that are either over-engineered and non-competitive (too costly), or under-engineered and unlikely to meet the SLAs with consequent financial penalty.

[2] There are increasing numbers of SLAs with substantive reward mechanisms for exceeding service levels

SmartFrog

Patrick Goldsack

SMARTFROG (the *Smart Fr*amework for *O*bject *G*roups) is a general-purpose framework for deploying and managing software components based on a specification of their configuration. SmartFrog can be applied to multiple problem domains, though recently we have focused our efforts on the use of SmartFrog to instantiate complex services and applications on utility infrastructures.

The realization of utility computing requires that we are able to rapidly and repeatedly repurpose utility infrastructure to offer different services and applications. A necessary condition is therefore that we are able to instantiate services:

- *automatically* – and hence repeatably
- *flexibly* so that we can configure the service easily for different demands
- *correctly*, so that there is some checking as to whether we are instantiating a service with a correct configuration
- *securely*, so that automatic service instantiation does not introduce new security deficiencies.

Once a service is deployed, we want to be able to manage it through its lifecycle. In particular, we want to enable adaptive behavior, so that our automatically created services can change their configuration to accommodate changing circumstances such as workload variation and failures. At the end of the lifecycle of a service instance, we wish to cleanly remove it from our utility resources so they can be repurposed.

The approach adopted by SmartFrog is to think of services as composed of distributed components that must collaborate to deliver the complete service. We keep the configuration details of each component as separate as possible from the functionality of the component. This means that components can be as general-purpose as possible, and can adopt different, *configuration-driven* behavior when deployed. To express the configuration details of individual components, and groups of components, mechanisms are provided to facilitate highly flexible description, manipulation and composition of configuration data. It is this configuration data that specifies the behavior required from the group of components that comprise the service. The configuration data is used at run-time to orchestrate the activities of component groups to deliver correctly configured, running services. We refer to this as creating *configuration-driven systems*.

The SmartFrog framework consists of three major elements: the configuration language, the runtime deployment engine, and the set of components to be deployed. SmartFrog is now readily available as an open source project – see http://www.smartfrog.org.

Model-Based Assurance – The Trust Record

Adrian Baldwin

A critical factor in ensuring that corporate IT systems are functioning correctly, securely and supporting the business is to ensure there are appropriate controls in place to ensure that corporate policies are met. In this case, a *control* (to use auditing terminology) is a mechanism that ensures that appropriate actions are undertaken in running an IT system. For example, a control may be the process via which a user gains an account which ensures that appropriate users can gain access in a timely manner, whilst others cannot access these accounts.

Many IT environments are audited against a number of such controls to ensure that they are operated in an appropriate manner. An auditor will first check that appropriate controls are in place to manage corporate risk in an appropriate manner and then they sample various controls, often around account management and change control, to check that each is running correctly. Any failures will be investigated further and lead to recommendations.

The correct and secure operation of a system is highly dependant on having the appropriate set of controls that are correctly run but where occasional audits are often insufficient to maintain these controls. Recent legislation such as Sarbanes Oxley stresses the importance of having appropriate controls that encourage a more continuous monitoring environment. Such an environment could consist of the collection of a set of system statistics that are fed into the audit department to help them focus there investigation. The *Trust Record* **[iTrust05]** is a continuous monitoring, analysis and reporting system that collects and analyses system data from key controls or that indicate key controls are being achieved and then allows a compliance or governance model to be built to report on how well the controls are working and their relationships to the business systems, risk impacts and corporate policies.

Secure Virtualisation of systems services

Chris I Dalton

Networking and systems ICT infrastructure today is hugely dependent upon hardware configuration and deployment, consisting of large numbers of independent and distributed servers, networks and storage devices. The prevailing issue is how can customers effectively bring all this complexity under their control in an effective way? Given that there is a continuing business need to reconfigure systems and maintain availability, a more strategic approach to providing this control has become necessary.

A radical and fundamental approach to this problem re-architects the hardware/software systems interface to exploit technologies that are inherently more manageable, more trustable and more configurable. A particularly promising example of that is *Secure Virtualisation* of systems services, such as operating systems and their network stack.

Secure virtualisation technology provides the ability to operate, for each server, several "sandbox" Virtual Machines that each runs a separate instance of the host operating system. These Virtual Machines provide strong compartmentalisation in the sense that there is controlled isolation of software applications running under each OS [TL01]. In particular, each OS can make independent progress, without affecting the others. This approach permits each server to efficiently host management applications in a distinguished Virtual Machine whose integrity can be trusted and assured. This technology also includes the ability to virtualise networking and storage components to make them more configurable and manageable.

Virtualised infrastructure offers the capability for richly mapping a logical architectural view onto the particular distributed resources that are currently available. This means it is easier and more convenient to replace faulty hardware and software components by remapping the logical view as necessary. From a security point of view, the virtualisation provides a security enforcement point where the view of the actual underlying hardware can be more tightly controlled. In particular, this technology makes it easier to provide a consistent, standardised view of the system which can more readily match the technical requirements of the applications software.

This technology has been implemented as a part of the SoftUDC project in HP Labs [SoftUDC].

Integrated Security Management and Security Modelling

William Horne & Brian Monahan

The fundamental issue with coordinating security across an enterprise is that multiple security mechanisms are deployed at different technology layers; each mechanism is painstakingly configured and maintained using legacy user interfaces, most likely by different administrators in different organizations and at different sites. This piecemeal approach makes security management labour-intensive and, therefore, expensive, error-prone and slow to adapt.

The goal of the Integrated Security Management (ISM) project is to address this problem. In ISM we formally model applications and networking components, and then reason about how these components compose. We then compare this composition against a centrally specified set of high-

level access policies. This allows us to automatically validate and configure device and application security mechanisms.

Other recent work in security modelling of enterprise systems considered how logic-based object modelling techniques may be used to help service providers and their customers obtain insight concerning the security characteristics of utility infrastructure and networked systems. In [Mon05], we briefly describe two modelling tool prototypes that were built and the underlying technology they used. In earlier work, [Mon03], we gave an outline proposal for a security modelling framework, called SSML, to help describe the way that security protocols are used within distributed systems, software applications and services.