ARTICLE:

TECHNOLOGY AND BANKING: LESSONS FROM THE PAST

By **Ken Lindup**

Introduction

Anyone who spends a long time in a job builds up a fund of stories, some of which are true. Those of us involved with information security are no exception, and in this article, I want to share some of the events I have come across. I can vouch for their truth, but for obvious reasons I will not identify the names or nationalities of any of the organisations that were involved. What matters is not that the organisation was the victim, but the lessons that can be learned. Why are these failures relevant today? To paraphrase George Santayana, 'Those who cannot remember security failures are condemned to repeat them.'[1]

Poor product security – can we rely on what the vendors tell us?

The first two examples concern security features that appeared at first glance to be very effective, but which in practice were almost worthless. In both cases the equipment involved was supplied by well known multi-national vendors. I include them because they demonstrate that when building systems, it is important to verify vendor claims about security measures on which you rely.

Keys that are not all they seem to be

There are a number of generally accepted good security and control practices covering the use of terminals for financial transactions.

1. The user should always be positively identified using a combination of user-id and password.

2. Sensitive functions such as system administrator; security officer, etc. should be split between different people.

3. Encryption should be used either to encrypt a message or to generate message authentication codes to 'lock' the message electronically.

4. Control of encryption keys should be subject to stringent control, and ideally the key should be split into two halves, each half being controlled by a different individual.

One such system was concerned with electronic funds transfers with transactions involving sums around US$0.5 billion, which are not uncommon. All the normal precautions were taken, the organisation was set up to allow the separation of sensitive duties between different individuals; strong encryption was selected to protect the messages being transmitted. The most critical elements were even subject to a third party review.

At the heart of the terminal security was control of the transactions involving the following: the entering of the encryption keys, and the setting up of sensitive functions such as system administrator and security officer.

In accordance with the very best security practice, the separation of duties and dual control were enforced by the use of a physical key in addition to the user-id and password. The vendor advised that the key locks (built into the keyboard) functioned as electronic switches, and that the locks used high security barrels. A number of checks were made, which established that removing the locks or substituting another keyboard would not work, because the operation of the keys sent a signal to the processor. All the barrel numbers were known and a record maintained of which user had which barrel. A system was set up using a trusted locksmith whereby duplicate keys could be obtained under special circumstances. It appeared that a secure system was in place.

Unfortunately the locks were not high security locks; they were ignition locks from a popular brand of motor

1 *George Santayana, The Life of Reason: Reason in Common Sense (Scribner's, 1905), p 284.*

car. Duplicate keys could be obtained by walking into any service bar and asking for one to be cut, and this was exactly how the weakness was discovered.

When this was discovered, further checks were carried out which revealed that the switch mechanism was a contact breaker set from a car manufacturer. Inserting a thin sliver of plastic (from a coffee cup) between the lock barrel and the keyboard opened the points and bypassed the lock completely. A security measure at the heart of a high value electronic funds transfer system was worthless. To my knowledge, nobody took advantage of the weakness, and no one else ever discovered it. If it had been exploited, there could have been some interesting litigation to establish who was liable for any losses.

This example demonstrates the importance of taking nothing at face value and checking the substance of every claim in detail. The lesson is that just because it is encrypted, does not mean it is secure.

My second example concerns an early range of ATM machines. Any ATM requires a mechanism for identifying the customer and establishing that they are authorised to use the ATM. Customers are issued with a plastic card which contains a magnetic stripe on the reverse. The stripe contains many items of information including information to identify the issuing bank, the customer account number, and a means of validating the PIN. This problem occurred before the days of chip and PIN.

The PIN could be validated in one of two ways: they are still called online and offline. In an online system, the PIN is entered on the pin-pad by the customer and encrypted in the ATM. It is then transmitted to a central computer where it is verified. With offline working, the PIN entered by the customer is validated in the machine using information read from the stripe. This piece of information is called the offset. It is not the PIN; it is not even an encrypted version of the PIN. It is a number that can be combined with the PIN (using a cryptographic algorithm) to arrive at a predicted result. The key to this security mechanism is the strength of the encryption algorithm.

The vendor of the ATM in question had given assurances about the strength of the algorithm, which were accepted (it was a very large supplier of ATMs). The system was used for several years until one bank wanted to resolve a quality problem in the card production process. Cards were made on an embossing machine that took the information about the card holder from a reel of computer tape. The bank was having problems with errors that were being recorded in the magnetic stripe. The solution adopted was to add a further stage to the embossing machine. The additional stage took the information being fed to the stripe recorder and compared it with what had actually been recorded in the stripe. The embossing process required knowledge of the encryption method, and this was also required by the checking stage. The ATM vendors decline to release this information on the basis that it was crucial to the security of the whole ATM system and was a closely guarded secret. Their point was accepted, and it was decided to go ahead with the change, but not to validate the offset. The electronics engineers then discovered that the encryption process was described in detail in the field engineering manuals for the ATM.

These had been released to one set of engineers by the ATM manufacturer in the belief that they contained no information relating to security. What they actually contained was information that could have been used by criminals to build a device that would have enabled them to find the PIN of any plastic card that came into their hands.

The strength of an encryption algorithm is only part of the story. Security should not have to rely on the algorithm being secret – the best example of this is DES, where the algorithm is a published standard. Good algorithms have a variable component called the key. This must be kept totally secret. One method of attacking cryptographic systems is to collect the same information in encrypted and unencrypted form and analyse this to derive the algorithm and key. The trick is to make the amount of data and computing effort so large that it is not worth doing. This is called the work factor.

The algorithm in these particular ATMs had a variable key that was entered in two parts by two different people. The engineers discovered that by using about 15 ATM cards three times a week for about three months, they would have collected sufficient clear text and cypher text to find the encryption keys in use (they already had the algorithm from published sources). The algorithm had a very low work factor. When confronted with the claim, the vendors denied its truth. The solution adopted by the bank was to stop the offline PIN verification and rely on online verification only. Fortunately, these particular ATMs are no longer in use, and the banking industry had a lucky escape.

If these weaknesses had been exploited, any system based on the ATMs in question would have to have been shut down, the systems changed, and new cards issued to every customer. The cost would have been measured in many millions of dollars. It is also likely that claims

regarding unauthorised withdrawals from customers' accounts would have been difficult to refute.

Since these events occurred, much has been learned about security, and it is improbable that they could happen again. That is not to say that other weaknesses are not being implemented in secure commercial systems.

The security of the total system is what matters

Another example concerns early ATM systems and the card production process. The problem hinged on the differing security requirements of offline and online PIN verification systems. The PIN is used to authenticate the customer. All authentication schemes can follow one or more of three principles:

Something you are

Something you know

Something you have

A signature and fingerprint are examples of something you are, a password is something you know and a car ignition key is something you have. The PIN is 'something you know', the question is, is the plastic card you put in the terminal 'something you have'? With the exception of cards using specialised techniques such as magnetic watermarking or high hysteresis, the answer is it depended on whether PIN verification was conducted online or offline. The information on the back of a card is easily read – it consists of three tracks and complies with a published standard. Some of it is encrypted and some is in clear (unencrypted).

Offline PIN verification

When the PIN is verified in the terminal, it is essential to have the correct card and to know the PIN. The PIN entered through the keypad is processed together with the offset recorded in the magnetic stripe. The stripe is part of the encrypted information, and even if it is read, it is not meaningful. If a card is lost or stolen, the question is whether it can be used. If the card holder has stored the PIN with the card, it is simple for a thief to use the card. If not, the thief has to obtain the PIN, which is not particularly simple – although I described, above, how one system made it easy to obtain the PIN. Consider the other question; if someone obtains the details of a customer and somehow discovers the PIN, it is possible for them to manufacture a duplicate card? With offline verification it is difficult. A criminal would have to calculate the value of the offset and write it in encrypted form in the stripe. In such a case the thief must possess not only the account details and PIN, but must also possess the actual card. However, even this has its vulnerabilities, because a thief could steal a card and copy the offset from his or her own card and write this into the encrypted track. The thief would know the PIN associated with the new offset.

Online PIN verification

For countries that have not implemented chip and PIN, if the PIN is verified online, the account information is read from the unencrypted part of the card and transmitted with the PIN (entered on the keypad) to a central computer where validation takes place. The card plays no part in the security process. It is there to save the customer the bother of entering account information through the keypad. If a criminal knows the PIN and account details, it is a simple matter to take any card with a stripe and make it into any other. The manufacturer of my car gave me a plastic card containing its details for the convenience of the garage. I could have made it into a credit card or ATM card. Unless the issuer is using other specialised techniques, such as magnetic water marking, there is nothing that links a particular card to a particular account or customer. The card is therefore not 'something you have'.

One bank had originally decided that to maintain customer service when communication lines were down, they would do offline PIN verification. At the time, it was using a third party to manufacture or emboss the cards. This involved sending magnetic tapes containing all the account details and the PIN to the card manufacturer. The tapes were transported using security courier. None of the information was encrypted in any way. The original system specification called for the information to be split over two tapes, so that the PIN and account information did not come together until the embossing process. Somehow in the building of the system, this was changed so that only one tape was used. Since the PIN verification was offline, it did not matter too much. As we have seen, knowledge of account information and PIN was of little use because without the encryption process to generate the offset, it is all but impossible to produce a (cloned) fraudulent card.

Then one day the bank decided to stop offline verification and switch to online. This changed everything. The tapes with PINs in clear were suddenly a major weakness. With online verification, there was no link

between the card and the customer. In this case, possession of a copy of one or more embossing tapes would provide information to make copies of tens of thousands of genuine cards. Since the PIN was known, they were simple to use. Each card need only be used once, and the account details changed. The criminal is under no pressure to use the fraudulent card immediately as they are with a stolen card. The customer has the card in their possession, so no alarm is raised until the customer notices a transaction they did not make. This all takes time, and even then, the customer has the difficulty in convincing the bank that the transaction was not genuine.

There are two lessons from this example. The first is that evaluating a key control on the basis of a requirements specification alone is not valid. There is no guarantee that the final system will contain the control. To be certain, it is necessary to test that the system as delivered, contains the control. The second is that any change may introduce a weakness. Any security and control mechanism must be documented to show:

1. Why it is there.
2. What it depends upon for its effectiveness.
3. What other controls in turn depend on it.

When any change is made, it is essential that any control mechanisms are reviewed to determine whether the changes may make the controls ineffective.

Failure of 'end-to-end' security

My final example concerns another national funds transfer system used for both high and low value transactions. Like all such systems, great attention was paid to security with the use of strong encryption, and smart cards to control and manage encryption keys (that are divided into two parts to require collusion for their compromise). Like all such payment systems, it required three people to make a payment: one to input the details, a second to check and a third to release the payment. A proprietary database management system was used, and on the face of it everything was secure. Except that it was not! The application system used by the banks in the country concerned enforced the security rules. Messages sent between the banks and the central bank were encrypted. The database was held by the central bank, and it could only be accessed by the payments system program code. On some days in the year, a bank might have a very high volume of payments, and so they had a facility to input payments some days before they were due to be processed. They would be stored on the database as 'input' or 'checked', which meant that on the processing day the bank had only to release them.

The security measures contained in the payment system were very good, but the controls by which the database management system knew that it was communicating with the payment system were poor. It acknowledged only one user, namely the payment system, and it sent a password to the central bank to authenticate it. The developers building the payment system had taken no steps to conceal the password. Anyone reading the source code would find the password clearly labelled. The password was never changed because it would require an amendment to the program. What this meant was that if a criminal (or a member of staff) knew the password and caused the payment system to fail, they could use the internet to obtain access the database, enter the password and then be free to amend any payment details in storage, and to release the payment.

Conclusions to be drawn

We may not know what weaknesses exist in systems used by financial institutions. Indeed, it may be that in some country somewhere the same faults described in this article are still to be found. I have not seen changes in the way systems are built that gives me confidence that we are not making the same mistakes. I would go further, given the pressures on time and costs and the fragmentation of the development processes, I think it quite likely that insecure systems are still being implemented. I also think that anyone in the legal profession involved in cases concerned with financial systems should put every effort into analyzing all aspects of such systems and independently verifying any claims made about security. It requires a particular mindset, some might call it devious, to question how security measures fit together and whether there are weaknesses at the interfaces. There is one group of people who are likely to be doing this now: the criminals.

Since 1973, Ken has worked on the problems of implementing new systems and networking technologies without incurring unacceptable risks. He believes that developments in both these areas have created major changes such that many of the existing principles by which we control business operations are no longer valid.

