The Effects of Breaches: A Study of Merchant Programs

Branden R. Williams

Abstract

Breaches happen on a global scale with serious impact to many of the firms who suffer in a post-breach way of life. The media glamorizes hacking as more prevalent than petty theft while payment brands stand behind the Payment Card Industry Data Security Standard (PCI DSS) as the solution to keep merchants safe. With PCI DSS turning 10 in 2014—the same year which saw multiple high-profile breaches—members of the Merchants Acquirers' Committee wanted to understand how breaches financially impact both merchants and acquirers. This paper reports on primary research performed on the MAC members, specifically aimed at understanding the financial impact of breaches among merchants and acquirers. *Please note, components of this research appear in an industry whitepaper authored by Branden R. Williams as released in partnership with MAC.*

1 Introduction

Information security breaches are more of a reality for firms than they ever have been. Breaches involving payment cards were especially prominent in 2014. While we can discover how these breaches affect large, public companies, there is less data available on small to medium businesses. The literature is full of research on security breaches to varied degrees; however, there is not any specific research that attempts to look at the Level 3 and 4 communities of merchants. This research specifically aimed to uncover connections between different portions of the merchant community, as well as to discover more of the financial affects of breaches among these entities.

High profile merchant breaches over the last year such as Target, Micheals, Sally Beauty, and PF Changs clearly show that breaches are happening to large merchants. According to Trustwave and Verizon, breaches do happen to smaller companies too (1,2). Because of industry reports, we know investigations happen at all sizes of companies; but studying the effects of breaches on smaller companies cannot be done unless the affected firm is public. We await reports from both companies that will overlap in time with a significant period during which this research was conducted to add further color to these findings.

The goal of this emperical research was to gain understanding of the economic impacts of payment card breaches in the merchant community. The remainder of the paper will review the methodology, results, discussion, and future research. A bulleted list of high-level findings is in the final section.

2 Methodology

This research was conducted using a proprietary survey instrument (n=25) that went through a number of revisions as the goals of the research were refined and clarified. In addition, certain questions were dropped as the results would overlap with the existing studies mentioned above. A select group of MAC members piloted the survey prior to sending it out to the MAC membership to help with question interpretation issues and to boost validity and reliability. The final instrument was coded into SurveyMonkey and sent out to the MAC membership with the goal of understanding more about the consequences of breaches in the Levels 3 and 4 merchant communities.

One challenge with the research was the lack of complete responses from the 100 total respondents—only 25 could be used for this analysis. The survey instrument asked fairly complex and specific questions, which may have lead to the large number of abandoned surveys. One response was dropped, even though it was complete, as the merchant counts were clearly reversed (Level 4 numbers were reported as Level 1). The response did not make it into the final data analysis to avoid any additional skew based on the researcher altering parts or all of the response.

Although the data collected in the survey did not pass normality tests (with one exception), all of the tests performed in this analysis have certain robustness against non-normal data sets. As a disclaimer, even those tests that have robustness against non-normal distributions may generate skewed results that produce Type I and Type II errors.

3 Results

The results comprise of two general areas of analysis of the data. The first focused on breach trends and compliance rates among the different levels and the second analyzed the effects of managing merchant populations in compliance programs. The empirical analysis of the data does not take into account business goals, executive motivation, or penalty-based revenue streams (and their potential decline). The Discussion section of this paper will attempt to connect business constraints with some of the findings in the research.

3.1 Compliance by Level

Visa reports on compliance rates in their various merchant levels and publishes this information on their website (3). The compliance rates as reported by the survey respondents was considerably lower for Levels 1 & 2 merchants when compared to Visa's latest report of 97% and 88% compliance for Levels 1 & 2, respectively. This anomaly could be explained by acquiring institutions reporting certain merchants as compliant, therefore accepting any risk on their behalf in the event of a breach. From Visa's perspective, advertising high compliance rates with PCI DSS would demonstrate that the programs are successful and that the merchant and acquiring communities have accepted the standard. Contrary to some of the figures in the latest report from Visa, compliance rates are still relatively low across all four levels (see Figures 1 and 2).

The responses indicate that while compliance efforts are under way, Acquirers still have quite a bit of work to do to boost these rates—if, in fact, high compliance rates are a desired outcome. This may not be the case where organizations have come to rely on penalty-based revenue streams to meet financial obligations to their shareholders. Non-compliant Level 4 merchants with penalty-based revenue streams are more profitable than a population with 100% compliance. In addition, there is no penalty to acquirers for non-compliant Level 4 merchants like there is for Levels 1-3 from multiple payment brands.

Aggregate compliance rates as reported by the respondents do not exceed 70% in any category, with Level 2 merchants showing the highest achieved compliance rates.

3.1.1 Is PCI DSS Effective?

Perhaps one of the larger questions we should be asking is how effective PCI DSS is in preventing breaches. Previous to the breach epidemic we saw in 2014, experts seemed to argue either perspective. Spokesmen for the PCI Security Standards Council often state that no breached company was ever PCI DSS compliant at the time of their breach–thus implying that companies that comply will be safe from breaches. But recent thought by key industry influencers seems to suggest that PCI DSS isn't effective in stopping modern hackers (**???**). Perhaps a more telling grouping of data points to consider is the combination of breach activity, less than perfect compliance rates, and modern hacking techniques. PCI DSS and the ecosystem that supports it, in its current state, does not appear to be adequate to combat malicious actors. Not only are compliance rates low, but breaches happen even when companies appear to comply with PCI DSS (2).

There is no question that PCI DSS has, through brute force, improved the security programs of anyone who handles payment card information. Given its inability to quell recent breach activity, however, its value seems to have diminished dramatically over the last five years. Merchants may be better off outsourcing all of their payment processing and abandoning PCI DSS in favor of more robust security frameworks, such as ISO 27000. As with all security frameworks, companies will only earn a benefit from the program with

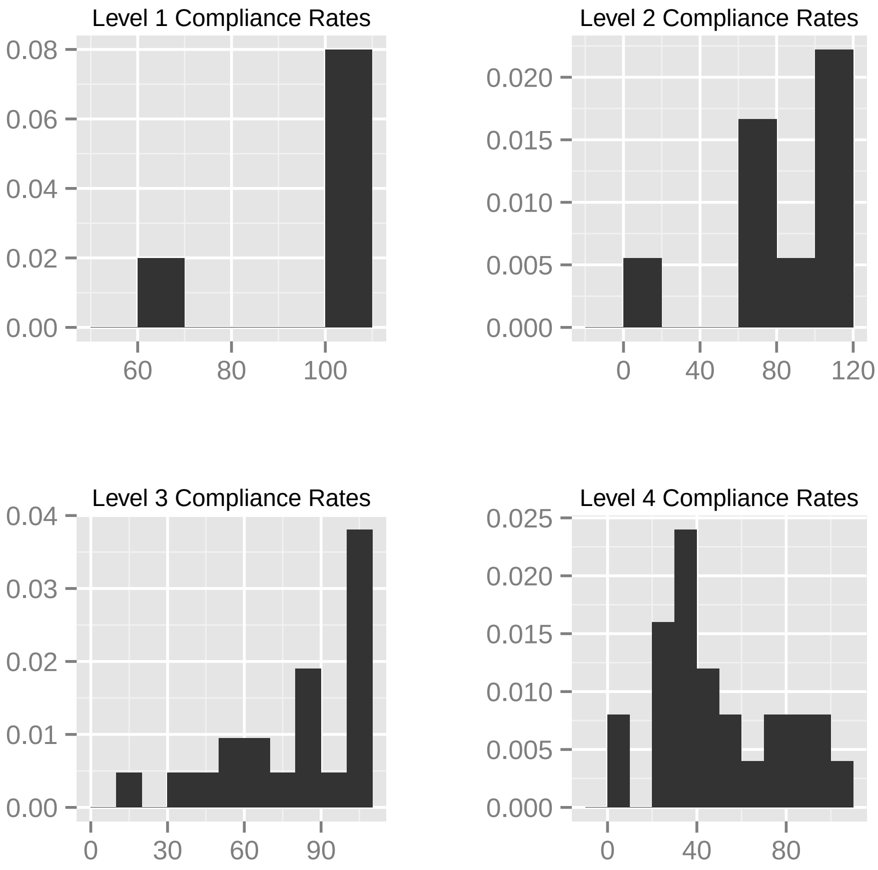

Figure 1: Compliance Rate Density, by Level.

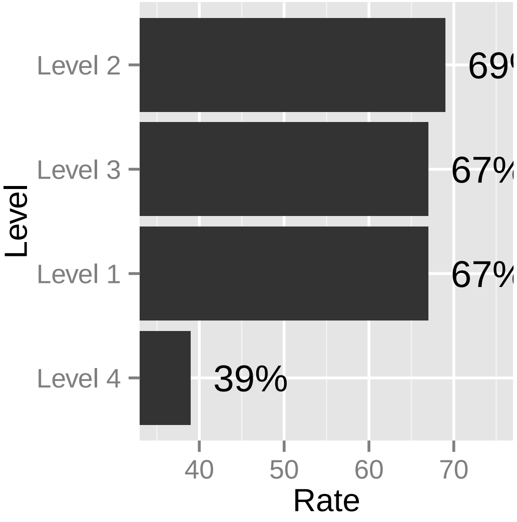

Figure 2: Compliance Rates, by Level.

an honest emphasis on excellence in information security. Firms that only go through the motions and put forth a minimum effort waste their time and money.

3.1.2 Program Analysis

There are challenges with performing correlation analysis on the data gathered in the survey. First, only the Level 4 compliance rates passed a Shapiro-Wilk normality test–or a test to determine if the collected data would be part of a normal distribution. Neither the Pearson nor the Spearman correlation methods explicitly require normal distributions; however, anti-normal tendencies can skew some of the results. The Spearman rank correlation method is acknowledged as a more robust correlation method for non-normal distributions (4). Both analyses methods are below for the sake of transparency.

Table 1: Pearson Correlation of Compliance/Breach Values.

	L1Comp	L2Comp	L3Comp	L4Comp	L1Breach	L3Breach	L4Breach
L1Comp							
L2Comp	0.686***						
L3Comp	0.228	0.189					
L4Comp	0.121	0.399*	0.187				
L1Breach	0.246	0.168	-0.125	-0.093			
L3Breach	0.442*	0.327	-0.101	-0.096	0.796***		
L4Breach	0.020	-0.038	0.190	-0.127	-0.047	0.098	
RepeatBreach	0.120	0.040	0.209	-0.015	0.258	0.161	0.859***

Table 2: Spearman Correlation of Compliance/Breach Values.

	L1Comp	L2Comp	L3Comp	L4Comp	L1Breach	L3Breach	L4Breach
L1Comp							
L2Comp	0.696***						
L3Comp	0.178	0.087					
L4Comp	0.148	0.434*	0.148				
L1Breach	0.325	0.171	-0.173	-0.113			
L3Breach	0.331	0.167	-0.246	0.043	0.532**		
L4Breach	0.149	0.098	-0.085	-0.172	0.144	0.524**	
RepeatBreach	0.369	0.218	0.201	0.023	0.527**	0.199	0.417*

An unexpected finding was that none of the respondents reported any Level 2 merchant breaches in the prior year. After closing data collection, at least one Level 2 breach became public. Therefore, it is not possible to conclude if there is correlation due to e-commerce activity in that level.

Please note, that since the respondents reported *zero* Level 2 breaches, that category does not exist in the correlation analysis is it yields a divide by zero error. For the following tables, the asterisks indicate levels of significance. Single asterisks include p-Values < 0.05, double are p-Values < 0.01, and triple are p-Values < 0.001.

What these data indicate is a strong positive correlation between the Level 1 and 2 compliance rates in the respondents–meaning that the groups move in the same direction at the same rate. Given the initial focus on Level 1, then Level 2 merchants, the data suggests that acquirers manage those programs very closely and achieve similar compliance rates between the two groups. This intuitively makes sense as those groups represent the vast majority of processed transactions. In addition, there is weak to medium positive correlation between Level 2 and Level 4 compliance programs. This many manifest itself as a byproduct of card-present compliance programs.

The plots below visually depict the various correlated values with a linear model estimate and confidence bars.

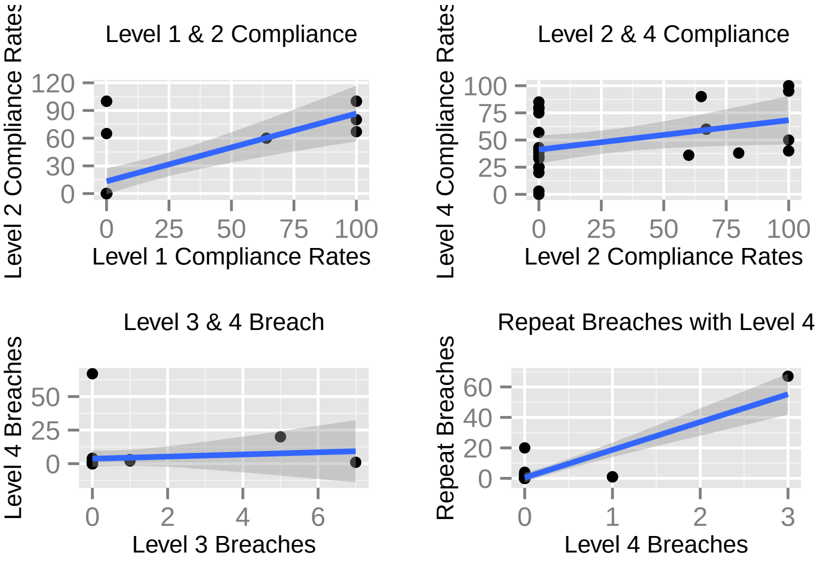

Figure 3: Correlation Graphs, with Linear Regression Estimates.

Some of these models may not be highly accurate given the width of the confidence area around the linear regression. In fact, the correlation between Level 4 breaches and repeat breaches looks suspiciously like the 4th panel of Anscombe's quartet (5). In fact, all of these variables have extremes with respondents giving 0 or 100% compliance rates in some areas. These graphs are simply for illustrative purposes. Future research should focus on coding the instrument to yield a larger number of complete responses in these areas, and eliminating the possibility for a zero response to be considered in the data.

The Pearson correlation matrix suggests that there is a strong positive correlation between Level 1 breaches and Level 3 breaches. A linear regression suggests that either Level 1 or Level 3 breaches could be seen as a predictor of the other (p-Value of 0.0000019). The common factor between the two groups is an E-Commerce, which could suggest that these kinds of breaches represent a significant portion of the respondents.

The Spearman correlation matrix confirms the above, but also suggests significant (but not strong) positive correlation between Level 3 breach data and Level 4 breach data. This could indicate that breaches may have common elements in those communities—possibly due again to e-commerce related breaches. Thus, acquirers should look for common elements in those breaches and deploy preventative (or detective) measures to quell the activity.

Keeping the discussion related to the visual appearance of these plots in mind, perhaps one of the more telling correlations is that of Repeat Breaches and Level 4 breaches. This data indicates that Level 4 merchants may be the most likely to see a repeat breach within twelve months of the initial breach. A linear regression analysis of Level 4 breaches being a predictor of Repeat Breaches indicates a strong positive correlation. Essentially, if you have a high number of Level 4 breaches, there is a high likelihood that one or more of those will be breached again within 12 months ($r2 = 0.7382387$).

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	0.03325	0.07058	0.4711	0.642
L4Breach	0.04047	0.005025	8.054	0.00000003821

One final result of the correlation analysis suggests that Level 1 breaches may be associated with repeat breaches. What this may indicate is that acquirers that have merchants who suffer more than one breach in a 12-month period may be more likely to see their Level 1 merchants also suffer a breach. This could be an indicator of a weak compliance enforcement program and overall low compliance rates. There were no statistically significant results from a linear regression of those two data points, and the data did not distinguish between the kind of repeat breaches that happened—meaning, the data does not tell us which category of merchant a repeat breach came from.

3.1.3 Program Management

A final outcome of this study is further understanding of how compliance rates were affected based on how compliance programs are managed. The data has a clear trend that suggests that a larger number of merchants in the program will yield a lower overall compliance rate. The systems and programs currently in use to manage large merchant populations may be inadequate to promote higher compliance rates. Alternatively, there may not be a true desire to promote higher compliance rates in the event that a non-compliant merchant generates revenue in the form of fines collected on a monthly or yearly basis. In the case of "Other" in the plot below, the respondent mentioned they use a combination of outsourcing and in-sourcing.

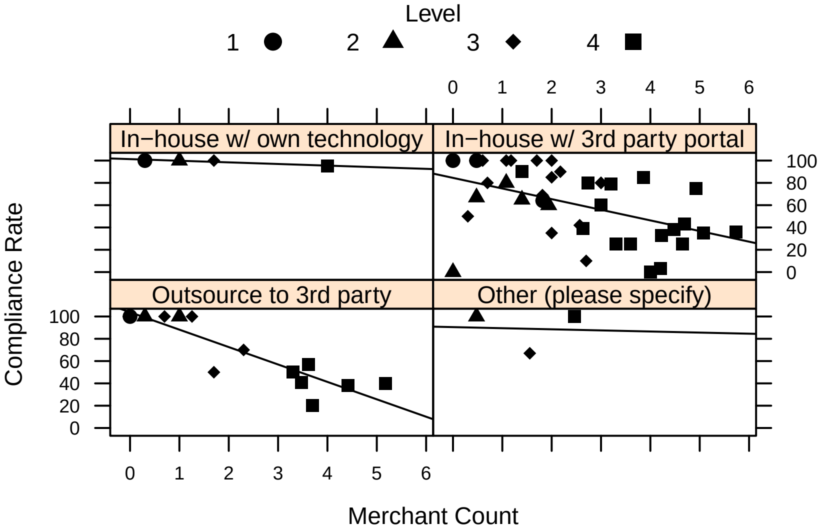

Figure 4: Compliance Rates by Merchant Count, Grouped by Merchant Program Operations.

Large populations are largely found in the Level 4 community. Given that Level 4 merchants have the lowest overall compliance rates, it is not surprising that Level 4 merchants would drag the trend downward. If we expand by level and look at compliance rates, it appears the more merchants present in the program the lower the overall compliance rates (Level 2 merchants being the exception).

3.2 Effects of Breaches

One of the key goals of this survey was to understand how bad the breach problem actually is for smaller merchants. With 119 total breaches reported, only two respondents reported on fines assessed. Using the data associated with just those two respondents, breaches on average are coming in at just under $18,500 per incident. *Please note, this information is for illustrative purposes only. Under no circumstances should any risk-based calculation could be made such that any breach would be subject to that arbitrary average amount.*

One piece of data that is useful when viewing this problem is how transaction volumes are impacted by breaches. The majority (69%) reported unknown changes in transaction volumes, while 27% reported no

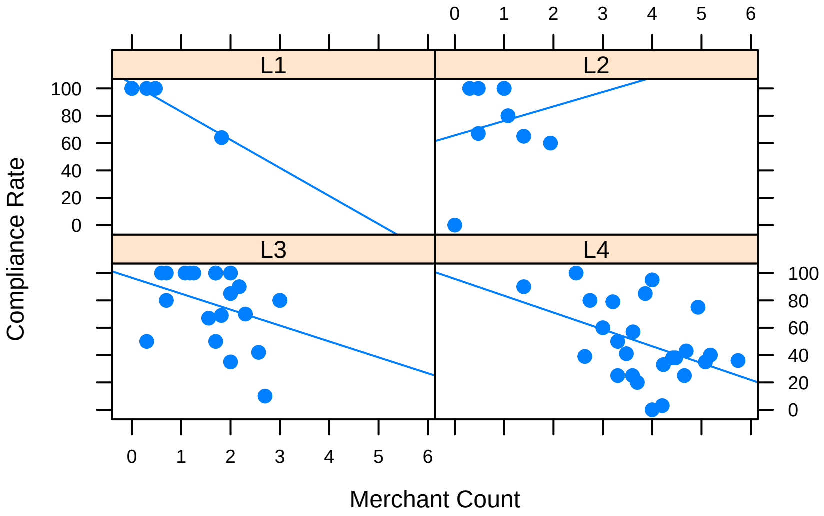

Figure 5: Compliance Rates by Merchant Count, Grouped by Level.

change at all. Unknown transaction changes could be related to a lack of data, visibility, or simply no metrics to track it because it is assumed to remain steady. Only 4% reported a decline.

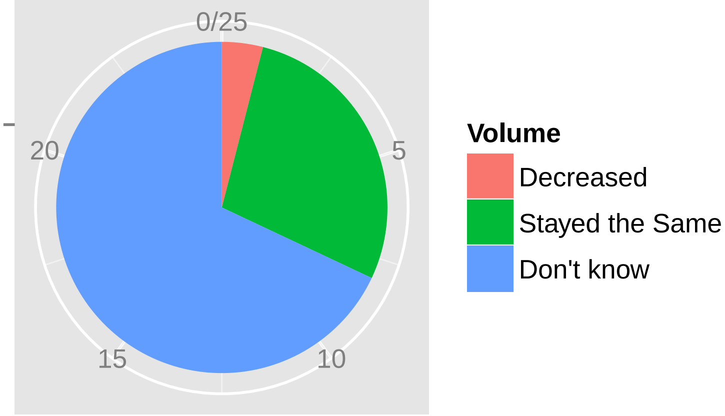

Figure 6: Post Breach Transaction Volume.

What this data suggests is that the general public does not change their shopping behaviors at breached merchants. Given that breached merchants continue to have customers coming in the door and paying with payment cards, there is little motivation to change behavior based on the fear of declining sales. If the trend was strongly negative, acquirers would have strong incentives to invest money in technology and security aimed at preventing breaches as acquirers depend on transaction volumes for revenue.

Our respondents reported a total of 119 breaches in all levels, and of those breaches, only 6 were actually terminated (or about 5%).

4 Discussion

This research focused on the MAC membership as a basis for analysis as it provides a great cross-section of US card processing market without bias to payment card brand or forensic investigator. The main bias is limited to MAC membership (over non-MAC members) and the limited responses received.

Several interesting findings came from this data analysis. First, the breach problem does not appear to be as severe as perceived or as advertised (6,7). Level 1 breaches make headlines and have significant impacts to those companies; however, of all of the merchant counts reported in the survey (1,144,681), only 119 were reported as having a breach, with 5 of them having more than one breach in the previous twelve months. This equates to 0.01% of the total merchant count suffering a breach.

Is any one level of merchant more likely to suffer a breach over another? A common test to check for this is called an Analysis of Variance (ANOVA). Running this test across the three groups reveals no statistically significant difference between the breaches as reported by level, regardless if Level 2 breaches are included in the analysis. A p-Value of 0.1145216 supports no statistically significant difference among which group is more likely to be targeted over another. ANOVA has robustness against non-normal distribution. There is substantial literature supporting the robustness of ANOVA in situations where the normality assumption is violated (8–12). Further analysis using Tukey's method does not reveal any particular interactions among the four groups as significant, further strengthening the ANOVA results.

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Level	3	290.8	96.92	2.032	0.1145
Residuals	96	4579	47.69	NA	NA

Does a lack of variance between the groups suggest that any one group has a higher probability of suffering a breach? The Level 4 population suffered the highest count of breaches, but due to the large Level 4 populations, they were the least affected by breaches when looking at the size of the population. If you have a level 4 merchant population, there is a high probability that any given breach will be from that population (86.6%) with Level 3 merchants coming in second (11.8%). In addition, Level 4 communities come with the highest probability of an acquirer dealing with a breach at all (40%). Level 3 merchants come in second at 16%.

4.1 Implications for Practice

Breaches do happen, but their hype may not be nearly as big of an issue as it seems. Merchants are responsible for their own security in the PCI DSS ecosystem. If a firm chooses to accept payment cards, its managers must be prepared for attacks—some of which may result in a breach. While this may seem to be common sense, informal polls of merchants in the US and Western Europe suggest that merchants believe their banks are responsible for keeping them safe.

PCI DSS is incredibly complex. Most merchants do not understand the inner workings of the standard, how it applies to them, and how to ensure their technology partners are keeping them safe. If the stated goal of the payment card providers is to reduce breaches, incentive structures must change. As an example, taking the economic concept of Least Cost Provider in detecting and addressing fraud, Issuers would be the most optimal player to shoulder the burden of fraud prevention (13). That said, acquirers are best positioned to stop merchant fraud and terminal or POS misuse. Acquirers could affect the breach problem by investing in technology that protects merchants while they process payment data.

Merchants should consider outsourcing all of their payment processing to a third party and devote their information security efforts to defending the sensitive information they collect and use. Stolen payment card information is still fairly easy to monetize, thus any firm who chooses to handle this data in a way where they can negatively affect the security of the information is a target. The build-or-buy business cases around plastic card processing largely deployed in the 1980s do not apply anymore with information technology becoming increasingly open and connected. If refactored using today's dollars, more companies will find that outsourcing is ultimately cheaper than insourcing and staffing a large security and fraud department to maintain the systems.

Given the relatively low number of breaches and the small amount of fines assessed, acquirers and processors have no incentive to quell breaches through proactive measures. It's simply too easy to either take the losses, self-insure (premiums for current offerings clearly outpace losses), or pass them along to merchants instead of proactively working to address the issue. Firms looking to reduce their risk of a breach should focus on outsourcing as the current incentive structure stacks the deck strongly against the merchant. Without advanced knowledge of the risks in operating their payment systems, merchants may over spend in areas for an event that does not have a measured liklihood of occuring.

Acquirers and processors may consider investing in tools that effectively remove the merchant from scope to reduce breach incidents and significantly mitigate the risk from payment processing. Merchants may see this as a reason to continue their current processing relationship, or it potentially could be a feature that gives an acquirer competitive advantage.

5 Future Research and Summary of Findings

The biggest limitation of this study was the low number of responses to the survey instrument inside the MAC membership. Future research may focus on specific components of this analysis and develop standard data collection mechanisms that effectively remove the need of an instrument. Further examination on the connections between repeat breaches and merchant breaches, by level, could yield better predictors on where additional risk lies.

One area of study that is missing in this research is the affects that the Account Data Compromise (ADC) and Common Point of Purchase (CPP) requests (and others like them) affect resource allocation among acquirers. If this activity is increasing with ever decreasing batches of cards, acquirers may see significant resource drain in their operations while completing these requests.

This research yielded several interesting findings that have impact for practitioners and scholars alike. In summation, the high level findings from the research were:

- Compliance rates, by level, are lower than many entities suggest (3).
- Level 1 and Level 2 program success seems to be strongly positively correlated.
- Level 1 and Level 3 breaches share correlation, positively due to e-Commerce connection.
- Repeat Breaches and Level 4 breaches share correlation.
- There is no one group more likely to be breached than another.
- A larger merchant population is directly correlated with lower compliance rates (in most cases).
- Breaches and fines are relatively small and localized.
- A lack in transaction level change, post breach, indicates that consumers do not care about breaches enough to change their spending habits.

6 About the Author

Branden R. Williams, DBA, CISSP, CISM, has almost twenty years of experience in technology and information security both as a consultant and an executive. His specialty is navigating complex landscapes—be it compliance, security, technology, or business—and finding innovative solutions that save companies money while reducing risk and improving performance. Along the way Dr. Williams was Director of the PCI Consulting Practice for VeriSign, a CTO at RSA, and served on the PCI Board of Advisors. He is a co-author of three books on PCI Compliance, and his other publications can be found at www.brandenwilliams.com.

References

1. Trustwave Holdings Inc. 2014 TRUSTWAVE GLOBAL SECURITY REPORT executive summary [Internet]. Trustwave Holdings, Inc. 2014 p. 123. Available from: http://www2.trustwave.com/rs/trustwave/images/2014/_Trustwave/_Global/_Security/_Report.pdf

2. Verizon Business. 2014 data breach investigations report. Verizon Business Journal [Internet]. 2014;2014:1–60. Available from: file:///C:/Users/Edward S. Forde/Downloads/rp_Verizon-DBIR-2014_en_xg.pdf

3. Visa I. U.S. pCI dSS compliance status [Internet]. Foster City, CA; 2014 p. 1. Available from: http://usa.visa.com/download/merchants/cisp-pcidss-compliancestats.pdf

4. Bishara AJ, Hittner JB. Testing the significance of a correlation with nonnormal data: Comparison of Pearson, Spearman, transformation, and resampling approaches. Psychological Methods. 2012;17(3):399–417.

5. Anscombe FJ. Graphs in Statistical Analysis. The American Statistician. 1973;27(1):17–21.

6. Brennan R. Businesses too lax with data, says privacy commissioner; Identity theft feared as watchdog cites 'epidemic' of breaches. Toronto star. Toronto, Canada, Toronto; 2008.

7. Mirabella L. Retailers face threat from cybercriminals. The buffalo news. Buffalo, NY; 2014.

8. Wilcox RR. How many discoveries have been lost by ignoring modern statistical methods? American Psychologist [Internet]. American Psychological Association; 1998;53(3):300–14. Available from: http://library.capella.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true/&db=pdh/&AN=1998-00122-003/&site=ehost-live/&scope=site

9. Box GEP. Non-normality and tests on variances. Biometrika. 1953;40(3-4):318–35.

10. Glass GV, Peckham PD, Sanders JR. Consequences of Failure to Meet Assumptions Underlying the Fixed Effects Analyses of Variance and Covariance. Review of Educational Research [Internet]. 1972;42(3):237–88. Available from: http://rer.sagepub.com/content/42/3/237.short

11. Lindquist EF. Design and analysis of experiments in psychology and education. 1953;

12. Boneau CA. The effects of violations of assumptions underlying the t test. Psychological Bulletin [Internet]. American Psychological Association; 1960;57(1):49–64. Available from: http://library.capella.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true/&db=pdh/&AN=1960-06756-001/&site=ehost-live/&scope=site

13. Levitin AJ. PRIVATE dISORDERING? PAYMENT cARD fRAUD lIABILITY rULES. Brooklyn Journal of Corporate, Financial & Commercial Law [Internet]. Brooklyn Law School; 2010;5(1):1–48. Available from: http://ezproxy.library.capella.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true/&db=bth/&AN=62170734/&site=ehost-live/&scope=site