Meaningful Security SLAs

Brian Monahan, Mike Yearworth
HP Laboratories
HPL-2005-218R1

Keyword(s):
service level agreements, security, analysis

Abstract:
Service Level Agreements (SLAs) are the de facto method of managing IT Outsourcing (ITO) contracts. Negotiated during pursuit (pre-sales) phase and then used as a dashboard for performance management during delivery the SLA ultimately becomes both the lever and the measurable for revenue and margin performance on a contract. That SLAs should be meaningful, both for customers and vendors as defined by some objective criteria, seems obvious but evidence from procurement failures for large IT systems suggests otherwise. As a consequence of bringing a rigorous and analytical approach to negotiating meaningful SLAs for ITO deals we have encountered on two occasions a customer requirement for a performance oriented security SLA that was not meaningful by our definition. This has inspired an investigation into the possibility of offering alternative security SLAs that we believe would be meaningful to both HP and customers with potential for improved operational visibility into the cost of delivery that also differentiates HP's offering.

External Posting Date: October 10, 2008 [Fulltext] - Approved for External Publication

Internal Posting Date: October 10, 2008, 2008 [Fulltext]

© Copyright 2008 Hewlett-Packard Development Company, L.P.

Meaningful Security SLAs

Brian Monahan, Mike Yearworth

Trusted Systems Lab
Hewlett-Packard Laboratories,
Filton Road, Bristol, UK

November 2005 (Revised 8th October 2008)[1]

Abstract

Service Level Agreements (SLAs) are the de facto method of managing IT Outsourcing (ITO) contracts. Negotiated during pursuit (pre-sales) phase and then used as a dashboard for performance management during delivery the SLA ultimately becomes both the lever and the measurable for revenue and margin performance on a contract. That SLAs should be meaningful, both for customers and vendors as defined by some objective criteria, seems obvious but evidence from procurement failures for large IT systems suggests otherwise. As a consequence of bringing a rigorous and analytical approach to negotiating meaningful SLAs for ITO deals we have encountered on two occasions a customer requirement for a performance oriented security SLA that was not meaningful by our definition. This has inspired an investigation into the possibility of offering alternative security SLAs that we believe would be meaningful to both HP and customers with potential for improved operational visibility into the cost of delivery that also differentiates HP's offering.

1. Introduction and Objectives

The performance of IT Outsourcing (ITO) contracts is typical measured against Service Level Agreements (SLAs) that were agreed at contract signing. The degree to which the SLAs provide a meaningful tool for delivery management is largely a question of whether there was any rational basis for their formation. Typically, HP confronts a number of situations:

1. The customer has no specific service level objectives in mind and HP offers a standard service with HP defined SLAs
2. The customer has specific service level objectives in mind but they may not be wholly rational and HP is able to negotiate to a standard SLA offering
3. The customer has a rational basis for demanding a certain enhanced level of service due to the nature of their business and HP needs to be cognisant of the implications of agreeing to these (often demanding) SLA terms
4. The customer has ceded SLA setting to a third party with strict compliance requirements and where negotiation may be impossible or difficult regardless of the rationality of the SLA

Of these cases clearly 1 and 2 present no particular difficulty to HP although achieving 2 may require an analytical approach e.g. HP Open Analytics. Case 3 is only problematic if HP has not taken into account the delivery implications. Again, an analytical approach will help understanding of potential pitfalls and may help in the price negotiation with the customer. For example, this might assist in explaining exactly why the HP price is higher than expected in order to achieve a certain level of service. Case 4 is exceedingly difficult to manage for HP although HP Open Analytics may help in challenging the credibility of the intermediary and is discussed later.

Formal model based approaches such as HP Open Analytics clearly have value in helping HP's ITO business in achieving desired outcomes as far as SLA formation is concerned. Given that there is usually a service credit model associated with meeting SLAs on a month by month basis the value can be expressed as protecting margin; both up front, in the sense of helping to achieve the best possible deal for HP to sign from an SLA point of view, and ongoing, in that analytics can be used to predict

[1] This report was only internally available within HP as HPL-2005-218.

deviations from SLA compliance and help model potential corrective actions. This in particular helps encourage HP pursuit teams to negotiate service deals that are operationally effective and profitable.

Another value associated with HP Open Analytics is less easy to measure but has been articulated by pursuit teams on various services opportunities for HP as creating a degree of confidence in the customer that HP has a sound understanding of the alignment of customers' business requirements with the information systems that it will be providing and managing for a number of years.

1.2. Enterprise ICT Services

The enterprise ICT services business model is well-known and has become commonplace within the IT industry. Even so, we review this to establish the particular terminology used here, within this report.

The starting point is that, due to the widespread availability of networking, it has become both technically possible and economically viable to consolidate and commoditise ICT services into particular bundles or offerings that are governed under a contract containing SLAs. This means that they can be operated on behalf of the Service Owner in a cost-effective manner by Service Providers. The Service Provider may or may not be a part of the service owner's organisation. From the Service Provider's point of view, the Service Owner is also their principal Customer for the service delivered, and is typically known as the Service Customer.

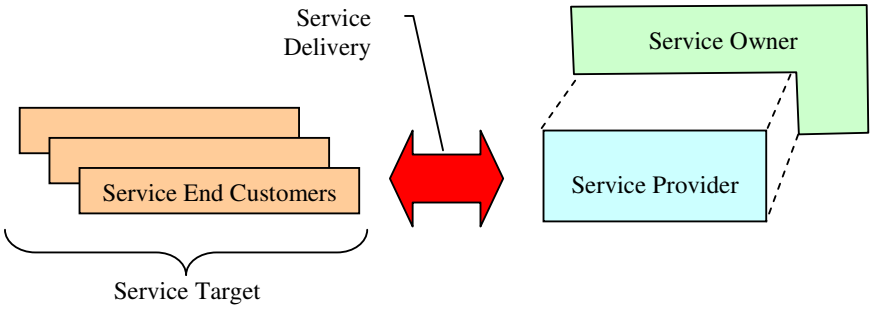

Figure 1 A simplified service model with external service end customers

It is assumed here that Service Delivery will typically involve networking technology, but may equally involve other kinds of underlying ICT service as well (e.g. database management, desktop management). It is most important to observe that the provider is in effect selling or trading a process in its own right. Service Delivery itself involves performing an overall *activity* or *process* by the provider on behalf of the intended recipients or Service Target: this may involve handling Service End-Customers, the customers of the service owner. These end-customers will have a service contract with the service owner, usually describing limited liabilities and so forth.

The contract or SLA we are more interested in expresses the agreement between the Provider and the Service Owner, defining the ICT services to be delivered to the Service Owner's End-Customers. The Service Provider is responsible for delivering the service on the owner's behalf, according to a number of service delivery performance targets, expressed as Service Level Objectives (SLOs). These SLOs are in turn defined in terms of physically measurable, statistically valid quantities known as Key Performance Indicators (KPIs). The idea is that the KPIs can statistically represent meaningful characteristics of service delivery and can be used to, for example, provide the basis for service billing as well as assessing how well the service is being delivered.

There will also typically be a service contract between service end customers and the service owner defining the expected service to be delivered by the Service Provider. This is depicted in Figure 2.

Of course, this diagram indicates only one of the possible arrangements between these participants. For example, the Service Provider may further subcontract some aspects of service delivery to other

providers. Additionally, the service owner may have more than one Service Provider, thus spreading the risk of a single point failure in service delivery by a single Service Provider.

Note that, from a legal point of view, the various contracts do not *compose* in a natural sense – and neither should we expect them to. Suppose that service end-customer Jane takes out a service contract with service owner, Acme Services Inc. In turn, Acme Services Inc. has an SLA with their service provider, called Quick-Serve Provider Ltd.

If, at some later stage, there is a legitimate breakdown of service for Jane, then she will have to contact the organisation she has the contract with, i.e. the service owner, to seek appropriate remedy or compensation. Now, it will most likely be that the service owner will have outsourced the customer services "helpdesk" function as well, meaning that Jane will in practice have to deal with this outsourced customer service agency to solve her problem. Because the helpdesk gets to interact with the Service Owner's end-customer base, the Service Owner will probably recruit the customer services helpdesk agency directly, rather than let the Service Provider take complete control of it.

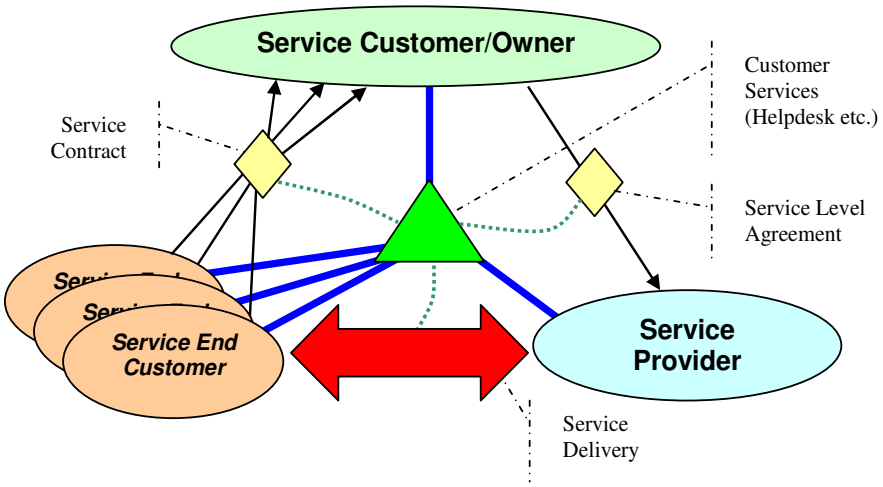

Figure 2 Complex service contract relationships between participants

In practice, the contractual situation is typically far more complex than this, as there will often be several Service Providers operating with the same Service Customer; equally, the Service Provider(s) may typically sub-contract to others for particular sub-services. Although none of these contracts will "compose" directly, there is a natural *chain of legal liability* in the sense that A may sue B who sues C who sues D – and so on. An important aim of service delivery is to minimise operational cost and risk, including legal liabilities. Hopefully, suitable *services engineering* can be put in place to mitigate these risks and liabilities.

1.2. Baseline Security Compliance Requirements for ICT Services

In the ICT service business, the baseline security compliance requirements are clear: each Service Provider has a number of customers to provide networking and compute for – what they use the provided networking and compute capabilities for is entirely their own business, and no-one else's. Therefore, the basic security requirements are for *separation* and *compartmentalisation*:

1. Protect and insulate the Service Provider from the activities of each customer and their end-customers.
2. Compartmentalise each Service Customer (and end-customers) from all the other Service Customers and *their* end-customers.

We can see this diagrammatically in Figure 3. Of course, the problem is that, although each Service Customer has similar needs and requirements, they aren't all completely identical, with varying degrees of *security risk* and *exposure* for all concerned. To help mitigate operational risks, some compliance certifications (as opposed to SLAs) are required. For example, the finance industry currently uses the

US auditing standard SAS 70 [SAS 70] for certifying that a data-center is well-run, This includes some security aspects such as provision of appropriate physical security (e.g. door locks etc.) as well as physical data resilience, such as raised flooring to protect against flooding

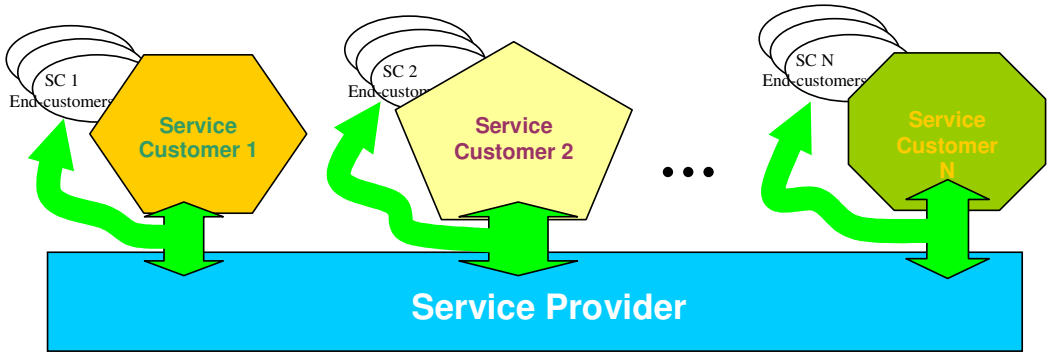

Figure 3 Baseline Security Requirements for ICT Services

A typical approach to implementing the baseline security requirement is to for each Service Customer to operate their Security Operations Centre[2] (SOC) whose function is to *co-ordinate* security-related functions with the Service Provider's own SOC (often combined with the Network Operations Centre or NOC). In this way, each Service Customer retains control of their internal business processes and communications and is also kept separate from the Service Provider's own business functions. We diagrammatically portray this in Figure 4.

When the customer is a Small to Medium Enterprise, the SOC functionality can then be minimal, perhaps offering a single flat sub-network, largely subsumed by a NOC. Alternatively, the SOC might be large and complex, and concerned with managing many different domains as sub-networks with multiple gateways and so on. Generally, the complexity of a customer's SOC will be largely dependent upon the size of the customer's organisation and the degree to which it requires networking for business purposes. There are interesting questions to be asked about the potential for customers to either operate a SOC for themselves or to attempt to outsource this functionality further

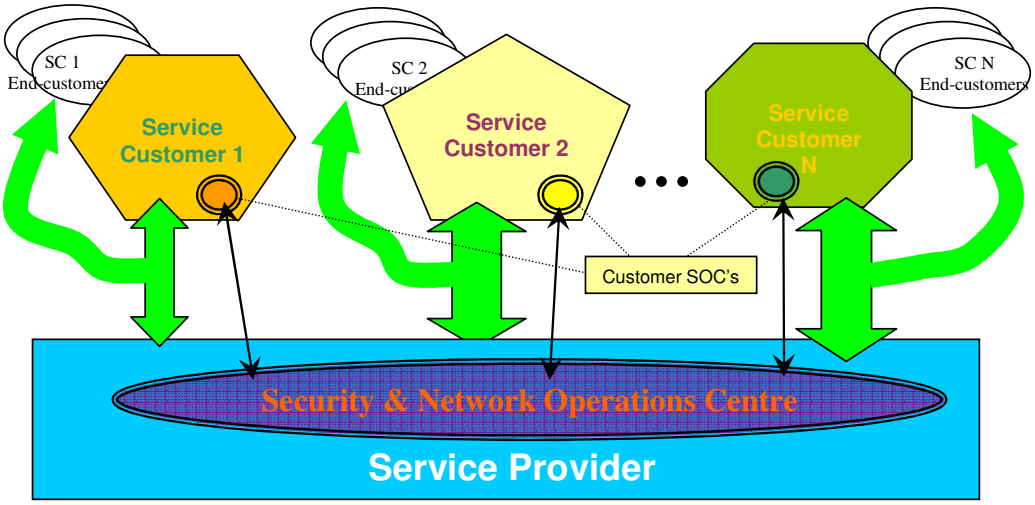

Figure 4 Using SOC's to implement the Baseline Security Requirement

[2] In the banking and financial services industry, the SOC may be incorporated as part of a Risk Operating Center (ROC) whose objective is to provide a strategic implementation of risk management for the business as a whole.

In addition to the SOC, the Service Provider will also control and maintain a number of Policy Enforcement Points (PEPs) on the Service Providers own network. These are configured and controlled by the Service Provider's SOC to provide appropriately compartmentalised networking for each of the Service Customers (e.g. VPNs). There are interesting risk assessment questions about how to securely support various kinds of systems access, such as access to the customers own machines which typically requires a "hole" in the compartmentalisation. Some of these issues are touched on in [DBSy].

2. Business Case

For most ITO contracts the SLA terms cover basic performance and availability requirements for the service elements that make up the overall contract. A typical example from a prospective contract would be:

> *"Resolve Priority 1 Incident – Response and Fix: 95% of Priority 1 Incidents resolved within 2 hours of the Receipt of Incident."*
>
> "*Personal Productivity UK Users: 99.95% Availability*"

These 'performant' measures should be contrasted with 'compliant' requirements which to first order can either be met; or not. This categorisation is based on standard RFP practice where vendors are requested to indicate yes/no for each compliance term, whereas SLAs are presented as minimum and expected level of performance with associated percentage revenue at risk

Customers' security requirements are usually expressed in terms of compliance and there is almost no evidence of any security performance SLA requirements from customers. However, there have been two HP pursuits where a simple and apparently trivial security SLA has been introduced – the distribution of virus definition files to devices within a certain time interval. There is suspicion that the inclusion of this SLA is a token effort; in our opinion the customer's representative has considered that there should be some security related SLA item but this example was the only one they thought they could readily measure and operate against.

Given that on these two occasions HP has been asked to meet a security SLA that in our analysis is not particularly meaningful for the customers' business, we propose the development and deployment of some well reasoned alternatives that would be of more *value* to both customers and HP. We can explore the value proposition through the following testable hypotheses

1. We will demonstrate HP's thought leadership in security and improve customer confidence in the HP solution.
2. We will have a negotiating position from which HP can attempt to avoid signing up to pointless security SLAs with consequent saving in associated measurement and operational costs.
3. We will have better visibility into security operations costs as a consequence of using the predictive performance models that support the new security SLAs.

Testing these will require a suitable customer engagement but clearly we also need to solve the immediate problem of deriving a candidate set of potential meaningful security SLAs and this is the primary focus of this paper. Core to solving this problem is the necessity of developing suitable predictive performance models for each of the candidate SLAs although this is beyond the scope of this paper and the subject of further work. We also need a useful working definition of *meaningful* is this context and this is explored in more depth in Section 0.

HP's Open Analytics focuses on using a model based analytical approach to answering these questions in a process framework that ideally leads to rational negotiation of SLAs in cases 2 and 3.

2.1. The effect of an intermediary broker for IT Outsourcing deals

Services has been engaged in a number of pursuits that have been brokered and negotiated under the management of intermediary broker's;

1. A senior executive, whether CIO, COO or IT manager, contemplating outsourcing all or part of a company's IT function needs powerful support from credible advisors in order to sell to the board,
2. The complexity of assessing disparate bids from many different vendors can be reduced by a strong intermediary imposing standard RFP templates on vendors' service definitions to better enable a customer to make like-for-like comparison on price, and
3. Profusion of IT procurement failures leading to realisation that such projects are high risk.

Essentially, customers outsource their corporate IT because of their relative lack of maturity in understanding IT issues which are not core to their business. However, even when one outsources IT, the business impact of that IT still has to be grasped, understood and eventually negotiated with IT providers.

2.2. Ingredients of an SLA

Each SLA defines a service to be delivered, as part of the overall legal contract between the Service Customer and the Service Provider. It obviously needs to define financial attributes of the transactions between Customer and Provider e.g. price for delivered service, penalties for non-compliance, and rewards for exceeding expectations. Overall, the SLA should say:

- What the customer *expects* to receive …
  - As long as the provider behaves well, by delivering according to schedule, not being disruptive, etc.
- What the provider *promises* to deliver …
  - As long as the customer behaves well, by paying their bills promptly, keeping to the by-laws, etc.
- How the service processes are to be *measured*:
  - Need to specify *observable* Key Performance Indicators (KPI's)
- Penalties (for non-compliance, by either side)
- Rewards (for exceeding expectations, by either side).

We can see that, although an SLA constitutes both a legal and financial instrument, it has also to serve as a *process* description. An interesting subtlety is that SLAs typically constitute a single statement of what the service is and what each participant has to do. This means that the same SLA statement will be interpreted by the Customer and Service Provider from their own point of views. Clearly, such documents need to be carefully drafted to avoid ambiguities and opportunities for creative misunderstanding and malicious compliance.

2.3. Assessing SLAs – Defining what makes a SLA meaningful

For the purposes of this document, we shall define and present "Meaningful SLAs" in the following terms. We first state what the SLA is in natural language (typically presented inside a box), followed by these subsections:

Description: This gives a statement of the *intent* behind the SLA and discusses what the various *expectations* might be from the point of view of each of the participants – Service Provider and Customer.

Technical Development & Engineering: The purpose of this section is to outline some of the issues concerning the technicalities of development and engineering. In particular, this should broadly say what options the Service Provider has to deliver the service. Each of these options should define the various technological and engineering approaches to be taken and the assumptions that each makes and relies upon.

Analysis: The purpose of this section is to throw some light upon the various options available to the Service Provider with the objective of trying to rationally predict performance characteristics. Based upon these predictions, it is then possible to make better informed decisions concerning the SLA – for example, whether to reject it and renegotiate, or to accept it. An SLA should only be accepted if it provides value to the Customer, a good profit margin for the Service Provider, with an acceptable downside risk.

The analysis proceeds by answering questions like those below:

Value to the customer: What is the value of the service being offered to the Service Customer? Does the SLA capture something that customers can directly appreciate, understand and value in terms of their own business?

Predictable: The consequences of honestly meeting the SLA for the Service Provider and Customer should be predictable. In particular, the Provider needs to investigate what the engineering implications are to meet this SLA. The Provider should also determine an *estimate* of the *operational risk* – that is, the impact of **not** meeting the SLA, and how often this might happen in practice. To do this, an investigation is needed of how the assumptions for delivering the service could fail (i.e. **service delivery threat analysis**).

Based upon this, we can determine which options (if any) have acceptable downside risk for the Service Provider. In this way, we can try to determine the *impact* of a failure to meet an SLA on the Customer's and Service Provider's business.

Measurable: How can the Service Provider demonstrate to the Customer an accurate and true statement of progress and the current status? What statistically significant values need to be routinely and regularly captured? What Key Performance Indicators need to be evaluated and assessed?

Affordable: Can the Service Provider and Customer afford to sustain a service at the levels proposed? At what level can the service be operated in a cost-effective way for both the Service Customer and Provider? Will meeting this SLA be profitable to the Service Provider? What about the downside risk to the Service Provider?

Understandable and Unambiguous: This means that each of the Service Provider and the Customer can arrive at a common, shared account of what the contract effectively entails. In particular, the contract terms should be clearly stated and not be subject to differing interpretations.

In operational terms, it should be possible (at least theoretically) to put both the Provider and Customer in separate rooms and ask them to produce independent accounts of what the contract entails operationally, including penalties. When these accounts are later compared, there should be a clear correspondence between the contract clauses that each have produced. If this is so, then an agreed shared understanding has been achieved – if not, then further negotiation effort is needed to bridge the gap.

In effect, this means that both participants have an appreciation of what the *other* participant needs to do and expects to happen.

2.4. Some characteristics of well-posed SLAs

Briefly, here are some general characteristics of a well-posed SLA. An SLA described in these ways is more likely to be meaningful and sensible to both sides, although it should always be tested:

- SLAs should be expressed in terms of observable, measurable quantities (KPI's), so that there is a clear statement of what kind of data has to be gathered on a regular basis by each participant.
- SLAs should always be stated in terms of *outputs* of *relevant* processes – for example, "the number of transactions of this or that type performed over the reporting period" – and not the *inputs* – for example, "the number of servers of this type that are available for transaction processing".

Otherwise, the SLA can all too easily end up saying that some "particular state" was intended and that some effort was made to achieve this "particular state". This is still a valid statement, of course, but probably not the one that was intended. This illustrates an interesting point – merely making a valid statement is *not* enough to make a useful and meaningful SLA, since perfectly valid statements can also be completely *irrelevant* and *ineffective*.

- SLAs should not contain "implicit functions" that have been left hanging and undefined – this is a very dangerous source of ambiguity, since a lot of complexity could arise hidden away inside them. In other words, it is important to clearly define the lexicon of general terms available to express the SLA and, where necessary, to define any specialised terms as they are introduced.

3. Current Security SLAs (circa 2005)

Examples of the current state of the art (circa 2005) in Security SLAs have included the following kinds of requirement:

- **"Critical" updates**

  for vulnerabilities regarded as critical by the customer the anti virus update file must be installed on connected computers within one hour of its availability. Performance is measured as the percentage of computers that achieve this target.

- **"Non-critical" updates**

  as for critical except that the time limit is relaxed to five hours.

This SLA specifies that the Service Provider should provide critical releases of the Anti Virus definition file to the organisation within 1 hour of its release by the anti-virus vendor. This service must be provided 24x7. The expected fulfilment of this SLA is typically in excess of 99%, with a minimum of not less than 97%.

Typically, the same service requirement and fulfilment criteria are specified for non-critical releases, but only during business hours.

This requirement raises a number of issues. The anti virus update file, which could be ~1MByte, needs to be distributed to all the machines on the network, which for a large customer could number many hundreds of thousands. Furthermore, these devises are connected by a heterogeneous network comprising many different types of circuit. In addition each device will need to go through an update, reboot and validation/scanning cycle and whether this time is included in the measurement of the SLA could be open to interpretation depending on wording..

Knowledge of current security operations processes suggests that 1 hour almost certainly impossible to achieve and probably requires re-engineering (at the application level, or caching), or validation outside 1 hour. If the process could be re-engineered to achieve 1 hour then we must ask why achieving 5 hours is any easier or substantially cheaper? What is the actual threshold and what is the cost of meeting the technical barrier?

If meeting this SLA is such a serious concern then one way of achieving it would be to mandate processes that shutdown other traffic to improve network performance and improve security, although of course this opens the door for a sophisticated denial of service attack on the customer as a side effect.

3.1. But is this SLA sensible?

We can critique some of the issues raised above in the form of questions about the *intent* of the SLA and its implications for *delivery* such as:

1. Who should decide whether a particular anti-virus signature is critical? Who should decide which anti-virus vendor to use?
2. Which network nodes actually require the anti virus update file within 1 hour of release – and how many are there? What would fulfilling this requirement do to network bandwidth?
3. Why is a fixed number of 1 hour *necessary*? Would fulfilling this requirement on, say, a more relaxed time-scale be just as *effective* from a security point of view?
4. Shouldn't there be another SLA saying something about what the distributed anti-virus data file is used for? In particular, the virus data file needs to be *used* by the anti-virus scanning process on each host – how do we know that this happens?
5. Why is the particular fulfilment number as a percentage specified? What is the *impact* in security effectiveness of these numbers?
6. The critical and non-critical variants of this SLA are very similar in terms of what the Service Provider needs to do. In terms of process, these variants are pretty much identical, except that critical releases of the virus data file *must* be expedited urgently. In practice, it is highly likely that these two variants of the SLA would actually be met by the same underlying systems engineering, rather than two independent mechanisms.

   The upshot of this is that, to meet the more urgent variant, the system overall will be engineered to a *higher* performance spec. than if just the less urgent variant were required. This means that the customer will typically get better performance overall, particularly for the less urgent case. Since the Service Provider is now more likely to significantly over-perform, shouldn't they be rewarded for their effort?

These are all natural questions to ask about an SLA like this. The rest of this report shows how asking questions like this can contribute to refining the offering – to the mutual benefit of the Service Customer and Provider.

The kind of example quoted above is not an isolated example. There are *numerous* examples of SLAs that have been encountered where the *lack of precise definition* leads to radically different outcomes for customers and vendors depending on interpretation. Greater precision of SLA definition allows more science to be applied to gain more accurate prediction of outcomes. Resolving this kind of ambiguity and performing predictive service estimation are some of the main reasons why HP Open Analytics was originally developed. This example is a typical case.

3.2. Making security into a differentiator for ICT services

It is clear that the baseline security proposition described in section Chapter 0 is something that all Service Providers need to sign up to. As such, this will merely become a matter of contractual *compliance*. Consequently, the provision of the IT delivery capability (i.e. running the Security Operating Centre and associated Policy Enforcement Points) represents a pure cost upon the Service Provider. Because every Service Provider is contractually obliged to deliver this generic security capability, it naturally cannot *by itself* provide a differentiator over our competitors. Thus, the best that can be done is to reduce the cost of providing this capability – and then use these savings to make our service offerings more competitive on price.

The above statement is true, but only if the baseline security proposition remains the *only* security-related value proposition to be delivered. HP should be able to do a good deal better and *create* differentiation by making a virtue of this necessity and offering *additional* value-added security-related services.

Any such additional services should exploit the existing investment needed to deliver the baseline security compliance requirement. If we can also find ways in making these additional services *measurable* in some natural way, so that we can associate SLAs with them, then so much the better.

4. Basic examples of security-related services and SLAs

We begin here by presenting some basic examples of security related services and some putative SLAs.

4.1. Network resilience to virus/worm attack

Within X hours of initially detecting a level-K virus/worm attack, at least Y% of the network infrastructure will be available and accessible.

4.1.1. Description:

The SLA conveys the idea that the impact of a virus/worm outbreak on the network will be constrained, that a sizable percentage of the network infrastructure will have sufficiently recovered within a certain period of time. This SLA is carefully restricted to a statement concerning network infrastructure (e.g. routers, switches and firewalls) and such like. Note that the SLA also provides potential for qualification by the *severity* of virus/worm outbreaks (e.g. "*level K*") to allow for a more graded response. However, this parameter would have to be specified, along with penalties etc., before any prediction can be performed. We assume that another part of the overall contract will tackle *client node* defences against viruses and worms (e.g. anti-virus).

4.1.2. Technical Development & Engineering:

The technical development and engineering needed to meet an SLA like this could usefully exploit the Virus Throttle technology developed by HP Labs and now incorporated within HP Proliant servers. Viruses and worms tend to infect leaf nodes rather than the network infrastructure directly and many cause a massive burst of connection requests in an attempt to infect other leaf nodes. The idea would be to deploy virus throttle into non-backbone servers and routers with the effect of rate limiting outgoing connections from leaf nodes onto the backbone. The effect would be to help maintain traffic levels by limiting the spread of virus/worm infections to other uninfected parts of the network. This helps to contain infection to isolated pockets.

We should recognise that this SLA has a non-zero downside risk – there are circumstances, with the best effort and best practice in the world, where a previously unknown worm/virus is released and then the outbreak is harder to contain – more damage will be caused and will take longer/more effort to recover from. Technology for rate-limiting message transmissions can help here whenever the new worm/virus attempts to spread in an explosive manner. Of course, if that is not the case, as may happen with a stealthily spreading worm/virus, then rate-limiting won't make much difference. One may thus argue that rate-limiting technology has the longer term consequence of encouraging virus-writers to write more stealthily spreading viruses. . This suggests that the SLA should be drafted to recognise when the Service Provider acts in the best interests of the customer and does the best job possible under the circumstances.

4.1.3. Analysis:

Value to the customer: We assume that the Service Customer has a significant requirement for network availability. If this SLA fails, then the network is unavailable for business use.

Predictable: As this is a value-added service, prediction of costs would have to be based upon standard costs such as software licensing and the need for additional trained personnel to operate the service outlined here. In general, there will be a need for more compute capacity – which requires more machines and thus increases the manpower to conduct operations and systems administration activities. (This amounts to a general statement that will apply to all the SLAs we shall explore).

Measurable: There might be difficulties in identifying and classifying a particular virus/worm attack in real time. However, network availability/accessibility is certainly measurable.

Affordable: Rate-limiting technology, such as HP virus throttle, is already commercially available and built into servers and switches. There are circumstances where this technology will not work and the SLA will be violated – due to novel, previously unknown attacks and weaknesses being exploited for the first time. This represents a significant downside risk for the Service Provider.

Understandable and Unambiguous: Ambiguities lie in what constitutes a "level-K" virus/worm outbreak. This needs to be defined appropriately so that the extent of exposure due to failing to deliver can be estimated.

4.2. Detection of unpatched and unmanaged machines

At least X host system scans completed upon Y% of the specified network IP address range. Each scan delivers a database report enumerating the following categories:

(a) Managed systems that are up-to-date for patch level (OS & apps).

*(b) Managed systems that are **not** up-to-date for patch level.*

*(c) Managed systems that are **unexpectedly** unavailable.*

*(d) Systems that are detected accessible to the network, but which appear to be **unmanaged** (e.g. "ghost IT"/"shadow IT").*

Each report delivered has fewer than N misclassifications of leaf node status.

4.2.1. Description:

This SLA is akin to "patrolling the beat" – the idea is that a number of scans are regularly performed over specified IP address ranges to determine configuration properties of each host leaf node detected. In particular, the scan should determine if the system is "under management" and, if so, determine its patch level concerning the OS it should have and other critical applications. The patch levels for the OS and critical apps software on each managed machine detected should be reported appropriately – as being up-to-date or not. Managed machines that are unexpectedly unavailable should be reported – their status could not be determined, and may thus need further investigation. Finally, any other machines that turn up on the network and appear not to be managed need to be reported as such.

The classification tree looks like the example shown in Figure 5.

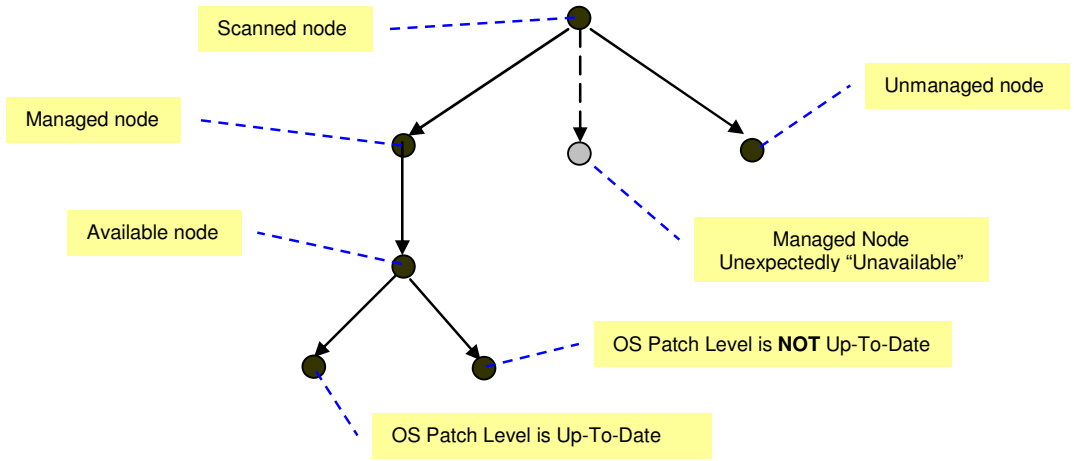

Figure 5 Potential Classification tree

Naturally, to be able to determine if there are any managed machines that should have been available that are not, it is necessary to know what to expect (i.e. the set of the managed machines that should be

available at this time). This means that the property of "unexpected availability" cannot be computed until the scan has either timed-out or has otherwise completed for the node in question.

4.2.2. Technical Development & Engineering:

The technical development and engineering required to meet this SLA could include both the HP Active Countermeasures and the Trust Record technologies, developed within HP Labs. The Active Countermeasures technology uses a variety of methods (including using known systems vulnerabilities) to gain entry to discovered systems, whereupon tests for patch level and whether the system is managed or not can be made. The advantage of this approach is that may also be possible to gain entry to any system discovered on the network, even if they are not currently managed. Assuming that each managed systems necessarily has a standard means for (secure) remote systems administration, this means that any system found on the network that is resistant to entry can also be classified as being "unmanaged".

Another advantage of this approach is that because the vulnerability testing uses known systems weaknesses to check for potential penetration, the system has a very low rate (essentially zero) of *false positives*. This will significantly contribute to the quality of the report produced, at least in terms of correctly classifying machines.

The Trust Record technology can also be used to provide secure, time-stamped records of event information which would be hard for either Service Customer or Service Provider to tamper with or falsify.

4.2.3. Analysis:

Value to the customer: HP demonstrating superior capability by discovering and locating machines that are unpatched and unmanaged. Helps customers understand the security exposure due to patch level of current systems, and to understand the current state of their networked system

Predictable: As for 4.1.3.

Measurable: The number of scans is clearly measurable. From the customer's point of view, they are paying for the quality content of the reports produced. This quality is expressed in the SLA in terms of the number of misclassifications

Affordable: It is difficult to assess how to price regular scans in a managed operating environment that is relatively threat free.

Understandable and Unambiguous: Potential for ambiguity lies in what constitutes a "misclassification". This SLA only makes sense in a regime where the results of a system scan are examined and tested as part of a larger systems administration process.

5. Service disaggregation: Anti-virus signatures

As we saw in section 3, the SLA quoted there tried to cover a complex security-related process as though it were a single process. An alternative approach is to break these up into more manageable service "chunks" or phases that the customer can appreciate and understand the security benefits of. Broadly, this also allows the Service Provider some greater flexibility, since each phase can be independently measured and charged for. Pragmatically, this flexibility is more likely to help accommodate a working, operational solution and thus deliver more reliably the value that the customer wanted in the first place.

At first glance, it is perhaps unusual to consider such a basic security process as 'Anti-Virus' as being a "value-added security service". It should become evident that the particular value-added by this proposal for an 'Anti-Virus' solution lies in knowing the extent to which the latest Anti-Virus signatures have been captured and used to reduce overall corporate exposure to the latest virus/worm. This simply translates into not only acquiring knowledge of the status of the corporate network and systems with

respect to worms and viruses, but also knowing that there is an effective process in place to deal with the remaining exposures.

5.1. Anti-virus signatures: Capture

Anti-virus signatures of criticality level N are made available on nominated internal servers within X_N hours of publication by anti-virus (A-V) vendor.

5.1.1. Description:

Upon release by the anti-virus vendor, the anti-virus signatures are assessed by the Service Provider[3] to have a particular criticality level, *N*. Having determined the criticality level, the signatures are made available to key internal servers within a defined amount of time, X_N hours. Instead of having different clauses in the SLA for different criticality levels, we can introduce a *function* (e.g. presented as an explicit table) giving the various limits required.

5.1.2. Technical Development & Engineering:

The Service Provider needs to have working relationships with their anti-virus vendors so that they get signature updates served efficiently according to criticality level, N. This may involve push-technology from the anti-virus vendors. The Service Provider will need to inform the Service Customer (via SOC) that the updates are in the process of being captured.

5.1.3. Analysis:

Value to the customer: This process is an essential part of virus threat management.

Predictable: As for 4.1.3.

Measurable: Logging of anti-virus signature capture events.

Affordable: Yes – the number of nominated internal servers that receive the signature updates is negotiable and scalable. In this way, the Service Provider can assess the bandwidth required and ensure that this requirement can be met.

Understandable and Unambiguous: Potential ambiguity arises in dealing with the criticality level, *N*. An appropriate scale needs to be negotiated and agreed with the Customer. Once the range of levels has been agreed, it is also necessary to specify how the criticality level is assigned in any particular instance, thus allocating responsibility for doing so.

5.2. Anti-virus signatures: Distribution

Distribute anti-virus signatures from nominated servers to all nodes in its anti-virus zone within X hours of receipt from A-V vendor.

5.2.1. Description:

This stage involves the distribution of freshly arrived anti-virus signature files from the nominated servers to all the nodes belonging to its *anti-virus zone*. This zone is a collection of nodes that represents the portion of the customer's view of the network that it coordinates and administers for anti-virus purposes. Clearly, the software has to not only initiate anti-virus transfers, but also accurately log and audit these transactions performed on behalf of the customer.

---

[3] The Service Provider should *negotiate* with the Customer which anti-virus vendor to use and deploy, perhaps by offering a selection of anti-virus vendors that the Service Provider supports. After all, the Service Provider is bearing the weight of *technical* risk for ICT service provision - this is presumably why the Service Customer hired the Service Provider in the first place!

5.2.2. Technical Development & Engineering:

For this automated administration to happen, appropriate software needs to be available to schedule transfer and distribution of the anti-virus signatures from the nominated servers in a way that will not overwhelm the network. Furthermore, the status of each of these transactions has to be recorded and audited in a secure, tamper-resistant manner. Secure record-keeping tools for auditing and data aggregation, such as the Trust Record technology under development within HP Labs, could be useful in this context.

5.2.3. Analysis:

Value to the customer: This process is an integral part of virus threat management.

Predictable: As for 4.1.3.

Measurable: Requires technology for coordination and secure record keeping.

Affordable: The Service Provider can control the size of each anti-virus zone (i.e. number of nodes) and so ensure that the impact on network performance of meeting this SLA is within acceptable limits.

Understandable and Unambiguous: The anti-virus zones need to be specified and kept up-to-date i.e. managed. In particular the anti-virus zones must partition the set of all systems to be scanned – thus, every system must belong to at least one of the anti-virus zones. However, it is wasteful if there is any overlap of anti-virus zones.

5.3. Anti-virus signatures: Scanning

All systems must perform a full system A-V scan at least once every X days, but no more than Y times within Z days.

5.3.1. Description:

The objective of this agreement is to ensure that the anti-virus signatures are used (by scanning) to detect and eliminate known virus/worm threats. This scanning has a direct benefit for both Service Customer and Service Provider – it helps to reduce viral/worm infection and thus reduces the potential for damage to the (shared) network.

However, one can have too much of a good thing – it is clearly true that if the nodes spent all their time scanning for viruses and worms, then there would be little time for nodes to do productive work on the customer's behalf. This agreement therefore places an upper limit on the number of scans within a given period - essentially an upper bound on the *rate* of scanning.

Records need to be kept on when scans were performed, which anti-virus signature files were used and so on. In particular, the *result* of performing a scan needs to be processed as it may indicate that there is a need to remediate and/or restore the system to a known safe configuration. These indications will need to be suitably flagged to trigger the following stage (see section 0).

Secure record-keeping tools for auditing and data aggregation, such as the Trust Record technology under development within HP Labs, could again be useful in this context.

5.3.2. Technical Development & Engineering:

Conventional anti-virus technology can be configured to run scans on particular local disks and scheduled to run at particular times of day, at particular days of the week or on particular dates. However, this also needs secure remote administration technology for this to be economically effective.

5.3.3. Analysis:

Value to the customer: This phase is the primary part of virus threat management. This phase delivers the primary security value – both of the previous phases exist to enable this one to operate.

Predictable: As for 4.1.3.

Measurable: Need to provide aggregated summary logs to customer (c.f. Trust Record)

Affordable: Yes – bounds ensure that the overhead of anti-virus scanning uses resources appropriately.

Understandable and Unambiguous: As usual, the parameters need to be unambiguously determined.

5.4. Anti-virus signatures: Remediation

Each system requiring level K remediation (as detected by routine A-V scan) will be remediated within X_K hours.

5.4.1. Description:

This agreement describes a conditional process that is intended to run only if required, as determined by the previous stage. Remediation involves restoring systems to their full operational state as efficiently as possible. The kind or degree of remediation effort can be graded (level K) and, where possible, scheduled as part of standard maintenance. The agreement asserts that any system requiring remediation at the appropriate level will be remediated within a certain number of hours. Depending upon how remediation is done, this could have a significant impact upon network performance.

5.4.2. Technical Development & Engineering:

We assume that any customer-specific data that any machine may process is persistently stored on a dedicated storage system accessible over the network. This assumption ensures that it would be technically possible to wipe any local OS images etc. without danger of destroying customer-specific data.

Since the remediation process is triggered on a per-machine basis, some automation software is needed to assess what remediation is required (if any) following each scan. The anti-virus scan status information may or may not be sufficient to help determine this requirement. In practice, such software could be very dependent upon a number of disparate factors – in particular, the OS & patch level, the particular applications required, and so on. This dependency makes determining exactly what needs to be remediated, and what does not, difficult to make.

However, one way to simplify this entire process is to say that if the estimate of the apparent damage is above a certain threshold, then the Service Provider should simply remediate the entire machine into a standard configuration agreed in advance with the Service Customer.

There are several degrees of freedom possible here – for example, variations in both threshold value and the particular configurations to be remediated. Each of these variations will have associated costs and security benefits which can be traded off with the customer.

Substantial remediation over the network is going to cause network performance problems, especially if a large number of machines are going to need remediation within a few hours. In practice, one solution might be to nominate local servers that can (concurrently) provide remediation to a number of their local nodes, based upon its own local `Golden' backup copies of these standard configurations. A further advantage of this strategy is that it allows for routine remediation of the configuration, in case the customer is concerned that configurations are being stealthily mutated in otherwise undetectable ways.

Secure record-keeping tools for auditing and data aggregation, such as the Trust Record technology under development within HP Labs, could again be useful in this context.

5.4.3. Analysis:

Value to the customer: During this phase, the Service Provider can conduct remediation as a part of systems recovery and needs to be done as quickly and as efficiently as possible to resume business processing as soon as possible.

Predictable: (Conditional phase) As for 4.1.3.

Measurable: Provide aggregated summary logs to customer (c.f. Trust Record)

Affordable: Yes – there are qualified limits on the kind of remediation (e.g. threshold levels, agreed kinds of configuration), which limits liability for the Service Provider to fixing certain kinds of problem.

Understandable and Unambiguous: The parameters that specify threshold levels, level K, and what remediation activity must be carefully specified to a particular level of operational effectiveness. The conditions that trigger this phase are defined by the scanning phase given earlier.

5.5. Anti-virus signatures: Overall SLA composition

The processing of Anti-Virus signatures consists of the following pipelined stages:

1. *Capture* *(see §6.1 above)*
2. *Distribution* *(see §6.2 above)*
3. *Scanning* *(see §6.3 above)*
4. *Remediation [as required]* *(see §6.4 above)*

- *Each stage is costed and charged for independently.*
- *Immediate penalties are charged only for the specific phases that originally failed. Any failures in subsequent phases directly caused by these are discounted. Thus, `knock-on' failures do not contribute to penalties here.*
- *However, an overall penalty is charged for the number of phases that failed.*

5.5.1. Description:

In the above, we describe the ingredients of a 'composition' agreement that covers how the various stages described earlier fit together. Specifically, we determine how charges and penalties for the overall anti-virus process are computed from the component stages. The basic idea is that the customer can, at least in principal, see the overall process more transparently and be aware of where charges are arising. At the same time, the above states that the Service Provider will not be unfairly penalised for inevitable `knock-on' failures.

At the same time, it would be equally unfair to the customer if there was insufficient acknowledgement that the service they had hoped for had not been achieved. One way that this might be done is to include an overall penalty proportional to the total number of phases that failed.

By way of an illustrative example, this might work as follows. Suppose that for some reason the distribution phase (see section 0) has not always performed adequately over the quarter (i.e. assumed reporting period).

Let us assume that the distribution phase failed F times in the 90 day quarter. This means that, occasionally, the subsequent phases were not done completely satisfactorily. In particular, although the scanning phase may proceed, it will do so using potentially out-of-date signature files.

The overall penalties charged might then be, for example:

*Penalty = F * (Penalty for stage 0 + charge for overall consequential failure)*

*= F * (PenaltyFor(0) + 3*PerStagePenalty)*

Note that the expression for the charge for overall failures has a factor 3. This value is made up by a contribution of 1 from each of the phases 0, 0, and 0. Even though phase 0 is dependent upon the earlier phases and may be unnecessary, we have to assume that the impact is the same. Significantly, this charge for overall consequential failure could be factored neatly into the penalty for each originating stage.

From the Service Customer's point of view, the penalties charged reflect the fact that a specific failure arose and that the overall service was degraded proportionately. From the Service Provider's point of view, the penalties reflect the actual situation and are proportionate to the extent that SLAs were not met in practice. A rational scheme for assigning penalties such as this is more preferable to either having a protracted reputation-damaging dispute between Service Customer and provider, or to the imposition of blanket charges that will be unfair to one or other of the customer or provider, and thus a source of future ill-will and distrust, leading to instability.

The benefit of a more transparent and rational assignment of penalties is to reduce risk of dispute and ensure smoother, more assured collaboration between customer and provider. This engenders mutual trust and encourages a more stable, sustainable relationship to develop between customer and provider.

5.5.2. Technical Development & Engineering:

To implement this agreement requires implementations of the sub-ordinate SLAs in such a way that stage failures can be isolated, tracked and logged. Dependencies between stages are clear in a pipelined architecture like this one – each stage depends upon earlier ones. Each failure incident should therefore be linkable back to the stage containing the originating failure and any penalties generated accordingly. Given appropriate tracking and process instrumentation technology, it should be straightforward to audit this process in an accountable manner.

5.5.3. Analysis:

Value to the customer: This agreement is necessary to show how all the subordinate SLAs relate to each other. Moreover it defines how to compute overall penalties arising as a result of SLA composition.

Predictable: As for 4.1.3.

Measurable: Tracking of component SLA failures involving process instrumentation – possibly using auditable Trust Record technology to aggregate process KPI's.

Affordable: Failure of this SLA or its non-implementation would mean that customers are incorrectly charged (either over or under).

Understandable and Unambiguous: The penalties need to be specified.

5.6. Discussion of anti-virus process disaggregation

This approach could work quite well. Anti-virus has the natural advantage of being a clear, linear process where each stage or phase has a defined objective, with clearly defined consequences of failure.

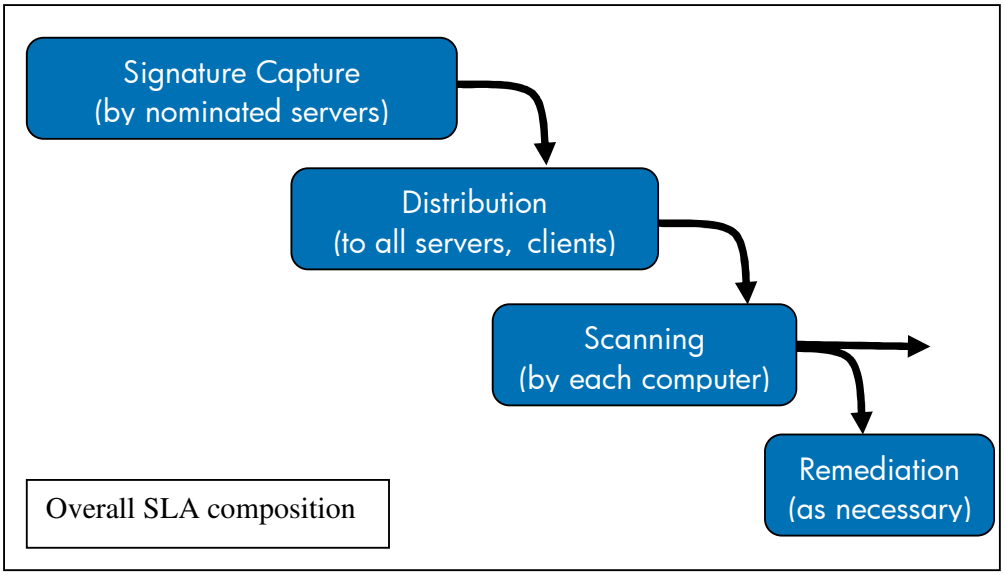

Figure 6 Process Disaggregation for Anti-Virus SLAs : Each stage has an SLA of its own.

5.6.1. Benefits of service disaggregation to provider and customer

- **Service Provider benefits:** Disaggregating the service permits the Service Provider to explain "knock-on" consequences in terms of phases and thus argue to reduce the impact of failures. The Service Provider may argue that a failure in one phase can lead to failures in later phases – and thus penalties should only be levied for causative phases and *not* for any subsequent "knock-on" failures. This would seem fairest, since otherwise the Service Provider would be penalised for phases that worked perfectly well, except that they had failed to meet their agreed objectives due to failures at earlier stages.

  The upshot is that the Service Provider may breakdown costs and charges in a manner that is understandable and transparent to the customer.

- **Service Customer benefits:** The Service Customer clearly benefits from service disaggregation because the overall process becomes more *transparent* and thus easier for the customer to understand and assess consequences to them and so estimate future performance.

A complex process is broken down into a succession of achievable phases dependent only upon the previous stage. The customer gains confidence and trust in the Service Provider because the overall process is predictable and liabilities (risk) are limited. This arises because the processes themselves are defined in simple, self-contained terms.

5.6.2. Finessing composite SLAs

It turns out, in practice, that composite SLAs are notoriously complex and difficult to adjudicate against – complex penalty schemes provide lots of opportunity for creative misunderstanding. Having developed a composite SLA that describes an operational process, the finesse is then to examine this process and then argue for a simplified SLA that captures the critical value to be delivered to the customer.

In the case of our Anti-Virus example above, the critical phase turns out to be Scanning – the two preceding phases of Signature Capture and Distribution are of course necessary for the Scanning phase to operate with up-to-date signatures. So, with the analysis in hand, the finesse is to propose that the SLA be based around the Scanning phase itself (i.e. Section 5.3), instead of proposing the overall composite statement given in Section 5.5. This means that the effect of the earlier phases has to be implemented in an effective way in order to meet the final SLA in the chain. As it is optional, the

Remediation phase following Scanning should operate under its own SLA, which can then more carefully qualify the circumstances under which it should operate.

Formulating the composite SLA is still useful as it represents an effective process decomposition that the Service Provider can use to plan and estimate their delivery. Furthermore, it helps to expose and uncover those goal(s) that the customer cares most about and thus helps to refine the SLA under negotiation.

6. Observations

This work noted earlier that the technical requirements of providing for *baseline security* for ICT services is an overhead, just like any other IT service. The security SLAs that were envisaged and proposed here are *value-added security services* that would help HP to further exploit the technical investment *already made* to implement this baseline security requirement. These proposed services were largely based upon security-related technologies currently emerging from HP Labs. The technical modelling work that would be required later focuses on *technical accountancy* and *management science* aspects - e.g. FTE manpower costings etc. Such modelling lies well within the scope of the existing financial/reliability modelling tools that the Open Analytics team uses (e.g. [Demos 2000]).

In section 5 we considered the problem of formulating an appropriate SLA for an Anti-Virus service, in response to the SLA discussed earlier in section 3. It is clear that the same methodology should equally apply to other security-related SLAs, including:

- **Patch Management**: Application of scheduled patches across the enterprise. The measure here would be the percentage of machines updated to a particular patch level.
- **Security Incident Management**: This involves diagnosis and response to security incidents. This is made more practical and realisable in service terms by the emergence of Security Incident Management technologies such as [netForensics].
- **Secure data transfer between endpoints within a corporate intranet**: This involves providing a data transport service between identified entities within a corporate network. The customer value would be availability (liveness properties) together with proposition that the information remains secure and will not be compromised during transport (integrity & safety properties).

Finally, further investigations into security operations are required in order to understand the issues involved in producing predictive costings (e.g. FTE manpower costs) that can say whether it would be economic to offer the kinds of value-added security services proposed here. Informal feedback from service delivery people within HP revealed that they are primarily focused on deploying standard security product offerings - i.e. delivering industry best-practice standards (ITSM), together with the use of market leading best-of-breed commercial products. The prevailing risk-averse attitudes implied that any vendor offering value added security services would unnecessarily incur significant levels of commercial liability and risk.

6.1. Implications for technology

We have investigated some potential candidates for "value-added security-related services" - by looking at possible services first and then later finding matching technologies that could be useful when deploying and running these services. Lots of security technology exists off the shelf as commercial offerings – however, there is still need for more systems with better integration. Perhaps because of familiarity, most of the candidate services considered made use of Labs technologies such as Trust Record and Active Countermeasures. It is clear that for deployment within a services context, Labs technologies like these would need to be further developed into commercial products.

This suggests that it might be fruitful to think *first* about the business service models, figure out which ones could be profitable and then determine what candidate technologies are necessary to support these business services. Importantly, by doing this, we get some idea of the cost/benefit profile that needs to

be met by these technologies so as to achieve profitable usage. This clearly helps give investment direction and provides a basis for deciding what technology developments to invest and disinvest in.

There is a clear need for secure remote administration technology to allow Service Provider staff to perform administration tasks from Security and Network Operations Centers, irrespective of location of equipment. Technologies such as HP Trust Record can provide necessary support for model-based assurance, which in turn encompasses secure record-keeping and audit of service processes. Process automation technologies, such as [SmartFrog], provides a means to help reduce excessive manpower costs by providing scalable automation of systems deployment and their administration tasks.

7. Summary and Conclusions

In this report, we explored the issue of meaningful SLAs in the context of security, starting from a security-related SLA (measurable distribution of Anti-Virus Signatures) that has arisen during the pursuit of a couple of IT Outsourcing deals. We discussed how a meaningful SLA embodies not only certain legal and financial contractual elements, but it is also necessarily associated with a *process view*.

This process view describes the service that the customer expects to receive and what the service provider promises to deliver. We illustrated this idea by "breaking down", or *disaggregating*, an SLA for an Anti-Virus service into several measurable processes. Such a disaggregation can be useful for both customer and service provider – it helps each understand what needs to be done and what each expects of the other in process terms. This in turn helps both sides to calculate expected returns – this means value for money for customers, and profitability for service providers. A further benefit is that this decomposition can help both parties to better negotiate the SLA for better and more assured delivery over the life of the contract.

All of this can be restated in the language of costs and benefits. However, to be meaningful, numerical measures need to be captured and understood in an appropriate *context*. Given that service delivery has to meet goals and deadlines defined by SLAs in numerical terms, that appropriate context is the process linking service delivery to the contractual SLAs. To manage costs, one must understand what affects them in the first place.

7.1. Acknowledgements

We are grateful for the various discussions, comments and considerations generously offered by our colleagues Adrian Baldwin, Chris Dalton, Jonathan Griffin, Keith Harrison, Chris Tofts, and Michael Wonham.

8. References and Bibliography

[Demos 2000] See http://www.demos2k.org/

[DBSy] See http://library.hp.com/techpubs/2005/HPL-2005-141.pdf

[SmartFrog] Smart Framework for Object Groups**.** See http://www.smartfrog.org/.

[netForensics] See http://www.netforensics.com/
http://www.sans.org/rr/whitepapers/tools/408.php and also
http://managementsoftware.hp.com/partner/isv/netForensics.jsp

[SAS 70] Statement on Auditing Standards (SAS) No. 70, *Service Organizations,* is an internationally recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). See http://www.sas70.com/

8. References and Bibliography

[Demos 2000] See http://www.demos2k.org/

[DBSy] See http://library.hp.com/techpubs/2005/HPL-2005-141.pdf

[SmartFrog] Smart Framework for Object Groups**.** See http://www.smartfrog.org/.

[netForensics] See http://www.netforensics.com/
http://www.sans.org/rr/whitepapers/tools/408.php and also
http://managementsoftware.hp.com/partner/isv/netForensics.jsp

[SAS 70] Statement on Auditing Standards (SAS) No. 70, *Service Organizations,* is an internationally recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). See http://www.sas70.com/