APPLYING ICS TIME METRICS

TO

GLOBALPLATFORM SMART CARDS

Dr. David Brewer
Gamma Secure Systems Limited
Diamond House, 149 Frimley Rd
Camberley, Surrey GU15 2PS, UK

William List *CA, Hon FBCS, CITP*
W[m].List & Co.
46 Snakes Lane
Woodford Green, Essex IG8 0DF, UK

INTRODUCTION

At eSmart 2003, Kekicheff and Brewer [1] reported on the development of the GlobalPlatform Card Security Requirements Specification (CSRS) [2]. They pointed out not only that the CSRS presents a *semi-formal* security specification for the entire card platform, but also that it proposes a method for selecting the most appropriate card configuration based on risk analysis. Particular attention was paid to the delay time between:

- The occurrence of an event and its detection
- Its detection and its subsequent correction/ rectification.

Subsequently, List and Brewer [3] have extended this concept of time delay to develop a mathematical theory (referred to as the Time Model) for measuring the operational effectiveness and cost-effectiveness of an internal control system (ICS). The objective of this paper is to demonstrate how this theory may be applied in the context of a GlobalPlatform smart card. In particular it shows how the most appropriate card configuration can be selected for a typical payment application. Furthermore, it shows how security considerations can be factored into a cost-benefit analysis for the selection of the optimum card configuration.

We begin by showing how the Time Model extends the concepts proposed in the CSRS by introducing:

- A set of time metrics
- A set of financial metrics
- An event-impact risk assessment methodology.

In accordance with that methodology, the paper presents a "risk treatment plan" (RTP) for a typical payment system. This RTP shows, using the Time Model, how the optional card security functionality (specified in the CSRS) should be chosen such that the residual risk is acceptable. Of importance, the RTP is expressed in terms that are designed to be understood by the application provider's business managers, and thus acts as a bridge between them and the information technologists. Finally, we show how the financial metrics are used to cost-justify the resulting choice of card configuration.

THE IMPORTANCE OF TIME

Since its inception, the specification of GlobalPlatform cards has always included some optional functionality. This has presented difficulty, in not only in how to specify it in languages such as the Common Criteria [4], but also and perhaps more importantly in how a card issuer should choose what functions are right for their particular purpose. Indeed, the Global Platform Compliance Packages document [5] identifies 37 optional packages. The CSRS proposes a way of addressing this problem, based on risk analysis. It argues that the acceptability of risk depends on a number of factors. It identifies time as a key issue, in particular the time taken to detect an event and subsequently to correct or otherwise rectify its effect. It also points out that the number of actors involved and their respective exposures are aggravating factors, e.g., when the actor able to detect the event (actor X) is different from the actor able to take the corrective action (actor Y).

The "Time Model" Concept

The Time Model [3] extends the time metrics proposed in the CSRS by introducing a third time parameter, the "time window", on the expiry of which some undesirable impact, often of a financial variety, occurs. The objective of the security controls embodied within a GlobalPlatform smart card and its associated off-card systems is therefore to detect an event in sufficient time for it to be satisfactorily countered before the expiry of the time window. The full set of time metrics are therefore:

- The time of detection (T_D if detected by the ICS, or if detected by some other means T_M,

e.g. reported by the cardholder or reported in a newspaper)

- The time that the damage caused by the event is fixed (T_F), should it be possible and appropriate to fix it, or otherwise resolve the problem

- The time limit after which (T_W), if the damage is not fixed, some impact penalty I_P (whether financial or otherwise) is incurred.

Their relationships are shown in Figures 1, 2 and 3.

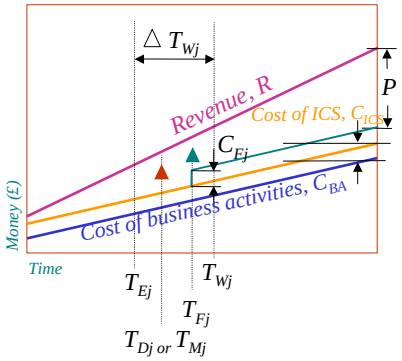

Figure 1: Detecting the event in good time to avoid the impact penalty

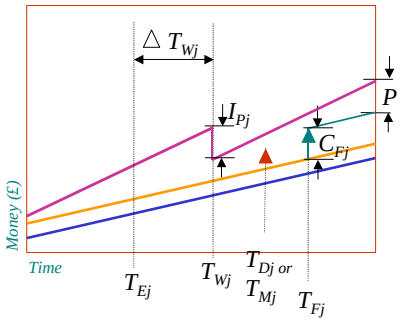

Figure 2: Detecting the event too late to do anything about it within the time window

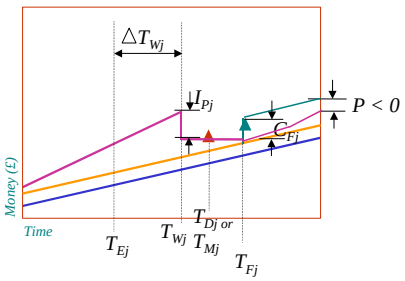

Figure 3: The onset of disaster

In the absence of the event, an application provider (say) would make some profit P, being the difference of some revenue R and the costs of providing the business activity, C_{BA}, and maintaining an ICS, C_{ICS}. The purpose of an ICS is to assist each actor meet their respective business objectives and manage their risks. Each actor therefore establishes its ICS with that purpose in mind. Each actor owns an ICS. Its senior business managers use and manage it. It consists of the on- and off-card functionality, procedures and people within scope of that actor's organisation. For example, an application provider's security domain[1] and applications would be in scope of their ICS, but not the issuer security domain. Note that the ICS relating to card activities will form part of the total ICS for each actor's business.

Figure 1 shows that some event happens at time T_{Ej}. It is detected, either on-card or off-card at time T_{Dj} or T_{Mj}, whereupon some correction or other remedy is made at time T_{Fj} at a cost C_{Fj}. For example, suppose the card had been lost and the cardholder reported that loss. The "cost to fix", in this case, ought just be the cost of blacklisting the lost card and issuing a new one. Note that T_{Fj} is less than T_{Wj}.

Figure 2 illustrates what happens if T_{Fj} is greater than or equal to T_{Wj}. In this case, not only is there a cost to fix, but the situation is aggravated by further financial loss due to the onset of some revenue penalty. For example, suppose instead that a thief had stolen the card and had been caught trying to use it after having made several fraudulent purchases. The cost to fix, in this case, would probably include the costs of any investigation/ prosecution as well as blacklisting and card replacement. The revenue penalty would correspond to the value of the fraudulent purchases that have to be made good.

Figure 3 illustrates the worst case where the losses due to the impact penalty are severe. In the general case of an ICS, this corresponds to a disaster and would usually require the invocation of a disaster recovery or business continuity plan. In the case of a system using cards one of the most damaging disasters would require the immediate cancellation and reissue of the card base.

The cost of some event is the cost of that part of the ICS that is designed to deal with the event (e.g. the cost of "chip and PIN"), plus the cost to fix plus the impact penalty, should it arise. It is possible that in many cases the cost of the ICS can be amortised over many similar events. Reference [3] cites particular cases and gives specific inequalities for

[1] The Card Specification [H2H] describes a security domain as the "on-card representative of an application provider", and the card issuer domain as the "on-card representative of the card issuer".

each, allowing the cost-effectiveness of the controls to be determined.

CLASSIFICATION OF CONTROLS

The Time Model classifies controls as belonging to one of seven classes, depending on their ability to detect an event and take corrective action in relation to the time window. Classes 1 to 3 achieve their desired objective before the expiry of the time window. Classes 4-7 do not.

In reality the control classes form a continuum, the seven classes corresponding to particular regions of that continuum. Their definition is given in Figure 4.

Class	Ability to detect the event and take recovery action	Type
1	Prevents the event, or detects the event as it happens and prevents it from having any impact	Preventive
2	Detects the event and reacts fast enough to fix it well within the time window	Detective
3	Detects the event and just reacts fast enough to fix it within the time window	
4	Detects the event but cannot react fast enough to fix it within the time window	
5	Fails to detect the event but has a partially deployed BCP	Reactive
6	Fails to detect the event but does have a BCP.	
7	Fails to detect the event and does not have a BCP.	

Figure 4: Control Class Definitions

The table further classifies the controls as being:

- *Preventive* - which seek to ensure the impact never materialises. This type of control either prevents the event from occurring or affecting the organisation, or detects the event as it happens and prevents any further activity that may lead to an impact.
- *Detective* - which identify when some event, or events, have occurred that could lead to a materialisation of the impact, and invoke appropriate actions to arrest (or mitigate) the situation.
- *Reactive* - which identify the impact has occurred and invoke appropriate actions to recover (or mitigate) the situation.

Classification of the on-card controls

The CSRS specifies the on-card security function requirements in the form of 51 tables (labelled in [2] as Table 5-1 through Table 5-51). There are four types of function, referred to by their colour coding: blue, green, yellow and tan. Each table may be conceptually regarded as an asynchronous process that is triggered when its pre-condition is satisfied. Yellow and tan tables are mapping functions and transform some input to some output. Blue tables are access control functions and authorise or deny access to some function that is represented by a green table or yellow table. Blue tables therefore have true and false outcomes (respectively request_authorised and request_ denied), the outcome being decided by evaluating a set of rules defined within the table. Green tables represent functions that have known failure modes. For example, the function that loads an application on to the card will fail if there is insufficient room on the card. Like blue tables, green tables therefore have true and false outcomes (respectively function_OK and function_failed). These too are determined by evaluating a set of rules defined within the table. In addition, a false[2] outcome in both blue and green tables may be accompanied by:

- Returning an error message (either an API error code or an APDU-R message) or by raising an exception as appropriate, *and/or*
- Invoking the "failure management" function (Table 5-45) to rollback the failed function, *and/or*
- Raising an alarm.

Note that some people may regard certain green and yellow tables as services rather than controls (for example Tables 5-42 and 5-43 which create audit data and Table 5-49 the random number generator). In this case these services may not have a control class but inherit the class of the security function using the service at the time of use. If for example Table 5-50 encryption/decryption will return garbage as the failure mode this will be recognised by the GlobalPlatform security features that use the service because there will an inequality on comparison of say a hash. Similar procedures will no doubt exist in applications using these services.

Blue tables

Blue tables represent access control functions. Their purpose is to control access to on-card functionality. The event of interest to these tables is the attempt of some entity, either on- or off-card, to access the controlled function. The blue table security features detect such events as they happen (see Table 1) and are able to prevent access if the

[2] For Table 5-46 Sensors and Alarms Security Feature, the alarm is raised on the true outcome.

entity does not satisfy one or more of the access control rules specified by the CSRS for that table. Thus the blue tables represent preventive controls, i.e., Class 1.

Green tables

Green tables represent functions, such as LOAD [for LOAD], LOAD [for INSTALL], which may fail under certain predetermined circumstances. Their failure is a security issue, as it might leave the card in an unsafe state. However, the success of these functions is not, as their execution will have been authorised by the relevant blue table[3].

In the majority of cases, these functions fail safe, leaving the card in a safe state. These green security features can therefore be regarded as detecting the event (i.e., function failure) as it happens and preventing the card from entering an unsafe state. These green tables are therefore also preventive controls, i.e., Class 1. However, this is not true for all green tables. Some are unable to leave the card in a safe state. Instead, they detect that the function has failed and invoke the failure management security feature (Table 5-45) to rollback[4] the function to the state before the invocation. These tables are therefore Class 2, 3 or 4, depending on how quickly the rollback function can be completed with respect to the onset of some adverse impact (see Figures 1 and 2). In the general case, rollback will be attempted immediately and unless the rollback function fails, we can regard the rectification of function failure to have been completed well within the time window. These security features can therefore be regarded as Class 2[5].

If the rollback function fails, the event action security feature (Table 5-40) determines what happens. This function is invoked whenever there is an alarm. The table is re-entrant. The action taken is pre-determined by the card issuer, who is cautioned to ensure that the feature does not cycle forever but will take some finalistic action such as blocking the card.

A particular failure mode is power failure, often caused by premature removal of a card from a terminal. The CSRS specifies how this is detected. The applicable security features cannot prevent the damage that may occur, but they do allow for it to be detected on subsequent power-up and for rectification then to take place. These tables are therefore Class 2 as well.

In summary, green tables are Class 2 if they invoke the failure management process or raise an alarm. Otherwise they are Class 1.

Yellow tables

With one exception, yellow tables pre-empt the event of an attempted attack, either as a control or service to other features. They do not prevent the event but take action to prevent the attack from succeeding. They may therefore be regarded as Class 1. The exception is Table 5-45 (failure management), which can raise an alarm. It is therefore Class 2 (and strictly speaking ought to be a green table).

Tan tables

These specify the interface between the card and the card manufacturer in the Pre-OP ready state, and serve the purpose of entering data, such as the transport key. From the perspective of the card, on completion of the data loading procedure the card will be placed into a secure state, thereby pre-empting particular types of attack. Tan tables are therefore Class 1.

Classification of off-card controls

The foregoing shows that all security feature tables specified in the CSRS are either Class 1 or Class 2. This conclusion is valid provided that the card meets the requirements of the CSRS for the chosen configuration and that the associated policies, also defined in the CSRS, are satisfied by the off-card systems. In Common Criteria [6] terms, these policies correspond to the assertions/assumptions made in protection profiles and security targets. It is also critical that an appropriate card configuration has been chosen.

However, what if the card does not behave as intended, or the associated policies are not properly implemented? Suppose a correct implementation, e.g. of Table 5-51 (Tamper resistance) fails in practice. What happens if the card blocks itself? In these and any other cases, how do we know that something untoward has happened? What do we do about it? Does it really matter if we do nothing? The process of answering these questions is dealt with through the process of risk assessment, which we discuss next. The purpose of the off-card ICS is, of course, to detect the event in sufficient time to avoid the impact penalty.

---

[3] File integrity, in the sense that the file received is the same as that sent by the card loader sent is assured through the secure channel (Table 5-24, which is a green table). If there is an integrity violation, it is that function that fails. Similarly, Tables 5-29 and 5-30 deal with DAP verification and mandated DAP verification. These, however, are blue tables.

[4] In the case of a delete failure, the failure management security feature will "roll forward" and complete the operation.

[5] Table 5-44 (Self test) is class 1 if scheduled to run on startup and class 2-4 otherwise.

RISK ASSESSMENT

Our risk assessment approach starts with the *events* and the *impacts*. The event is something that causes the impact. In business terms, the impacts that seem to capture the interest of senior executives include:

- Adverse press coverage
- Customer dissatisfaction
- Inability to carry out some or all of the organisation's business
- Loss of revenue
- Unanticipated costs
- Court action against an employee or the organisation itself.

Starting with these, we then ask what events might cause them. In practice, we have identified eight standard events, which we believe are common to most, if not all, organisations, to which we invite senior management to add those that are the special concerns of the organisation itself. The eight standard events are:

- Theft
- Acts of God, vandalism and terrorism
- Fraud
- IT failure
- Hacking
- Denial of service
- Disclosure
- Legal.

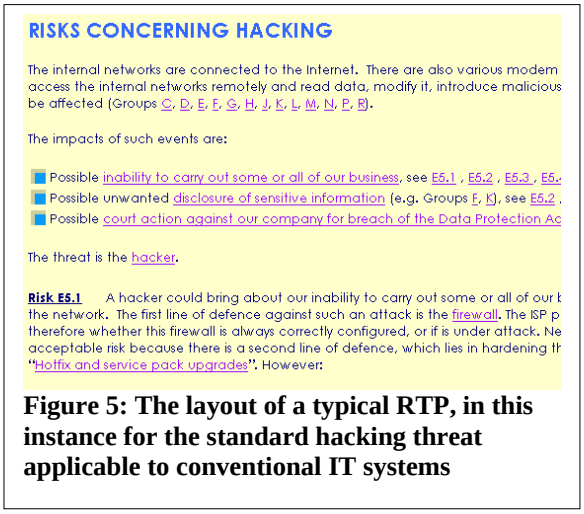
RISKS CONCERNING HACKING

The internal networks are connected to the Internet. There are also various modem access the internal networks remotely and read data, modify it, introduce malicious be affected (Groups C, D, E, F, G, H, J, K, L, M, N, P, R).

The impacts of such events are:

- Possible inability to carry out some or all of our business, see E5.1 , E5.2 , E5.3 , E5.
- Possible unwanted disclosure of sensitive information (e.g. Groups F, K), see E5.2 .
- Possible court action against our company for breach of the Data Protection Ac

The threat is the hacker.

Risk E5.1 A hacker could bring about our inability to carry out some or all of our b the network. The first line of defence against such an attack is the firewall. The ISP p therefore whether this firewall is always correctly configured, or if is under attack. Ne acceptable risk because there is a second line of defence, which lies in hardening th "Hotfix and service pack upgrades". However:

Figure 5: The layout of a typical RTP, in this instance for the standard hacking threat applicable to conventional IT systems

The risk treatment plan (RTP) concept

We have developed a technique for carrying out a risk assessment that uses events and impacts as its starting position. Application of the procedure, described in [3], identifies what preventive, detective and/or reactive controls are necessary to reduce the risk of an impact to an acceptable level.

The ISO term for the resulting document is a "risk treatment plan" (RTP). Some senior executives, who have produced these in accordance with [3] for their respective organisations, have likened them to writing down the story of their "worst nightmare", albeit at each twist of the story proudly showing how their controls save the day. Some have found some of their controls wanting, but the flaws we have been shown have been extremely subtle, and it is easy to see how they might have been missed before the application of this approach.

In developing an RTP, the idea is to ensure that wherever possible the *occurrence of the event can be detected in sufficient time* to do something positive about it before the impact occurs.

In some cases, it may be possible to *prevent* the event or detect it whilst it is happening. It is then necessary to consider the *detective* controls, if for no other reason than to appreciate that in practice the preventive controls may fail. Finally, it is necessary to consider the *reactive* controls.
The layout of a typical RTP is shown in Figure 5.

Example

As an example it is useful to consider a payment application, such as EMV [7]. Let us assert that the application provider for the EMV application is *not* the card issuer, but wishes to own the only payment application on the card. A further concern is that other applications on the card, their application providers or the card issuer might do things that damage the reputation of the EMV application provider. The EMV application provider therefore prepares an RTP to reduce these risks to an acceptable level, identifying in passing the optional GlobalPlatform card functionality packages [5] that they require. The application provider then uses this RTP to assist in identifying and selecting a card issuer with which to partner. Clearly other considerations will be important, but these are outside the scope of this paper. In developing the following fragment of a RTP, will we take an extreme position and not consider costs. We reconsider that position in the next section.

Note: for the remainder of this section *text for inclusion in the RTP is in italics*, whilst explanations and comments are in normal text. Words *we* and *our* in the RTP text refer to the EMV application provider. We will illustrate how the RTP might develop for some event called "ADVERSE ACTIONS BY CARD ISSUSER, OTHER APLICATION PROVIDER AND THEIR APPLICATIONS", which could be briefly described as follows:

The card issuer may make take actions, for example in managing card content and attending to our requirements, that in practice are not in our best business interests. Likewise, other application providers, or their applications may also act in a manner that is prejudicial to our business.

Having defined the event, the RTP would then proceed (see Figure 5) to enumerate the assets, impacts and threats:

The information assets that could be affected are (a) our application (b) our customers' data[6].

The impacts are:

- *Customer dissatisfaction, and thereby loss of revenue, and possible court action and adverse press coverage*
- *Unanticipated costs.*

The threats are the card issuer, other application providers, other applications, cardholders and attackers.

Starting with the requirement to be the sole provider of the EMV application, the first risk might be developed as follows:

Risk 1a. Our contract with the card issuer is that there will only be one EMV application on the card, that one being ours. Suppose post-issuance there were two or more. We can take action to prevent a second application provider from loading such an application by having a veto on application downloads. The veto is established though inclusion of Package C (Mandated DAP Support, see [2]).

Package C is a preventive measure. We now deal with the possible failure of that control.

Risk 1b. Notwithstanding the GlobalPlatform on-card controls, should they fail or some enterprising attacker find a way around them then other applications may be downloaded contrary to our wishes. These, should they exist, could be detected by periodic audit, but that can only be done if we have control of the card, i.e. our application or security domain is SELECTED. Also it might take too long compared to performing the transaction that the cardholder wishes to transact thereby giving rise to customer dissatisfaction. In monitoring the EMV transactions, we could look for a gross drop in the total number of transactions for all customers, but that would not distinguish between a customer using our card with another payment application on it, using a different card, or a true drop in business activity. We therefore have to accept this risk[7].

The RTP would continue by discussing what to do if a problem is discovered. Other risks should then be dealt with in a similar manner until all residual risks are deemed acceptable and management is confident that the assessment is sufficiently comprehensive.

To provide further examples, some of the other risk treatments might include:

Risk n. The card issuer or another application provider might attempt to delete our application. Attempts by other application providers are prohibited by the GlobalPlatform card access control rules. The card issuer has overall control over the card content and could therefore delete our application without our authorisation. There is a contract in place, so such deletion is likely to be in error. For individual cards, the cardholder would complain. The time taken to repair the damage could, however, be unacceptably long if we need to rely on the card issuer to do the job. We must therefore be able to have delegated management responsibility, i.e., Package D (Delegated Management), and the ability to load applications, i.e., Package E (Delegated Loading). Should our application be deleted and a customer complain, the customer will be given instructions on what to do to allow us to rectify the situation quickly, thereby restoring customer confidence. Our failure to do this is an acceptable risk.

Risk m. The card issuer or another application provider might attempt to extradite[8] *our application. This is prohibited by the GlobalPlatform card access control rules. If an attacker were to find a way around these, it would be as if someone were to find a way around the veto described in Risk 1b, and would detect the problem that way. This is therefore an acceptable risk.*

[6] In practice, this list may be longer but would certainly include these two.

[7] In practice, the application provider will be monitoring transactions for a variety of reasons. One reason would be to help ensure customer loyalty. There could therefore be a RTP that considers the event "CUSTOMERS DISAFFECT TO A COMPETITOR". Risk 1b, rather than accepting the risk would refer to this other RTP for the risk treatment.

[8] LOAD [for EXTRADITION] allows an application to be downloaded onto the card just once but then instances of it installed (and thereby associated with) two or more application providers.

Cost benefit analysis

Figures 1 to 3 illustrate the costs involved in providing controls and in not providing controls. A particular consideration in deciding whether it is cost effective to introduce a control or not is the frequency of attacks. Reference [3] uses the new requirement for *chip and PIN* as an example. It points out that at one time the payment associations appeared to tolerate the loss of significant sums of money due to fraud, although when challenged they were quick to point out that these sums very, very small compared to the total volume of transactions. Thus, at that time credit card fraud appears to have been an acceptable risk to the card issuers[9]. With the widespread introduction of chip and PIN, it would appear that that risk is no longer acceptable. The cost of introducing and maintaining chip and PIN now outweighs the losses due to fraud.

By itself chip and PIN will not, and cannot, reduce the set of attempted fraudulent transactions to zero. It will not stop the thief who guesses the PIN, or found it conveniently written down in the person's wallet. It will not stop the genuine cardholder from spending more than the application provider is willing to lend them and deliberately not paying. Other controls, which already exist such as authorisation limits, are necessary to do that. What is does do, however, is (a) decrease the time between the event (attempted unauthorised use) and its detection; (b) increase the reliability of that detection.

The "cost balancing" inequality is therefore:

Number of cards times *cost of chip & PIN* must be (significantly) less than the (estimated) reduction in *number of frauds* time the average *cost of dealing with a fraud*.

With regards to our RTP example (risk 1a), it is all right to deduce that we need a veto (mandated DAP support) but how much will that cost? Will that outweigh the risk of having another EMV application on the card? Remember, if the EMV application is preloaded onto the card, or better still installed and personalised pre-issuance, any other EMV application has to be downloaded. Thus, a competitor's application will have to be downloaded. Remember also, that a security domain with mandated DAP privilege must approve every application, not just those which particularly concern us. Thus it might not be worth going for the veto option in practice.

There will also be a cost associated with delegated management. However, in this case we would need to factor in some other benefits, which may also contribute to customer satisfaction. These ought to be developed in some other RTP that concerns problems attributable to the EMV application provider. For example, if a problem is discovered in the implementation of the application or if the EMV application provider wishes to upgrade it, for example to interface with new loyalty applications, there will be need to exchange the current application on each card for a new one. It may be faster and more reliable if this activity is performed by the application provider, rather than by the card issuer. In turn, this may give rise to greater customer satisfaction. The alternative may be less expensive, but could this bring about customer dissatisfaction? The answer may well lie in the frequency of occurrence.

[9] In order to simplify the argument we have excluded the views of the regulators.

CONCLUSIONS

The CSRS introduces the concept of time to detect and time to fix as a means for selecting optional Global Platform card functionality.

We have extended that concept by first introducing the concept of a "time window" upon the expiry of which some undesirable impact occurs. This allows a taxonomy of seven control classes to be developed, spanning the traditional preventive, detective and reactive controls. We have shown that the controls on-board a GlobalPlatform card are predominantly Class 1, although some by necessity are Class 2. However, the performance of these controls in practice will depend on a variety of external factors, such as the quality of engineering, the ability of actors to fully implement the required off-card policies and the ingenuity of attackers to find new ways of successfully attacking the card. It is therefore always necessary for each actor to ask the questions:

- What if the on-card controls do not work?
- How will I know?
- How quickly will I know?
- What do I do about it?
- Does it matter?
- What is the overall effect on the bottom line?

In order to answer these questions, we expanded the Time Model to address costs and presented an event-impact led approach to risk assessment.

This approach leads to a risk treatment plan, which can be produced by the senior business people and expressed in a language they (and everyone else) can understand. In the process of drawing up the RTPs, those that impinge on card issues will assist in identifying the most appropriate card

configuration. The answers can be checked through a cost benefit analysis. Overall, the approach leads to a determination of the most appropriate card configuration based on sound business arguments, backed by the work already done in ensuring that the GlobalPlatform card provides a sufficiently secure platform from which to work.

REFERENCES

[1] "*The GlobalPlatform Card Security Requirements Specification*", Kekicheff, M., Brewer, D.F.C, Proceedings eSmart 2003, Sophia Antipolis, September, 17-19, 2003

[2] "The GlobalPlatform Card Security Requirements Specification", GlobalPlatform, May 2003, http://www.globalplatform.org/

[3] "Measuring the effectiveness of an internal control system", Brewer, D.F.C., List, W., March 2004, http://www.gammassl.co.uk/topics/time

[4] "Smart Cards: The Open Platform Protection Profile (OP3)", Kekicheff, M., Kashef, F., Brewer, D.F.C, Proceedings of the Second International Common Criteria Conference, Brighton, UK, 2001

[5] "GlobalPlatform Card Specification 2.1 Compliance Packages", Version 1.1, June 2002, http://www.globalplatform.org/

[6] "Common Criteria for Information Technology Security Evaluation", ISO/IEC 15408:2000

[7] "Integrated Circuit Card Specification for Payment Systems" EMV2000, Books 1 – 4, Version 4.1, June 2004, www.emvco.comT