




The PCI DSS Appendix B: Compensating Controls

An IRM Research White Paper by
Michael Owen



IRM Research

Information technology constantly changes and advances. IRM is dedicated to keeping pace with new technology and continuing to innovate in the field of information security. This ensures that we are well informed of new issues and technologies, expanding our knowledge and providing world class services to our clients.

Executive Summary

The Payment Card Industry (PCI) Data Security Standard (DSS) is a standard for IT security mandated by a number of payment brands operating as the "PCI Security Standards Council." Initial members included American Express, Discover Financial Services, JCB, MasterCard Worldwide, and Visa International. The standard is a broad based security standard which describes a range of cross-business security functions and procedures which are required to ensure the security of payment cardholder data.

A key aspect of the PCI DSS is its focus on rendering cardholder data "unreadable," typically through the use of cryptographic mechanisms such as encryption or hashing. This requirement (requirement 3.4) has caught out a number of organisations, as data encryption is a much-neglected mechanism for securing data in databases.

Encryption tends to be discounted during design phases for a number of reasons including cost, speed to market, and the frequently seen issue of a "test or pilot system" becoming an operational system once a product foothold has been gained in a developing marketplace. Many organisations are now finding that to add encryption back into systems after many years of operations is unfeasibly expensive, and may in many cases require significant re-architecting of systems. Faced with potential for a complete redesign to be performed with funds that have not been allocated, many organisations are wondering how to deal with requirement 3.4 in a cost effective manner.

Fortunately, the PCI DSS version 1.1 includes a section called Appendix B: Compensating Controls. This appendix section describes a set of additional controls which are needed in organisations that have performed an analysis of the feasibility of data encryption, and determined that such encryption is not feasible. The controls take the form of four technical requirements, all of which must be met to allow compensating controls to be deployed in lieu of data encryption.

The requirements are all technical in nature, but can be summarised as being a set of control requirements which ensure that databases holding cardholder data are protected from an organisation's internal systems through the creation of a secure perimeter around the storage of the data. This is very similar to the secure perimeter an organisation employs to protect itself from the Internet. Many of the requirements can be met using existing technology and networking techniques which are generally already deployed within SMEs and major enterprises.

It cannot be said that Appendix B of the PCI DSS is a pain-free, get out clause for organisations which find it impossible to implement cardholder data encryption. There are a number of areas affected by the controls described, particularly with respect to basic connectivity and operational support. Careful planning and documentation is required to ensure that these changes do not negatively impact the operational processes which make use of cardholder data, and to ensure that the security benefits of the new controls are maximised. It is recommended that organisations perform a comparative cost-benefit analysis of the implications of performing data encryption and of implementing the compensating controls described in Appendix B before making a final decision on the strategy to pursue compensating controls over mechanisms which would render cardholder data unreadable. Implementers are strongly advised to seek the guidance of a formally qualified QSA before designing and implementing a solution, as a QSA will not approve any solution which does not entirely meet the requirements of this Appendix.

1. Introduction

The Payment Card Industry (PCI) Data Security Standard (DSS) is a mandatory requirement on all organisations that process credit card or debit card payments, develop products for payment card transactions and/or store payment card details. PCI DSS defines the requirements for payment card security, sets out the levels of compliance that organisations will need to meet and the way in which that compliance will be assured.

A key aspect of the PCI DSS is its focus on data encryption as a core mechanism for the security of cardholder data at rest. This requirement has caught out a number of organisations, as data encryption is a much-neglected mechanism for securing data in databases. Encryption tends to be discounted during design phases for a number of reasons including cost, speed to market, and the frequently seen issue of a new "test system" becoming an operational system once a product takes hold in the marketplace.

As a result, there is a clear need for either a large number of systems being completely redesigned or for some additional level of protection to be approved for remedial controls over a non-encrypting system.

Fortunately, the PCI Security Standards Council has anticipated this issue, and laid out supplementary controls which are required for organisations that cannot or will not meet the requirements around the encryption of cardholder data. These supplementary controls are laid out in Appendix B: Compensating Controls of PCI DSS version 1.1, under the heading Compensating Controls for Requirement 3.4, and will be explored in some detail in this document.

2. The Appendix

2.1 Qualification for Compensating Controls for Requirement 3.4

It should first be noted that the text around the compensating controls for requirement 3.4 state that the compensating controls cannot simply be applied as an alternative to making cardholder data unreadable. A risk analysis must have been performed to produce documented technological or business constraints which prevent the business from rendering cardholder data unreadable.

Assuming that such a risk analysis has taken place, the requirements can be considered;

Compensating controls may consist of either a device or combination of devices, applications, and controls that meet all of the following conditions:

1. *Provide additional segmentation/abstraction (for example, at the network-layer)*
2. *Provide ability to restrict access to cardholder data or databases based on the following criteria:*
 - *IP address/Mac address;*
 - *Application/service;*
 - *User accounts/groups;*
 - *Data type (packet filtering).*
3. *Restrict logical access to the database*
 - *Control logical access to the database independent of Active Directory or Lightweight Directory*
 - *Access Protocol (LDAP)*

4. Prevent/detect common application or database attacks (for example, SQL injection).

We will now examine each requirement in turn, discussing the nature of the requirement, any interpretation which might be necessary, possible technologies, and any potential issues to bear in mind.

2.2 Requirement 1 – Additional segmentation / abstraction

“Provide additional segmentation/abstraction (for example, at the network-layer).”

2.2.1 Interpretation

The objective of this requirement is an additional level of connectivity separation, placing a level of logical distance between the front-line systems and infrastructure holding the cardholder data. Cardholder data must not be directly and easily accessible by all users on the system network, and more crucially, the level of segmentation or abstraction must be higher than the default for the network and systems within the corporate environment. This type of control does not serve as a huge step up for security from deliberate attacks, but provides an additional layer of protection from accidental discovery or misuse.

In this context, we can take “segmentation” to mean logical compartmentalisation with appropriate controls such as VLANs with ACLs, whereas “abstraction” refers to the placement of an additional software-level proxy between two or more entities such as a client and a server, or a database and a web server.

2.2.2 Implementation Possibilities

For this particular requirement there is scope for performing either a separation or abstraction exercise, and we are not restricted to network level tools. It should also be noted that we are not restricted to only one mechanism or the other; the more controls employed in a compensating controls exercise, the more likely the signoff by your Qualified Security Assessor (QSA).

When considering separation, the most obvious strategy is one of network-layer separation, as mentioned in the PCI DSS requirement. For such a case, mechanisms would have to be deployed to hold cardholder data systems in a separate networking environment. This environment would at a minimum be a protected sub-network which contained only the cardholder data systems, and had restricted connectivity to other networks within the organisation.

For abstraction solutions, the key is to increase the degree of separation through the use of software level abstraction. If a system is currently able to make calls directly to the database, an additional layer of abstraction could be added through the use of middleware components. Configuring controls around the middleware and the database could ensure that for example, only the web servers could call the middleware in place, and the database would only accept calls from trusted middleware components.

2.2.3 Recommendations

When considering this requirement, it is important to select the mechanisms in the context of the current architecture of the systems being secured. Current data flows at both the middleware and network layers should be examined to determine where it is most useful to implement additional segmentation / abstraction; a security risk assessment considering both internal and external threats is an ideal technique to support the selection of appropriate technologies. Designers must keep in mind the potential impact of such changes, as a change at this level is intended to affect all connectivity to the databases, and thus the performance of payment processing systems could be affected.

2.3 Requirement 2 – Ability to restrict access to cardholder data based on additional criteria

"Provide ability to restrict access to cardholder data or databases based on the following criteria:

- *IP address/Mac address*
- *Application/service*
- *User accounts/groups*
- *Data type (packet filtering)"*

2.3.1 Interpretation

The objective of this requirement is to place access controls in place for data-containing systems which are above and beyond the standard networking controls in place for other systems. Note that there are two layers being discussed in this requirement – there are both network and application level criteria to consider. Crucially, this network-level control is focused on the databases themselves, and not the interior of the enterprise as a whole.

This is not to imply that the main perimeter of an organisation should be neglected – the key focus here is to establish a new perimeter internally, protecting databases holding cardholder data from internal threats in addition to external threats.

2.3.2 Implementation Possibilities

Restriction based on "User Accounts / Groups"

This is quite possibly the most straightforward of the requirements to be met. For the purposes of user-based controls, user accounts and user groups must be used both in the OS and database to ensure that people can only view information for which they have a documented need.

Restriction based on "IP address / MAC address," "Application / Service," and "Data type (packet filtering)"

This section clearly states a requirement for network-level access controls to be implemented. A variety of mechanisms can be employed, including firewalls and Network Access Control (NAC) technologies such as 802.1x.

Caution is advised in the selection of lower-end mechanisms, as simply port-filtering will not likely satisfy this requirement. While applications tend to be tied to a given port, this is a matter of convenience and standards, as opposed to a matter on which to base security. Attackers quite commonly tunnel connections over conventionally allowed ports to transmit data, and a port-filter will not detect this.

2.3.3 Recommendations

Choices for this requirement are more limited, however selection of products and vendors should still be considered. The trade off between ease of administration with a known vendor and the additional security of a secondary security vendor must be considered, along with the impact of new security technologies on the day-to-day operational load of administrators. To maximise both security benefits and the likelihood of having mitigating controls approved, strong reasons must be found for avoiding a second firewall vendor.

As with any firewalling or network level access control exercise, a data flow examination must take place before any implementation takes place. Placing new network and application level controls in place around a

live system is a delicate operation, and the business must be entirely certain that all relevant network traffic has been identified, accounted for, and included in the relevant firewall configurations.

There will be a significant need for policies and procedures covering the independent user administration of the cardholder databases. Care must be taken to ensure that when the switchover from centralised user administration to local administration is performed, key operational tasks are not impacted.

Additionally, standard operational tasks such as backups and log gathering will likely need re-engineering or at the least validation to ensure they will continue to function in the new user model.

2.4 Requirement 3 – Restrict logical access to the database

"Restrict logical access to the database

- *Control logical access to the database independent of Active Directory or Lightweight*

Directory Access Protocol (LDAP)"

2.4.1 Interpretation

The goal of this requirement is straightforward, though the wording is deceptive; while placing controls around the database to control access at a user level is already defined as a component of requirement two, this requirement specifies that access must be managed locally, as opposed to using one of the various centralised user management directories. The wording implies that all directories are thus inadequate, when in fact, a local dedicated AD or directory will meet this requirement, assuming it is sufficiently secured.

As central directories must allow large amounts of traffic to pass between themselves and other systems on the network, they are fundamentally difficult to secure, and are generally not to be trusted across security perimeters.

With the establishment of a new security perimeter around the databases, it is therefore no longer appropriate to trust the centralised user databases sitting outside this domain.

2.4.2 Implementation Possibilities

The implementation possibilities for this particular requirement are limited. Accounts used to control logical access to the database must be held local to the server(s) on which data is held, or in a dedicated AD/Directory server which is segregated from out of scope systems. Trust relationships which place trust on externally administered accounts cannot be used.

2.4.3 Recommendations

Depending on the user access control model in place, this may be one of the more onerous requirements for complying with the PCI DSS, and could well be the requirement which pushes system designers to consider the possibility of applying mechanisms to satisfy requirement 3.4 of the DSS more carefully. The use of local accounts or a dedicated directory will generally be a move away from the increasing use of centralised user administration.

It will place greater overhead on user management, and will have an impact on the efficiency of operational tasks.

Documentation is needed around this should the shift to local accounts be a significant one for an organisation; the procedural controls around databases must be updated, and operational tasks and scripting updated to reflect changes to the user model. Password policy enforcing controls must be reviewed

to ensure that the database controls reflect the corporate password policies as well as meeting the password requirements of the PCI DSS.

2.5 Requirement 4 – Prevent/detect common application or database attacks

"Prevent/detect common application or database attacks (for example, SQL injection)."

2.5.1 Interpretation

The most straightforward of the requirements in Appendix B, requirement four can be taken as it reads. Controls must be put into place to prevent or detect common application and database attacks. It should be noted that the "or" in the requirement should be interpreted as an inclusive, rather than exclusive "or."

2.5.2 Implementation Possibilities

Two main strategies are mentioned here; detection, and prevention. While this does not explicitly declare that both are needed, it is generally recommended that both strategies be pursued when attempting to implement compensating controls for encryption.

A prevention model for attacks could rely on a number of strategies taken from two key areas; technological controls, and procedural controls.

In terms of technological controls, a range of preventative controls could be considered. The most obvious choice would be application layer firewalling, however a range of other options such as network-based Intrusion Prevention Systems (IPS), or Host-based IPS (HIPS) systems can also be considered.

Procedural controls are not always an obvious mechanism, however they can compliment technological controls nicely. Code reviews are an ideal procedural control for the prevention of application and database attacks, removing vulnerabilities before they require protection by a technological control. While it would not be wise to rely solely on procedural controls, they add an additional layer of protection on top of protective technologies such as IPS.

Similarly, detection offers a range of possibilities. Both host-based and network-based IDS systems are an obvious means to detect these attacks when they are in progress. Further possibilities include technologies such as Security Event Management (SEM), a breed of tool which assembles views of security incidents by assembling and correlating security logs across a large number of systems.

2.5.3 Recommendations

While the selection of a mitigating technology for this requirement is too company-specific for a comprehensive recommendation to be made in this paper, a number of generalised recommendations can be made with regards to the technologies available, and the future of those technologies.

Firstly, there is the potential impact of IPS systems, either host-based or network-based. IPS systems are recognised as providing a greater level of security through the preventative mechanisms that they can provide, effectively allowing them to terminate suspect network or system level activity when certain internal criteria are met. This increased functionality does lead to potential concerns over automated decisions impacting on system functionality, such as when a required system is banned from connecting to a server due to a trigger in the IPS ruleset.

Secondly, data flows must be carefully considered before selecting any network-based mechanisms, including firewalls. The ability to act on network traffic is entirely dependent on the ability to actually see the contents of network traffic, and this is increasingly less common. Indeed, as soon as an organisation is using appropriately secured and encrypted protocols for administrative access, network level controls such

as network IPS and IDS systems are becoming more difficult to maintain. Hardware SSL terminators and encryptors can be considered to provide an unencrypted segment for network traffic monitoring, but costs must be carefully considered when assembling such a solution.

Thirdly, it should be remembered that this decision is not taken in a vacuum. The potential to leverage existing mechanisms or incorporate this requirement into other requirements to select a multipurpose appliance or solution must be considered. Additionally, the potential for other controls to limit the benefits of a mechanism must be considered. Taking the example of a NIPS system, network level IPS will be of little use with a set of databases already only able to accept connections from one set of hosts such as middleware servers. For cases where it can be identified that a mitigating control will bring all system activity to a halt, finer-grained controls (such as application level firewalls, in the case of an IPS) must be found.

3. Conclusion

The PCI DSS was put forward by the PCI Security Standards Council to raise the bar on cardholder data security across the industry. The standard spells out the basic requirements for security controls across systems holding, processing, or passing cardholder data, and includes a requirement to render cardholder data unreadable on systems storing such data. This requirement has caused some issues for vendors, as the use of such technologies (most typically, encryption technologies) is relatively uncommon.

Fortunately for vendors, there is an alternative set of mechanisms in the form of Appendix B of the PCI DSS version 1.1. This appendix lays out a number of technical requirements for systems which cannot satisfy the requirements of 3.4, describing a number of technical access control mechanisms and criteria.

The requirements effectively constitute the formation of a new security perimeter around databases containing cardholder data. This perimeter is much like the perimeter protecting an organisation from external attacks, but is intended to protect databases from any internal threats as well as providing additional protection from external attacks.

Control selection will require careful evaluation of the systems and data flows involved to maximise benefits while minimising the impact on the organisation as a whole. The requirements must be considered as a whole to ensure that appropriate selections are made to satisfy all requirements in a cost effective manner.

And finally, the new perimeter and its associated controls will require careful documentation both for compliance, and to minimise the operational impact of changes. The changes being made directly affect the use and maintenance of the databases in question, and thus all changes must be evaluated and detailed to ensure that operation tasks are not affected, and all live connectivity has been taken into account. This may well involve a level of reconfiguration of systems as well as some scripts associated with operational tasks.

The use of compensating controls from Appendix B provides organisations with an alternative to the use of data obscuring technologies such as encryption. These controls are not to be considered a quick and easy "get-out" clause for requirement 3.4 of the PCI DSS, and careful consideration is needed before they can be deployed. With appropriate documentation, planning, and design work, they can however provide a solid level of protection for cardholder databases which would otherwise be holding unprotected data.

IRM provides a range of services relating to the PCI DSS, including DSS health checks, documentation and policy consultancy (including policy and procedure writing), risk analysis / management, and technical consultancy for systems design and implementation. IRM is a PCI qualified QSA.

About the Author

Michael Owen is a Senior Security Consultant and Information Risk Management Plc, where he acts primarily as a risk and security management consultant, including performing PCI DSS health checks and providing advice on compliance programs. Before joining IRM, Michael worked with Egg plc, the online bank, where cardholder data security was a critical part of his work.

About IRM

Information Risk Management Plc (IRM) is a vendor independent information risk consultancy, founded in 1998. IRM has become a leader in client side risk assessment, technical level auditing and in the research and development of security vulnerabilities and tools. IRM is headquartered in London with Technical Centres in Europe and Asia as well as Regional Offices in the Far East and North America. Please visit our website at www.irmplc.com for further information.