verizon

2013 DATA BREACH INVESTIGATIONS REPORT

A global study conducted by the Verizon RISK Team with cooperation from:

AFP Australian Federal Police

CERT | Software Engineering Institute Carnegie Mellon

Consortium for Cybersecurity Action

CyberSecurity Malaysia An agency under MOSTI

Deloitte.

EC3 Europol

ES ISAC Electricity Sector Information Sharing Analysis Center

Danish Defence Intelligence Service

FE

G-C Partners

Guardia Civil

Homeland Security

IRISS Irish Reporting and Information Security Service

POLITI

POLITIE

ThreatSim™

2013 DATA BREACH INVESTIGATIONS REPORT

"Some organizations will be a target ***regardless*** of what they do,
but most become a target ***because*** of what they do.

If your organization is indeed a target of choice, ***understand as much as you can*** about ***what your opponent is likely to do*** and ***how far they are willing to go***."

– 2013 DBIR, pg. 48

2013 DATA BREACH INVESTIGATIONS REPORT

TABLE OF CONTENTS

Introduction 4

Methodology 8

Results and Analysis 11

Demographics 13

A4 Threat Overview 16

Threat Actors 19

External Actors (92% of breaches) *20*

Internal Actors (14% of breaches) *23*

Partner Actors (1% of breaches) *24*

Threat Actions 25

Malware (40% of breaches) *29*

Hacking (52% of breaches) *34*

Social (29% of breaches) *36*

Misuse (13% of breaches) *38*

Physical (35% of breaches) *40*

Error (2% of breaches) *41*

Environmental *41*

Compromised Assets 41

Compromised Data 44

Attack Targeting and Difficulty 47

Breach Timeline 50

Discovery Methods 53

Conclusions and Recommendations 56

Recommendations for Mitigating Highly Targeted Attacks 59

Appendix A: A Perspective from the New European Cybercrime Center (EC3) 60

Appendix B: Full List of 2013 DBIR Contributors 62

For additional information on the DBIR and access to related content, please visit
www.verizonenterprise.com/DBIR/2013

The Verizon RISK team is dedicated to **R**esearching the ever-changing risk environment, **I**nvestigating and responding to all manner of security incidents, developing **S**olutions based on credible data and analysis, and cultivating **K**nowledge within Verizon, its clients, and the security community. We'd like to thank our 2013 DBIR partners, all those who graciously allowed us to bounce some half-baked ideas off them, as well as everyone else who contributed in ways large and small to this report (there are many of you). We sincerely appreciate it.

INTRODUCTION

2012. Perhaps more so than any other year, the large scale and diverse nature of data breaches and other network attacks took center stage. But rather than a synchronized chorus making its debut on New Year's Eve, we witnessed separate, ongoing movements that seemed to come together in full crescendo throughout the year. And from pubs to public agencies, mom-and-pops to multi-nationals, nobody was immune. As a result—perhaps agitated by ancient Mayan doomsday predictions—a growing segment of the security community adopted an "assume you're breached" mentality.

Motives for these attacks appear equally diverse. Money-minded miscreants continued to cash in on low-hanging fruit from any tree within reach. Bolder bandits took aim at better-defended targets in hopes of bigger hauls. Activist groups DoS'd and hacked under the very different—and sometimes blurred—banners of personal ideology and just-for-the-fun-of-it lulz. And, as a growing list of victims shared their stories, clandestine activity attributed to state-affiliated actors stirred international intrigue.

All in all, 2012 reminded us that breaches are a multi-faceted problem, and any one-dimensional attempt to describe them fails to adequately capture their complexity.

The 2013 Data Breach Investigations Report (DBIR) corroborates this and brings to bear the perspective of 19 global organizations on studying and combating data breaches in the modern world[1]. The list of partners is not only lengthy, but also quite diverse, crossing international and public/private lines. It's an interesting mix of law enforcement agencies, incident reporting/handling entities, a research institution, and other incident response (IR)/forensic service firms.

What's more, these organizations contributed a huge amount of data to the report. All told, we have the privilege of setting before you our analysis of more than 47,000 reported security incidents and 621 confirmed data breaches from the past year. Over the entire nine-year range of this study, that tally now exceeds 2,500 data breaches and 1.1 billion compromised records.

ALL IN ALL, 2012 REMINDED US THAT BREACHES ARE A MULTI-FACETED PROBLEM, AND ANY ONE-DIMENSIONAL ATTEMPT TO DESCRIBE THEM FAILS TO ADEQUATELY CAPTURE THEIR COMPLEXITY.

We continue to learn a great deal from this ongoing study, and we're glad to have the opportunity once again to share these findings with you. As usual, we begin with a few highlights.

1 See page 62 for a complete list.

Who are the victims?

37% of breaches affected financial organizations (+)

24% of breaches occurred in retail environments and restaurants (-)

20% of network intrusions involved manufacturing, transportation, and utilities (+)

20% of network intrusions hit information and professional services firms (+)

38% of breaches impacted larger organizations (+)

27 different countries are represented

Victims in this report span restaurants, retailers, media companies, banks, utilities, engineering firms, multi-national corporations, security providers, defense contractors, government agencies, and more across the globe. A definite relationship exists between industry and attack motive, which is most likely a byproduct of the data targeted (e.g., stealing payment cards from retailers and intellectual property [IP] from manufacturers).

The ratio among organizational sizes is fairly even this time around, rather than tipping toward the small end of the scale as it did in our last report.

Who's perpetrating breaches?

Another year, another report dominated by outsiders. Another crop of readers shaking their fists and exclaiming "No—insiders are 80% of all risk!" Perhaps they're right. But our findings consistently show—at least by sheer volume of breaches investigated by or reported to outside parties—that external actors rule.

Pro-insider majoritists may see some justification in the results for all security incidents (rather than just confirmed data breaches), as insiders take the lead in that dataset.

State-affiliated actors tied to China are the biggest mover in 2012. Their efforts to steal IP comprise about one-fifth of all breaches in this dataset.

92% perpetrated by outsiders

14% committed by insiders (+)

1% implicated business partners

7% involved multiple parties

19% attributed to state-affiliated actors (+)

Legend:
- A plus (+) sign indicates either a 10% or greater increase from the previous year's report
- A minus (-) sign indicates a 10% or greater decrease from the previous year's report
- Measurements without an indicator showed no significant change

How do breaches occur?

52% used some form of hacking (-)

76% of network intrusions exploited weak or stolen credentials (-)

40% incorporated malware (-)

35% involved physical attacks (+)

29% leveraged social tactics (+)

13% resulted from privilege misuse and abuse

The one-two combo of hacking and malware struck less often this round, but definitely isn't down for the count. Filtering out the large number of physical ATM skimming incidents shows exploitation of weak and stolen credentials still standing in the ring.

The proportion of breaches incorporating social tactics like phishing was four times higher in 2012. Credit the rise of this challenger to its widespread use in targeted espionage campaigns.

Correlated with the 14% of breaches tied to insiders, privilege misuse weighs in at 13%. Insider actions ranged from simple card skimming to far more complicated plots to smuggle corporate IP to competitors.

What commonalities exist?

It's notable that the majority (but no longer a super-majority) of breaches result from simpler opportunistic attacks than from money-hungry organized criminal groups.

But don't miss the "un-commonalities" evident in this report. More determined adversaries tied to state-level and industrial espionage make their presence felt too. We contrast different motives throughout the report to give a better sense for commonalities among them.

All of the above still takes forever and a day to discover, and that discovery is rarely made by the victim.

75% driven by financial motives (-)

71% targeted user devices (+)

54% compromised servers (-)

75% are considered opportunistic attacks (-)

78% of initial intrusions rated as low difficulty

69% discovered by external parties

66% took months or more to discover (+)

What can we do about it?

- Eliminate unnecessary data; keep tabs on what's left.
- Ensure essential controls are met; regularly check that they remain so.
- Collect, analyze and share incident data to create a rich data source that can drive security program effectiveness.
- Collect, analyze, and share tactical threat intelligence, especially Indicators of Compromise (IOCs), that can greatly aid defense and detection.
- Without deemphasizing prevention, focus on better and faster detection through a blend of people, processes, and technology.
- Regularly measure things like "number of compromised systems" and "mean time to detection" in networks. Use them to drive security practices.
- Evaluate the threat landscape to prioritize a treatment strategy. Don't buy into a "one-size fits all" approach to security.
- If you're a target of espionage, don't underestimate the tenacity of your adversary. Nor should you underestimate the intelligence and tools at your disposal.

Picking over the remains of breach victims might paint a grim picture of our current state, but it's not a hopeless one. As we have said before—we have the tools; it's selecting the right ones and using them in the right way that challenges us.

To that end, we're convinced of the critical importance of understanding your enemy. If handling payment cards is your business, then there's a narrowly defined set of controls on which you can focus. If your IP is a hot commodity, you've got your work cut out for you, but knowing the attack patterns (and sharing them) can make that work more fruitful. Take steps to better understand your threat landscape and deal with it accordingly. See the Conclusion for tips on threat-centric control prioritization.

We strongly recommend readers consider the detection of failures (in a reasonable time frame) as a success. The security industry has long been overly focused on prevention. Let's keep preventing, but enhance our ability to detect threats that slip through our defenses (which they will inevitably do).

Questions? Comments? Brilliant ideas?
We want to hear 'em. Drop us a line at dbir@verizon.com, find us on LinkedIn and Facebook, or post to Twitter with the hashtag #dbir.

A BRIEF PRIMER ON VERIS

VERIS is designed to provide a common language for describing security incidents in a structured and repeatable manner. It takes the narrative of "who did what to what (or whom) with what result" and translates it into the kind of data you see in this report. Because we hope to facilitate the tracking and sharing of security incidents, we released VERIS for free public use. Get additional information on the VERIS community site[2]; the full schema is available on GitHub[3]. Both are good companion references to this report for understanding terminology and context.

METHODOLOGY

Based on feedback, one of the things readers value most about this report is the level of rigor and integrity employed when collecting, analyzing, and presenting data. Knowing our readership cares about such things and consumes this information with a keen eye helps keep us honest. Detailing our methods is an important part of that honesty.

With 19 contributors to the 2013 DBIR, a single methodology simply does not exist. It's true that all incidents included in this report were eventually coded using the Vocabulary for Event Recording and Incident Sharing (VERIS) to create a common, anonymous aggregate dataset, but the path taken to that end differed among all contributors. In general, three basic methods (expounded below) were used to accomplish this. The incidents were: 1) recorded by Verizon using VERIS, 2) recorded by contributors using VERIS, or 3) re-coded using VERIS from a contributor's existing schema. All contributors received instruction to omit any information that might identify organizations or individuals involved, since such details are not necessary to create the DBIR.

Verizon's data collection methodology

The underlying methodology used by Verizon remains relatively unchanged from previous years. All results are based on first-hand evidence collected during paid external forensic investigations and related intelligence operations conducted by Verizon from 2004 through 2012. The 2012 caseload is the primary analytical focus of the report, but the entire range of data is referenced throughout. Though the RISK (Research, Investigations, Solutions, Knowledge) Team works a variety of engagements (more than 250 last year), only those leading to confirmed security incidents are included in this report.

Once an investigation is completed, RISK Team analysts use case evidence, reports, and interviews to code a VERIS record of the incident(s) by entering relevant information into a VERIS-based form. Input is then reviewed and validated by other members of the team to ensure reliable and consistent data.

2 http://www.veriscommunity.net
3 https://github.com/vz-risk/veris

Methodology for contributors using VERIS

Contributors using this method provided incident data to Verizon in VERIS format. For instance, agents of the U.S. Secret Service (USSS) used an internal VERIS-based application to record pertinent case details. The Irish Reporting and Information Security Service (IRISS-CERT) and several others recorded incidents directly into an application created and hosted by Verizon specifically for this purpose. For a few contributors, we captured the necessary data points via interviews and requested follow-up information as necessary. Whatever the exact process of recording data, these contributors used investigative notes, reports provided by the victim or other forensic firms, and their own experience gained in handling the incident.

Methodology for contributors not using VERIS

Some contributors already collect and store incident data using their own framework. A good example of this is the CERT Insider Threat Database compiled by the CERT Insider Threat Center at the Carnegie Mellon University Software Engineering Institute.[4] For this and other similar data sources, we created a translation between the original schema and VERIS[5] and re-coded incidents into valid VERIS records for import into the aggregate dataset. We worked with contributors to resolve any ambiguities or other challenges to data quality during this translation and validation process.

SHARING AND PUBLISHING INCIDENT INFORMATION ISN'T EASY, AND WE APPLAUD THE WILLINGNESS AND WORK OF ALL THESE CONTRIBUTORS TO MAKE THIS REPORT POSSIBLE. WE SINCERELY APPRECIATE IT.

Security incidents versus data disclosure

In the past, the DBIR has focused exclusively on security events resulting in confirmed data disclosure[6] rather than the broader spectrum of all security incidents[7]. The 2013 DBIR will continue that tradition for the primary analysis, but will extend that focus to all submitted incidents in designated places throughout the report. The reason for this change is simple. We received more than 47,000 incidents from our contributors in 2012, but only 621 of those were confirmed data disclosures with enough detail to produce a reasonably complete VERIS record sufficient for DBIR-level analysis. We could either toss tens of thousands of records in the bit bucket, or we could "VERISize" them to the best of our ability and include them in this report. Because we like data and assume you do too, we chose the latter and highlight (separately) the larger pool of incidents throughout the text.

COMPLETE LIST OF 2013 DBIR PARTNERS

- Australian Federal Police (AFP)
- CERT Insider Threat Center at the Carnegie Mellon University Software Engineering Institute (CERT) (U.S.)
- Consortium for Cybersecurity Action (U.S.)
- Danish Ministry of Defence, Center for Cybersecurity
- Danish National Police, NITES (National IT Investigation Section)
- Deloitte (U.S.)
- Dutch Police: National High Tech Crime Unit (NHTCU)
- Electricity Sector Information Sharing and Analysis Center (ES-ISAC) (U.S.)
- European Cyber Crime Center (EC3)
- G-C Partners, LLC (U.S.)
- Guardia Civil (Cybercrime Central Unit) (Spain)
- Industrial Control Systems Cyber Emergency Response Team (ICS-CERT)
- Irish Reporting and Information Security Service (IRISS-CERT)
- Malaysia Computer Emergency Response Team (MyCERT), CyberSecurity Malaysia
- National Cybersecurity and Communications Integration Center (NCCIC) (U.S.)
- ThreatSim (U.S.)
- U.S. Computer Emergency Readiness Team (US-CERT)
- U.S. Secret Service (USSS)

4 http://www.cert.org/blogs/insider_threat/2011/08/the_cert_insider_threat_database.html
5 For instance, CERT has an attribute named "Motives and expectations" that maps very well to the actor.internal.motive in VERIS.
6 VERIS defines data disclosure as any event resulting in confirmed compromise (unauthorized viewing or accessing) of any non-public information. Potential disclosures and other "data-at-risk" events do NOT meet this criterion, and have thus not traditionally been part of the sample set for this report.
7 VERIS defines a security incident (or security compromise) as any event affecting any security attribute (confidentiality/possession, integrity/authenticity, availability/utility) of any information asset.

A quick (but important) word on sample bias

Although we believe many of the findings presented in this report to be appropriate for generalization (and our confidence in this grows as we gather more data and compare it to that of others), bias undoubtedly exists. Presumably, the merged datasets in this report more closely reflect reality than they might in isolation; it is still a non-random sample. Unfortunately, we cannot measure exactly how much bias exists (i.e., in order to give a precise margin of error). We have no way of knowing what proportion of all data breaches are represented because we have no way of knowing the total number of data breaches across all organizations in 2012. Many breaches go unreported (though our sample does contain many of those). Many more are as yet unknown by the victim (and thereby unknown to us). What we do know is that our uncertainty shrinks and our knowledge grows along with what we are able to study—and that grew more than ever in 2012. At the end of the day, all we as researchers can do is pass our findings on to you to evaluate and use as you see fit.

A few final remarks

Sharing and publishing incident information isn't easy, and we applaud the willingness and work of all these contributors to make this report possible. We sincerely appreciate it. Better information creates a more complete and accurate understanding of the problems we all face. While we're on this topic, if your organization investigates, reports, or handles security incidents and might be interested in contributing to future DBIRs, let us know. The DBIR family continues to grow, and we welcome new members.

IF YOUR ORGANIZATION INVESTIGATES, REPORTS, OR HANDLES SECURITY INCIDENTS AND MIGHT BE INTERESTED IN CONTRIBUTING TO FUTURE DBIR'S, LET US KNOW. THE DBIR FAMILY CONTINUES TO GROW, AND WE WELCOME NEW MEMBERS.

ON THE IMPORTANCE OF RESEARCH QUESTIONS

As we collected and analyzed data for this year's report, we realized more than ever before the vast differences that exist in the way different organizations approach incident metrics. A few hours of philosophizing about this led us to conclude that this state of affairs has a lot to do with the very different set of research questions driving these organizations to track incident data. For instance, law enforcement agencies are strongly focused on the "who" and need a level of detail and validity that can stand up in court. Forensic providers concentrate on the "how" and the investigation and containment-oriented needs of the client. Computer security incident response teams (CSIRTs) and information sharing and analysis centers (ISACs) frequently report the "what" when facilitating information sharing among their members and often need to do so in a quicker, less detail-intensive manner. None of these approaches is better or worse than the others; they are different use cases driven by different research questions that inevitably result in datasets that differ in focus and detail.

An additional challenge is that we cannot identify all the research questions prior to gathering data. Part of our challenge, as an industry and intelligence community, is to balance the investment in data collection and management against our ability to answer future questions. While it'd be great to collect and record ALL THE THINGS, it's naïve to expect all partners and organizations have an equal amount of technical ability and resources to devote to data collection. There are both critical data points and a diminishing point of return, and that's a balance we try to strike in this report and our data collection efforts.

Questions? Comments? Brilliant ideas?
We want to hear 'em. Drop us a line at dbir@verizon.com, find us on LinkedIn and Facebook, or post to Twitter with the hashtag #dbir.

RESULTS AND ANALYSIS

The 2012 combined dataset represents the largest we have ever covered in any single year, spanning 47,000+ reported security incidents[8], 621 confirmed data disclosures, and at least 44 million compromised records (that we were able to quantify). Over the entire nine-year range of this study, that tally now exceeds 2,500 data disclosures and 1.1 billion compromised records.

As you read this report, you should assume all charts, tables, and commentary pertain to the 621 "breaches" leading to confirmed data disclosure unless otherwise stated. Any stats or references to the larger, separate dataset of all 47,000+ "security incidents" will be clearly identified as such. One of the problems with looking at a large amount of data for a diverse range of organizations is that averages across the whole are just so...*average*. Because the numbers speak *for* all organizations, they don't really speak *to* any particular organization or demographic in great detail. This is unavoidable. We've made the conscious decision to study all types of breaches as they affect all types of organizations, and we can't possibly include dedicated views for every slice of the dataset someone might want to see. What we can do, however, is to present the results in such a way that they are more readily applicable to different groups and interests.

We could split the dataset a myriad of ways, but we've chosen to continue last year's tack of highlighting differences (and similarities) between smaller and larger organizations (the latter defined as having at least 1,000 employees). In addition to this demographic segmentation, we also take the threat-centric approach of contrasting state-affiliated espionage[9], financially motivated crimes[10], and, to a lesser extent, activism[11] throughout the report. These genres typically involve very different actors using different actions against different assets, and make for an interesting comparative study into the wide variety of threats facing global organizations today. We hope this helps to make the findings in this report both generally informative and particularly useful.

Our goal is to communicate multiple perspectives for different groups, so we break down the elements of breaches by organizational size and threat community. To minimize disorientation without sacrificing data density, basic figures in this report utilize a consistent (some will say repetitive) format explained below.

AS YOU READ THIS REPORT, YOU SHOULD ASSUME ALL CHARTS, TABLES, AND COMMENTARY PERTAIN TO THE 621 "BREACHES" LEADING TO CONFIRMED DATA DISCLOSURE UNLESS OTHERWISE STATED. ANY STATS OR REFERENCES TO THE LARGER, SEPARATE DATASET OF ALL 47K "SECURITY INCIDENTS" WILL BE CLEARLY IDENTIFIED AS SUCH.

8 The Methodology section discusses the difference between security incidents and data disclosures.
9 This will carry the label "espionage" hereafter in this report, and refers to state-sponsored or affiliated actors seeking classified information, trade secrets, and intellectual property in order to gain national, strategic, or competitive advantage. The only exception is when it is used for internal actors, where it refers to industrial espionage perpetrated by the employees of the victim.
10 This will carry the label "financial" hereafter in this report, and refers to all criminal activities driven out of financial or personal gain.
11 This will carry the label "activism" hereafter in this report, and refers to illicit activities tied to the motives of ideology or protest and/or fun, curiosity, or pride.

Figure 1: Example of charts used in this report

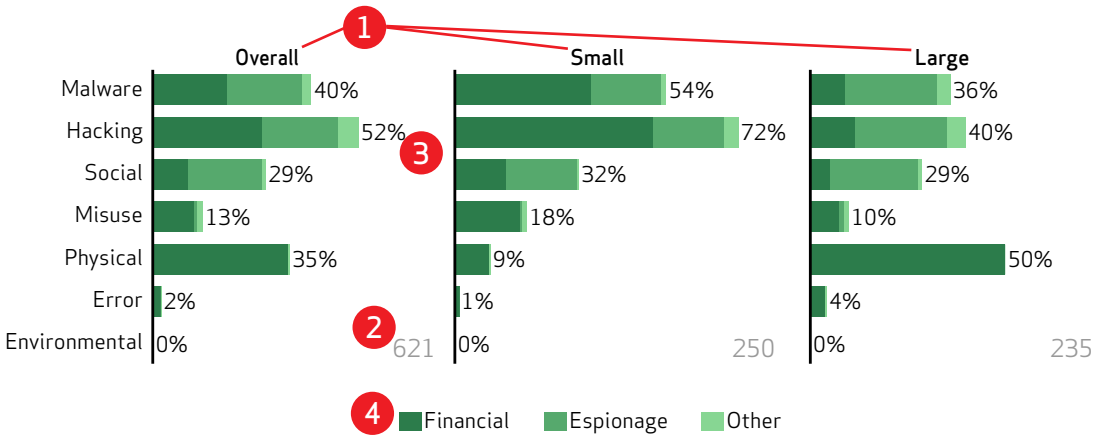

1. Demographic comparisons of "Overall" (all breaches of all organizations), "Small" (organizations with fewer than 1,000 employees), and "Large" (organizations with 1,000 employees or more). Note that not all organizational sizes were known to us for a variety of reasons, so there is a portion of unknown organizational size that is represented in the Overall result.
2. Number of breaches in each demographic. This is *n*; the denominator for percentages in each demographic.
3. Percent of breaches. Should be read as "52% of breaches affecting all organizations involved hacking." That figure changes to 72% of small organizations and 40% of large organizations.
4. Proportionality of motives. For example, nearly all breaches in the Physical category are financially motivated, while most Social actions are tied to espionage.

Many figures and tables in this report add up to more than 100%; this is not an error. It simply stems from the fact that items presented in a list are not always mutually exclusive, and thus, several can apply to any given incident. Figure 1 is a good example; many incidents involve malware and hacking and social actions in the sequence of events. Not all figures and tables contain all possible options but only those having a value greater than zero (and some truncate more than that). To see all options for any particular figure, refer to the VERIS framework[12]. Also, certain data points were only collected for Verizon IR cases, and these are identified in the text and figures. The "raw" stats for all figures in this report can be downloaded from the 2013 DBIR site[13].

As you study these findings, keep in mind that the dataset is anything but static. The number, nature, and sources of incidents change dramatically over time. Given this, you might be surprised at how stable many of the trends appear (a fact that we think strengthens their validity). On the other hand, certain trends are almost certainly more related to turmoil in the sample than significant changes in the external threat environment.

MANY FIGURES AND TABLES IN THIS REPORT ADD UP TO MORE THAN 100%; THIS IS NOT AN ERROR. IT SIMPLY STEMS FROM THE FACT THAT ITEMS PRESENTED IN A LIST ARE NOT ALWAYS MUTUALLY EXCLUSIVE, AND, THUS, SEVERAL CAN APPLY TO ANY GIVEN INCIDENT.

12 http://www.veriscommunity.net
13 http://www.verizonenterprise.com/DBIR/2013/

Demographics

Every year we begin by discussing the demographics of data breach victims. In the past we've treated demographic information as just another set of descriptive data points in a report that's chock full of them. The more we dig in and gain perspective, however, the more we become convinced it's not merely a victim-centered reference point, but serves as a pivotal indicator of underlying factors. In short, we're discovering that demographics may be one of the most critical and useful components of incident research.

ANY ATTEMPT TO ENFORCE A ONE-SIZE-FITS-ALL APPROACH TO SECURING OUR ASSETS MAY RESULT IN LEAVING SOME ORGANIZATIONS UNDER-PROTECTED FROM TARGETED ATTACKS WHILE OTHERS POTENTIALLY OVER-SPEND ON DEFENDING AGAINST SIMPLER OPPORTUNISTIC ATTACKS.

As we focus on demographics, we realize that the data tells a very different story than we hear in the industry—there is no one set of best practices that can be applied across industries and organizational sizes. Not all passwords are easily guessable, and we cannot make blanket statements about web applications being the most popular attack vector. Any attempt to enforce a one-size-fits-all approach to securing our assets may result in leaving some organizations under-protected from targeted attacks while others potentially over-spend on defending against simpler opportunistic attacks.

For example, small retailers and restaurants in the Americas should be focusing on the basics because attackers are leveraging poorly configured remote administration services to pull payment data from point of sale systems. But the basics won't be enough for the finance and insurance industry, which sees its ATMs targeted by skimming campaigns. And when we peel back that physical attack layer, we see a much higher proportion of attacks in its web applications than all other sectors. When we focus on manufacturing, engineering, consulting, and IT service firms, we see a whole different set of attacks exploiting human weaknesses through targeted social attacks to get multi-functional malware on internal systems. We discuss many of these differences in a series of industry-specific reports[14] we produced late in 2012.

Figure 2: Breach count by victim industry and employee count*

	Agriculture (11)	Mining (21)	Utilities (22)	Construction (23)	Manufacturing (31)	Wholesale Trade (42)	Retail (44)	Transportation (48)	Information (51)	Finance (52)	Real Estate (53)	Professional (54)	Management (55)	Administrative (56)	Educational (61)	Healthcare (62)	Recreation (71)	Accommodation (721)	Food Services (722)	Other Services (81)	Public (92)	Unknown	**Total**
1 to 100		1		2	10	1	79		5	18		14		3	1	3		3	38	6	2	7	193
101 to 1,000					13		3	1	8	3		5		1	2	1			13	2	4	1	57
1,001 to 10,000			1		7	1	3	22	10	12		6		1	2	1			2	1	2		71
10,001 to 100,000			2		13	1	4		2	93		5				1					1		122
More than 100,000		1	4		2					31		1					2		1				42
Unknown					1		7	1	14	73	1	5		1				1	2	2	5	23	136
Total		**2**	**7**	**2**	**46**	**3**	**96**	**24**	**39**	**230**	**1**	**36**		**6**	**5**	**6**	**2**	**4**	**56**	**11**	**14**	**31**	**621**

* Industries based on NAICS

14 http://www.verizonenterprise.com/products/security/dbir/verticals/

Figure 2 shows Finance leading the incident count this year, but that's mainly due to a large number of ATM skimming incidents. For an alternate view, Figure 3 gives a breakdown of industries resulting from network intrusions (incidents involving hacking or malware actions somewhere in the event chain). We see the Finance sector drops down the list since physical skimming attacks are filtered out. Introducing the additional dimension of attack motives reveals further evidence that limiting analysis only to high-level trends can be misleading. There is a very clear and important difference across industries with respect to motives that will factor prominently throughout this report.

We may not be able to gain much insight by looking at where victims base their operations, since most attacks can be launched from your mom's basement. But we may be able to infer some similarities among victims by looking at regions. For example, POS intrusions in Europe show up much less frequently in our data than in the Americas and Asia-Pacific regions (which may be related to payment card technology or sampling bias, or a combination of both). Additional regional attack trends are discussed more fully in the Threat Actions section. Overall, we recorded confirmed data disclosures from victims in 27 distinct countries, indicating we're not dealing with a simple localized problem.

THERE IS A VERY CLEAR AND IMPORTANT DIFFERENCE ACROSS INDUSTRIES WITH RESPECT TO MOTIVES THAT WILL FACTOR PROMINENTLY THROUGHOUT THIS REPORT.

Figure 3: Victim industry (filtered for network intrusions)*

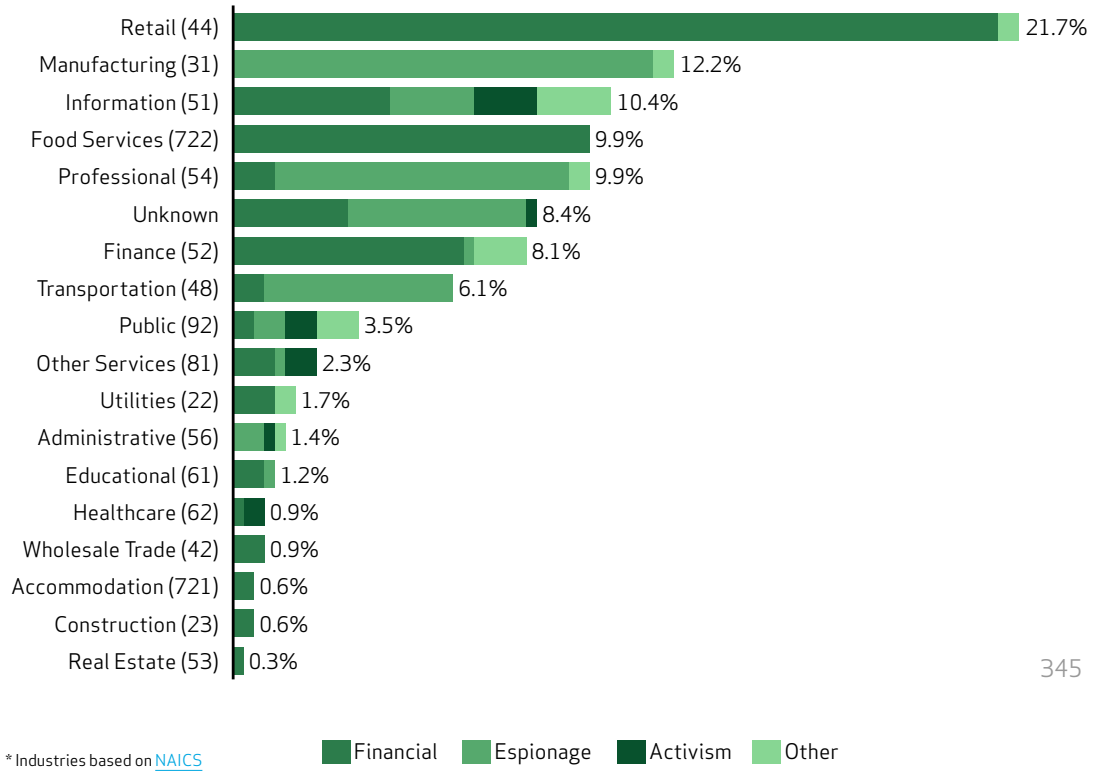

* Industries based on NAICS

Figure 4: Countries represented in combined caseload

Australia
Austria
Belgium
Brazil
Bulgaria
Canada
Denmark
Dominican Republic
France
Germany
Hong Kong SAR
India
Ireland

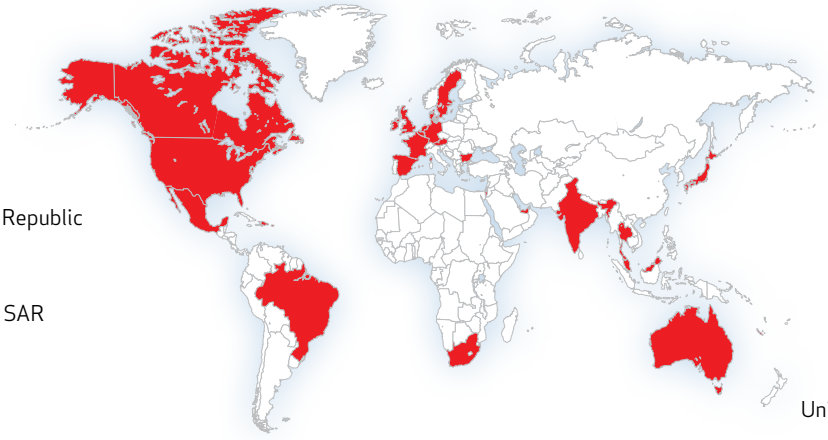

Israel
Japan
Luxembourg
Malaysia
Mexico
Netherlands
Puerto Rico (U.S. Terr.)
South Africa
Spain
Sweden
Thailand
United Arab Emirates
United Kingdom
United States of America

Am I a target of espionage?

Some may already know the answer to this question by firsthand experience. Many others assume they aren't or haven't thought much about it. Despite the growing number of disclosures and sometimes alarmist news coverage, many still see espionage as a problem relevant only to the Googles of the world. Unfortunately, this is simply not true, and we hope Figure 5 helps drive that point home.

Lesson one is that the "I'm too small to be a target" argument doesn't hold water. We see victims of

Figure 5: Victim industry by employee count (filtered for espionage)*

	Agriculture (11)	Mining (21)	Utilities (22)	Construction (23)	Manufacturing (31)	Wholesale Trade (42)	Retail (44)	Transportation (48)	Information (51)	Finance (52)	Real Estate (53)	Professional (54)	Management (55)	Administrative (56)	Educational (61)	Healthcare (62)	Recreation (71)	Accommodation (721)	Food Services (722)	Other Services (81)	Public (92)	Unknown	**Total**
1 to 100					9				1			11		1									22
101 to 1,000					13				3			4									2	1	23
1,001 to 10,000					5			18	4			5		1	1					1	1		36
10,001 to 100,000					12							5											17
More than 100,000					1					1													2
Unknown												3		1								16	20
Total					40			18	8	1		28		3	1					1	3	17	120

* Industries based on NAICS

espionage campaigns ranging from large multi-nationals all the way down to those that have no IT staff at all. Lesson two is that some industries appear to be more targeted than others. This often matches up with strategic objectives of the actors involved. Figure 5 draws from a few large campaigns, so industries falling outside those objectives are likely under-represented. But still, the blank columns associated with Retail and Food Services (which are the most targeted industries for financially motivated actors) are perhaps informative about the underlying motivations and goals.

Most organizations have some form of proprietary or internal information they want kept private. Without this confidential information it's hard to stay competitive. And because it's secret and competitively advantageous, others may want to steal it. Thus, "who wants my proprietary info?" is probably a better question than "am I a target of espionage?" Accurately assessing the varieties of actors that might and their capability to obtain it is crucial.

Another thing to keep in mind is that it might not be your data they're after at all. If your organization does business with others that fall within the espionage crosshairs, you might make a great pivot point into their environment. Make sure to take that into account when developing a well-considered and informed answer to this important question.

A⁴ Threat Overview

The Incident Description section of VERIS translates the incident narrative of "who did what to what (or whom) with what result?" into a form more suitable for trending and analysis. To accomplish this, VERIS employs the A⁴ Threat Model developed by Verizon's RISK Team. Describing an incident essentially means identifying all the actors, actions, assets, and attributes involved (the four A's).

Figure 6: VERIS A⁴ grid depicting associations between actors, actions, assets, and attributes

	External.Malware	External.Hacking	External.Social	External.Misuse	External.Physical	External.Error	External.Env	Internal.Malware	Internal.Hacking	Internal.Social	Internal.Misuse	Internal.Physical	Internal.Error	Internal.Env	Partner.Malware	Partner.Hacking	Partner.Social	Partner.Misuse	Partner.Physical	Partner.Error	Partner.Env
Server.Conf	35%	48%	23%	2%	.	1%		.	2%	2%	5%	1%	2%		.	.	.	1%		.	
Server.Integ	35%	41%	23%	2%	.	1%		.	2%	2%	3%	1%	2%		.	.	.	.			
Server.Avail	1%	2%	1%			.		.	.		.	.	.								
Network.Conf	.	.	.	.	1%	.		.	.		.		.								
Network.Integ	.	.	.	.	1%	.		.	.		.		.								
Network.Avail		.	.			.			.				.								
User.Conf	35%	36%	22%	1%	32%	.		.	.	.	3%	1%	.					.			
User.Integ	35%	34%	22%	1%	32%	.		.	.	.	1%	1%	.					.			
User.Avail	.	.	.	.	1%			.	.		1%	.									
Media.Conf	.	.	2%	2%	1%					2%	5%	2%	.					.			
Media.Integ	.	.	2%	2%	1%					2%	3%	1%	.								
Media.Avail			.	.	1%					.	.	1%									
People.Conf	22%	24%	29%	4%	1%			.		4%	4%	1%				.	.	.			
People.Integ	22%	24%	29%	4%	1%			.		4%	4%	1%				.	.	.			
People.Avail	.	2%	2%	1%	1%			.		.	1%	1%									

- **Actors:** Whose actions affected the asset
- **Actions:** What actions affected the asset
- **Assets:** Which assets were affected
- **Attributes:** How the asset was affected

It is our position that the four A's represent the minimum information necessary to adequately describe any incident or threat scenario. Furthermore, this structure provides an optimal framework within which to measure frequency, associate controls, link impact, and many other concepts required for risk management.

The next few sections provide separate analyses of the actors, actions, assets, and attributes, but first we'd like to present a big-picture view that ties them all together. Figure 6 shows *associations* between the four A's, and is our most consolidated view of the 621 breaches analyzed

IT IS OUR POSITION THAT THE FOUR A'S REPRESENT THE MINIMUM INFORMATION NECESSARY TO ADEQUATELY DESCRIBE ANY INCIDENT OR THREAT SCENARIO. FURTHERMORE, THIS STRUCTURE PROVIDES AN OPTIMAL FRAMEWORK WITHIN WHICH TO MEASURE FREQUENCY, ASSOCIATE CONTROLS, LINK IMPACT, AND MANY OTHER CONCEPTS REQUIRED FOR RISK MANAGEMENT.

in 2012. The proper way to interpret Figure 6 is "35% of all incidents involved an external actor AND a malware action AND a server asset AND the confidentiality attribute" (upper left intersection). It does NOT necessarily mean that an external actor installed malware that compromised the confidentiality of a server[15].

Figure 7: VERIS A⁴ grid depicting associations between actors, actions, assets, and attributes across 47,000+ security incidents

	External.Malware	External.Hacking	External.Social	External.Misuse	External.Physical	External.Error	External.Env	Internal.Malware	Internal.Hacking	Internal.Social	Internal.Misuse	Internal.Physical	Internal.Error	Internal.Env	Partner.Malware	Partner.Hacking	Partner.Social	Partner.Misuse	Partner.Physical	Partner.Error	Partner.Env
Server.Conf	.	2%	.	.	.	.		.	.	.	.	.	.		.	.	.	.		.	
Server.Integ	.	10%	.	.	.	.		.	.	.	.	.	.		.	.	.	.			
Server.Avail	.	.	.	.		.		.	.		.	.	.								
Network.Conf	.	.	.	.	.	.		.	.		.		.								
Network.Integ	.	.	.	.	.	.		.	.		.		.								
Network.Avail		.	.		.	.			.		.		.								
User.Conf	1%	.	.	.	.	.		.	.	.	.	.	20%			.		.		.	
User.Integ	20%	.	1%	.	.	.		.	.	.	20%	.	.		.			.		.	
User.Avail	.	.	.	.	.			.	.		.	.	18%			.				.	
Media.Conf	.	.	.	.	.					.	.	.	29%					.			
Media.Integ	.	.	.	.	.			.		.	.	.	.								
Media.Avail			.	.	.					.	.	.									
People.Conf	.	.	1%	.	.			.		.	.	.				.	.	.			
People.Integ	1%	.	1%	.	.			.		.	.	.				.	.	.			
People.Avail	.	.	.	.	.			.		.	.	.									

15 More discussion on the A⁴ Grid and its production can be found at http://www.veriscommunity.net/doku.php?id=a4grid.

The first observation is that much of the grid is blank or almost blank (<1% denoted by "."), meaning the intersecting A's never (or rarely) appeared together within a single incident. On the other end of the spectrum, a relatively few hot spots jump out. As we dive deeper into the report, it will become obvious why this is so, but for now let's take it at face value that the intensity is confined to a relatively few associations.

In our last report, we hinted that this grid can look a lot different depending on the sample from which it is derived. For instance, compare the grid for the 621 confirmed data compromise events to the one for all 47,000+ security incidents shared with us for this report (Figure 7). Striking, isn't it? Upon seeing that emerge from the data, we initially had something of an "Ermahgerd" reaction. Adding tens of thousands of incidents reveals more coverage, but results in fewer hotspots. Why? In short, it's because such a large proportion of all reported security incidents are rather mundane and repetitive. They represent the kind of things organizations deal with on a regular basis that don't involve data theft and aren't typically investigated by law enforcement agencies or external forensic firms.

This section is quite fun because we get to experiment with analysis and visualization methods a bit. In Figure 8, we simplify the A^4 grid to an A^2 format (actions and assets), and remove the percentages to enhance the contrasting patterns among breaches of different motives. You can almost see the underlying equation of Overall = Espionage + Financial + Activism + Other. We geeked out over these internally for some time and want you to have the same opportunity. Swipe, scroll, or flip whenever you're ready.

Figure 8: VERIS A^2 grid depicting associations between actions and assets (split by actor motive)

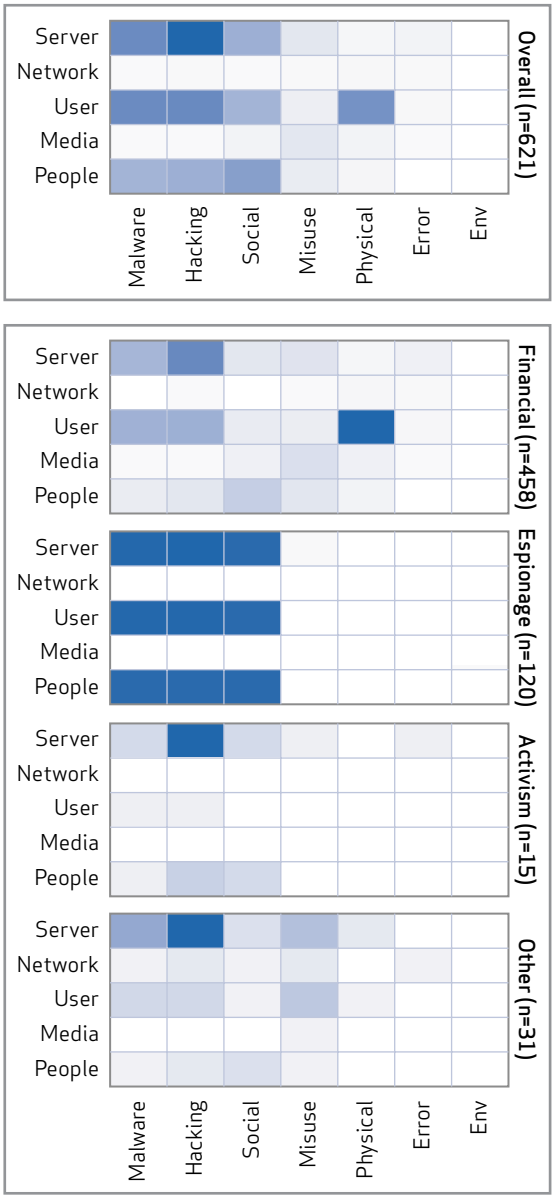

Threat Actors

Entities that cause or contribute to an incident are known as threat actors, and more than one can be involved in any particular incident. Actions performed by them can be malicious or non-malicious, intentional or unintentional, causal or contributory, and stem from a variety of motives (all of which will be discussed in subsequent actor-specific sections). Identifying actors is critical to immediate corrective actions and longer-term defensive strategies. VERIS specifies three primary categories of threat actors—external, internal, and partner.

- **External:** External actors originate outside the victim organization and its network of partners. Typically, no trust or privilege is implied for external entities.
- **Internal:** Internal actors come from within the victim organization. Insiders are trusted and privileged (some more than others).
- **Partners:** Partners include any third party sharing a business relationship with the victim organization. Some level of trust and privilege is usually implied between business partners.

Figure 9: Threat actor categories over time

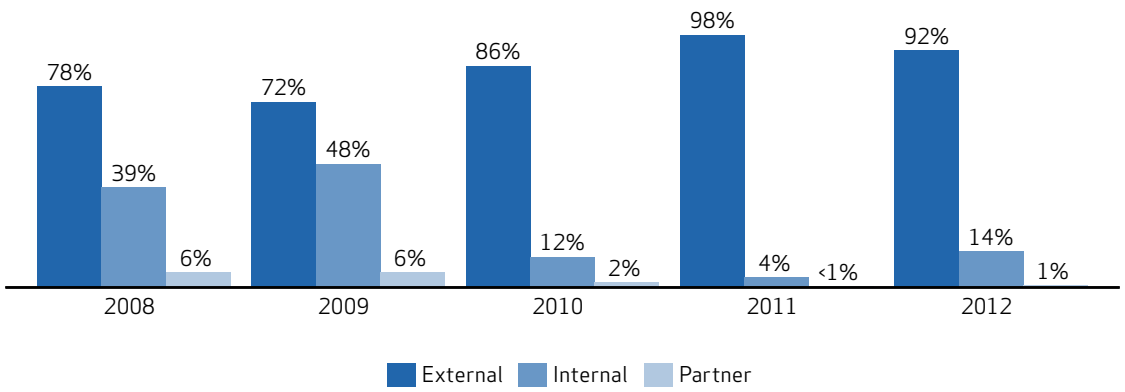

Figure 9, which shows the distribution of threat actors over time, should be familiar to our readers, as should the message it conveys. Keep in mind, however, that this is a volatile sample, with a different set of contributors each year.

The vast majority of 2012 breaches involve outsiders, though their exclusivity appears somewhat curbed when compared to 2011[16]. The two big reasons for the dominance of external actors are their numerical advantage and greater attack scalability. An organization will always have more outsiders than insiders, and the Internet connects criminals to a virtually limitless host of potential victims.

Figure 10: Threat actor categories

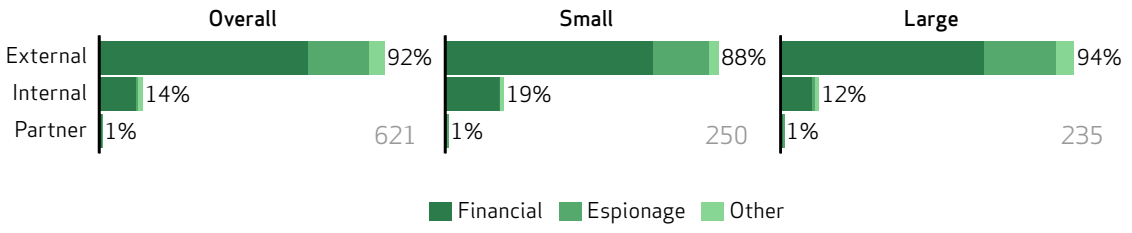

16 For more discussion on this trend, visit the Take a look back section of the main DBIR site to view the 2011 and 2012 reports.

Internal actors mustered a stronger showing, but this is more reflective of a changing sample than an evolving threat environment. The 2011 and 2012 DBIRs featured datasets teeming with highly scalable remote attacks that essentially overwhelmed the external-internal ratio. Fewer of these attacks, along with insider-heavy datasets from some of our partners (particularly CERT and G-C Partners) helped the proportion of insider breaches rebound to pre-surge levels. Breaches committed by business partners remain very low.

Moving on from time series analysis, we'll dig deeper into 2012 breaches to see what additional actor-related nuggets can be unearthed. Figure 10 reveals little difference between large and small organizations. A greater number of employees doesn't necessarily lead to more insider breaches, at least as far as we can tell from our last few years of data.

Examining motives across actor categories yields a gem or two worth tucking away. The majority of external actors exhibits financial motivations, but also reflects a strong pull toward espionage as well. Insiders showed mostly financial motives in both large and small organizations.

To mine much more than that out of this data, we'll need the pickaxes to chip away at each actor category in isolation.

Figure 11: Threat actor categories across 47,000+ security incidents

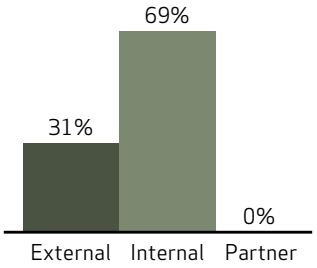

The list of modern myths just shrunk by one; insider majorities do exist! Bigfoot and Nessie remain unconfirmed. In all seriousness, though, this is an important reminder that confirmed data breaches are a rather exclusive subset amid the plethora of all types of security incidents. As will become more apparent in the Threat Actions section, most of these are insiders acting carelessly rather than maliciously.

External Actors (92% of breaches)

The reader may notice a bit of a "third verse, same as the first" motif in this section, but there is, at least, a fresh beat to the song this year, revolving around the "who" and "why" behind external breaches. VERIS differentiates between threat actor *variety* and *motive*, but these often correlate strongly. In other words, who you are and why you're doing something are not easily separated, and Figure 12 illustrates this point well.

In terms of the "who," more than half of all external breaches tie to organized criminal groups. This reflects the high prevalence of illicit activities associated with threat actors of this ilk, such as spamming, scamming, payment fraud, account takeovers, identity theft, etc. For professional criminals, the "why" is simple and consistent—money. As economic and social activities continue to go online, criminals will follow in order to exploit the soaring amount of data that can be (all too easily) converted to cash.

State-affiliated groups rise to the number two spot for the 2012 dataset, and there are several plausible explanations for this. On one hand, we saw a dip in financially motivated cases against small organizations in our dataset, and that dip allows other trends to become more pronounced. Another factor is the larger set of data sharing partners in this report that widens the population of incidents we can analyze. Furthermore, our own investigations comprised more espionage cases than any previous year, and this was bolstered by increased efforts to collect, share, and correlate IOCs that greatly improve the ability to uncover targeted attacks. So, it may be true that espionage activity is up,

Figure 12: Variety of external actor

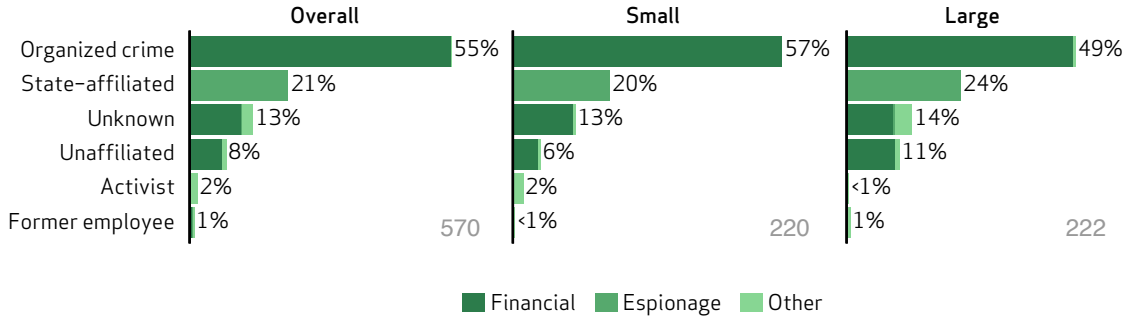

but it's also true that better sharing and improved detection capabilities result in more detections.

Threat actors engaged in espionage campaigns leave a completely different footprint than those motivated by direct financial gain. They seek data that furthers national interests, such as military or classified information, economy-boosting plans, insider information or trade secrets, and technical resources such as source code. They will generally not target payment systems and information, and according to our data, they aren't even targeting certain industries that have topped the charts for financially motivated attackers (e.g., Retail and Food Services).

It's important to point out that the process of attributing an attack to a particular person, group, or country is non-trivial. While we don't require evidence that will stand up in a court of law, we also don't guess or simply rely on low-confidence indicators like geolocation of IP addresses. Sometimes attribution is based on arrests and prosecutions, but it often comes down to the use of particular tactics, techniques, and procedures (TTPs) associated with known threat groups. Naturally, available information isn't always clear-cut, and thus "Unknown" features prominently in Figure 12. While having unknowns in the data is not all that informative, we've made the call that it is preferable to misattribution.

The proportion of incidents involving activist groups is on par with our previous report, but the amount of data they stole is down substantially (they nabbed more than any other variety of actor in 2011). Hacktivists generally act out of ideological motivations, but sometimes just for fun and epic lulz. These motives are often difficult to separate, and in any case, they just were not prevalent in frequency or record count in the data this year. Much of the activity claimed by these actors in 2012 shifted to other forms such as denial of service (DoS) attacks.

Country of origin

For the majority (>75%) of breaches in our dataset, the threat actor's country of origin was discoverable, and these were distributed across 40 different nations. From Figure 13, it's fascinatingly apparent that motive correlates very highly with country of origin. The majority of financially motivated incidents involved actors in either the U.S. or Eastern European countries (e.g., Romania, Bulgaria, and the Russian Federation). 96% of espionage cases were attributed to threat actors in China and the remaining 4% were unknown. This may mean that other threat groups perform their activities with greater stealth and subterfuge. But it could also mean that China is, in fact, the most active source of national and industrial espionage in the world today.

Figure 13: Origin of external actors: Top 10

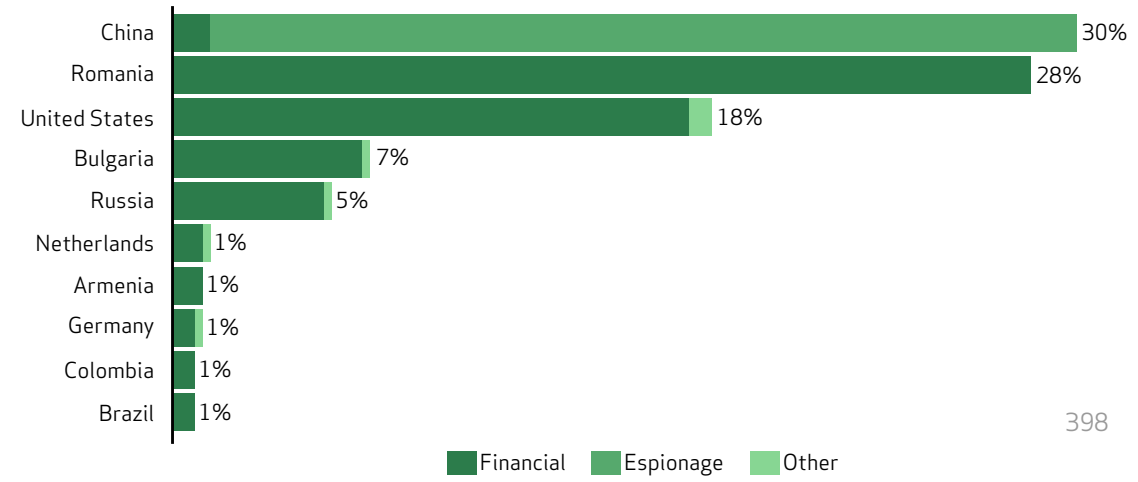

Table 1: Profiling threat actors

	ORGANIZED CRIME	STATE-AFFILIATED	ACTIVISTS
VICTIM INDUSTRY	Finance Retail Food	Manufacturing Professional Transportation	Information Public Other Services
REGION OF OPERATION	Eastern Europe North America	East Asia (China)	Western Europe North America
COMMON ACTIONS	Tampering (Physical) Brute force (Hacking) Spyware (Malware) Capture stored data (Malware) Adminware (Malware) RAM Scraper (Malware)	Backdoor (Malware) Phishing (Social) Command/Control (C2) (Malware, Hacking) Export data (Malware) Password dumper (Malware) Downloader (Malware) Stolen creds (Hacking)	SQLi (Hacking) Stolen creds (Hacking) Brute force (Hacking) RFI (Hacking) Backdoor (Malware)
TARGETED ASSETS	ATM POS controller POS terminal Database Desktop	Laptop/desktop File server Mail server Directory server	Web application Database Mail server
DESIRED DATA	Payment cards Credentials Bank account info	Credentials Internal organization data Trade secrets System info	Personal info Credentials Internal organization data

Table 1 pretty much speaks for itself, so we won't belabor the point except to say that it's based directly on our dataset rather than anecdotes. Items appear in order of prevalence among breaches attributed to each threat actor variety. Happy profiling!

Internal Actors (14% of breaches)

Consistent with prior years, most insider breaches were deliberate and malicious in nature, and the majority arose from financial motives. Of course, not all insiders are about malice and money. Inappropriate behaviors such as "bringing work home" via personal e-mail accounts or sneakernetting data out on a USB drive against policy also expose sensitive data to a loss of organizational control. While not common in our main dataset, unintentional actions can have the same effect. The broader collection of 47,000+ security incidents featured in this report offers ample evidence of this fact. These include "low-tech" events, such as sending sensitive documents to the wrong recipients, as well as less-frequent mistakes by system administrators and programmers. For instance, one incident in our caseload involved an errantly configured application debug setting that caused sensitive financial data to be stored insecurely and exposed to unauthorized parties.

Data theft involving programmers, administrators, or executives certainly makes for interesting anecdotes, but is still less common in our overall dataset than incidents driven by employees with little to no technical aptitude or organizational power. Per Figure 14, employees directly involved in the payment chain—like cashiers, waiters, and bank tellers—are most often responsible for breaches in our dataset. Such employees are often solicited to skim payment cards or provide customer account data to external parties that use the stolen information to fuel fraud schemes.

DATA THEFT INVOLVING PROGRAMMERS, ADMINISTRATORS, OR EXECUTIVES CERTAINLY MAKES FOR INTERESTING ANECDOTES, BUT IS STILL LESS COMMON IN OUR OVERALL DATASET THAN INCIDENTS DRIVEN BY EMPLOYEES WITH LITTLE TO NO TECHNICAL APTITUDE OR ORGANIZATIONAL POWER.

In larger organizations, the cashiers drop off significantly and administrators top the list. But before someone uses this detail in an eye-popping infographic about the scary administrator, we feel obliged to point out their role was accidental in eight out of the 13 incidents. Which means that infographic should be about how scary human error is (you're welcome, random vendor!). Perhaps the slightly less breach-prone regular users should seize the opportunity afforded here to start grumbling about the "stupid admins" for a change.

Figure 14: Variety of internal actors

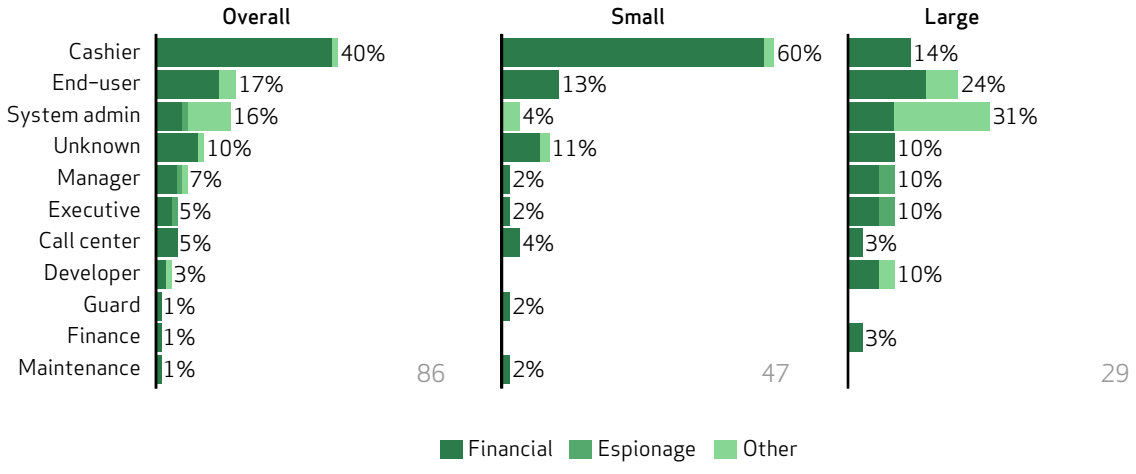

One last thing to note: the "espionage" here is specifically from insiders and we analyzed three cases from 2012. All three involved soon-to-be ex-employees on their way out the door trying to take proprietary information to a new employer. All three of those cases were at the manager or executive level as well.

Partner Actors (1% of breaches)

Partner breaches more than doubled from our last report! The prior sentence is factually accurate, but the increase was from only three incidents last year to an underwhelming seven this year. Partner involvement came in several forms, including a courier that lost a device with sensitive data, and point of sale vendors whose employees accessed customer systems in search of payment card information. There is even an incident involving a security consultant who used knowledge gained from a sanctioned penetration test to conduct a very unsanctioned breach of the victim's network.

Readers often misinterpret this low percentage of partner actors to mean that partner-hosted or managed assets are rarely breached. This is not so. The results above relate specifically to partners as threat actors, meaning they played a causal role in the breach. Incidents involving partner-hosted or managed assets are not reflected in this section (unless the partner's action(s) led to the information disclosure).

CERT INSIDER THREAT RESEARCH: CHARACTERISTICS OF A MALICIOUS INSIDER

"If you see something, say something" is the slogan for the Department of Homeland Security's campaign designed to raise public awareness about the signs of potential terrorist activity. The idea that everyone should be on the lookout for malicious behavior can easily be carried over into information security awareness as well. The CERT Insider Threat Center at the Carnegie Mellon University Software Engineering Institute has produced a body of research on the malicious insiders and their behavioral characteristics. According to their research[17], insiders intent on or considering malicious actions often exhibit identifiable characteristics and/or warning signs before engaging in those acts. They include:

- More than 30% of insiders engaging in IT sabotage had a prior arrest history. Note, however, this statistic may not be meaningful. For instance, a 2011 study found approximately 30% of U.S. adults have been arrested by age 23.
- Exhibiting concerning behaviors at work like bragging about the damage they could do to the organization if they so desired. This is often traced to a catalyst event like being passed over for promotion.
- Utilizing the organization's resources for a side business or having serious conversations with coworkers about starting a competing business.
- Attempting to gain employees' passwords or to obtain access through trickery or exploitation of a trusted relationship (often called "social engineering").
- In more than 70% of IP theft cases, insiders steal the information within 30 days of announcing their resignation. Changes in the pattern or quantity of information retrievals in that timeframe are potential indicators.
- More than half of insiders committing IT sabotage were former employees who regained access via backdoors or corporate accounts that were never disabled.

17 Silowash, G., Cappelli, D., Moore, A., Trzeciak, R., Shimeall, T., & Flynn, L. (2012). Common Sense Guide to Mitigating Insider Threats. In S. E. Institute (Ed.), (4th ed., pp. 17): Carnegie Mellon University. http://www.sei.cmu.edu/library/abstracts/reports/12tr012.cfm

Threat Actions

Threat actions describe what the actor did to cause or to contribute to the breach. Every incident contains one or more actions, often causing percentages to add up to more than 100%. VERIS classifies actions into seven categories; Figure 15 records the prevalence of each within our historical data, and Figure 16 isolates 2012 to show more detail.

Figure 15: Threat action categories over time

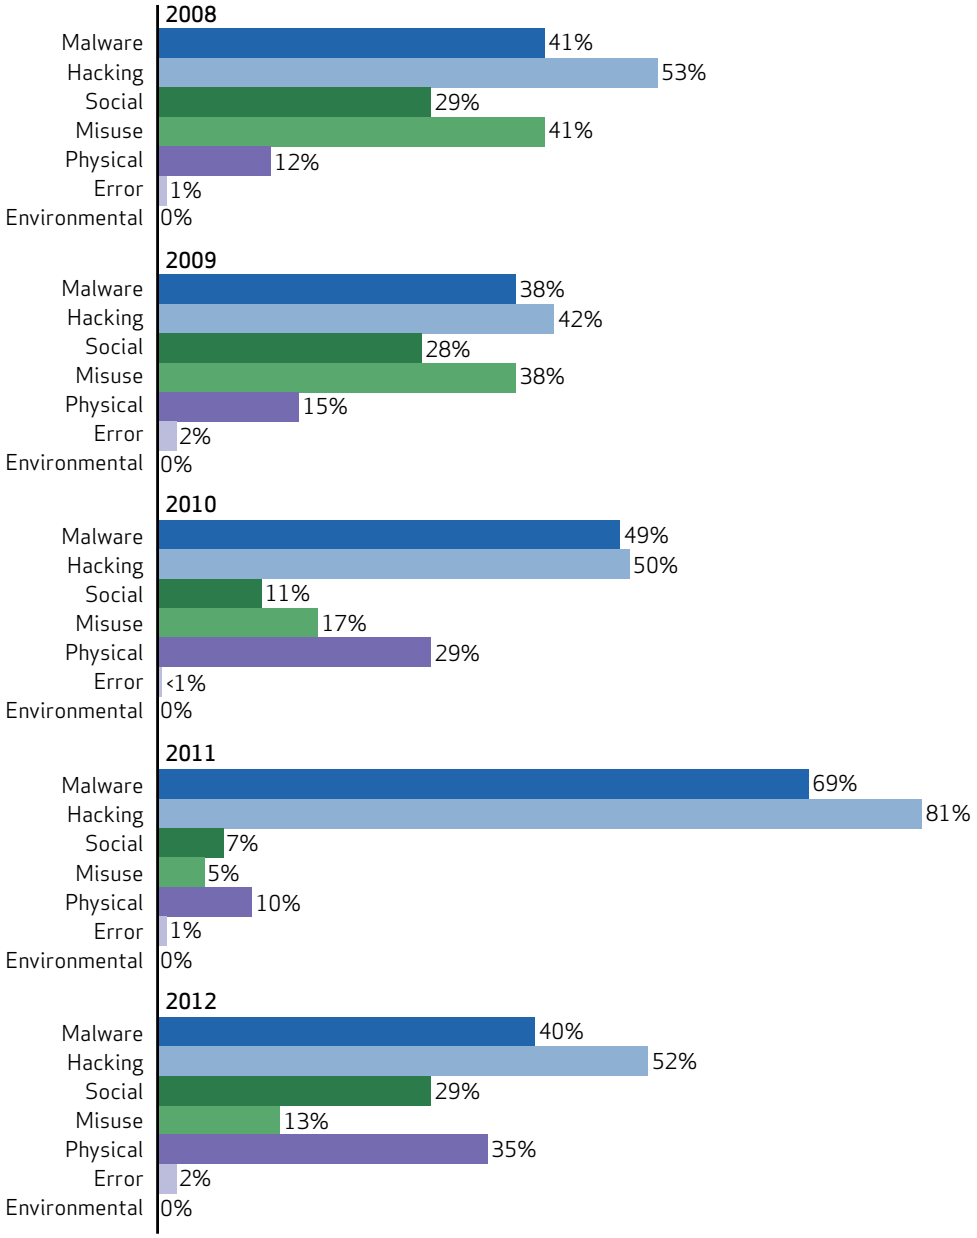

Before examining Figure 15, though, we should take a moment to admire the logos on the cover. The number of DBIR contributors has doubled each year since 2009 and tripled for this current edition. We want to stress the importance of this point because changes in overall trends (like threat actions observed year over year) may be caused by changes within the sample set or in the threat environment (or both). With that out of the way, we can get on to the data.

Malware and hacking still rank as the most common actions, but they scaled back rather significantly among 2012 breaches. Social quadrupled its proportion, physical reached its highest point ever, and thanks to some insider-focused contributors, misuse more than doubled. Overall, the threat profile of this dataset appears more balanced than its lopsided 2011 counterpart. And that probably has a lot to do with that growing list of partners.

Balancing act aside, the amount of repetitious attack patterns in these breaches (in any year) shouldn't be underestimated. There's the "POS smash-and-grab" we've described many times before that levies a brute force and malware combination against smaller franchises. Meanwhile, the "let's get physical" squad installs skimming devices on automated teller machines (ATMs) at larger banks. After adding in the "Assured Penetration Technique" of phishing-malware-hacking-entrenchment that is the staple of espionage campaigns, there's not a lot of room left for experimentation. While some may argue that we are dealing with an intelligent and adaptive adversary, the data tells us that adaptation isn't necessary for many of these attackers.

We know what you're thinking—that's a little vague. You'd like to see some substantiation. We can rectify that. Hang on, we need some scratch paper:

111	POS smash-and-grab
190	Physical ATM
+ 120	Assured Penetration Technique
421	
÷ 621	Total Breaches
68%	

Those three scenarios describe about two out of every three breaches in our 2012 dataset. While there is still some wiggle room for the baddies to be creative, this is an indication that treating our adversaries as random and unpredictable is counterproductive. We may be able to reduce the majority of attacks by focusing on a handful of attack patterns.

Figure 16: Threat action categories

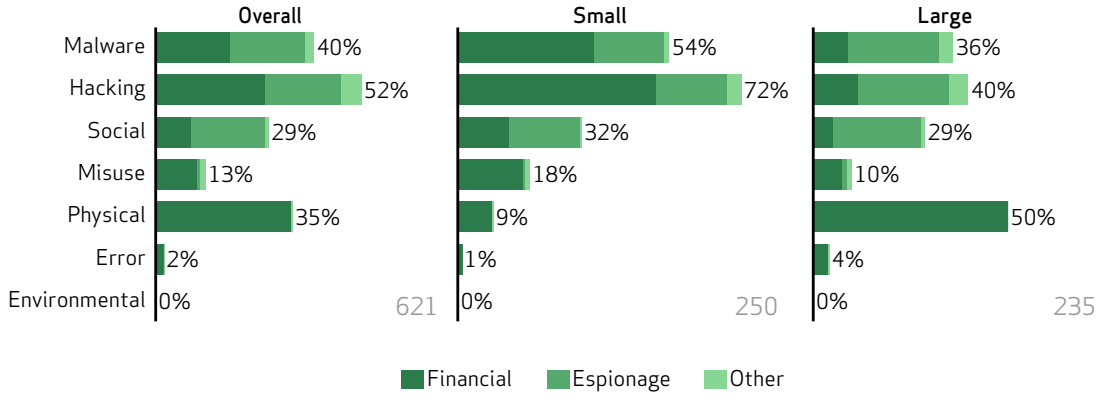

Figure 17: Top 20 threat actions

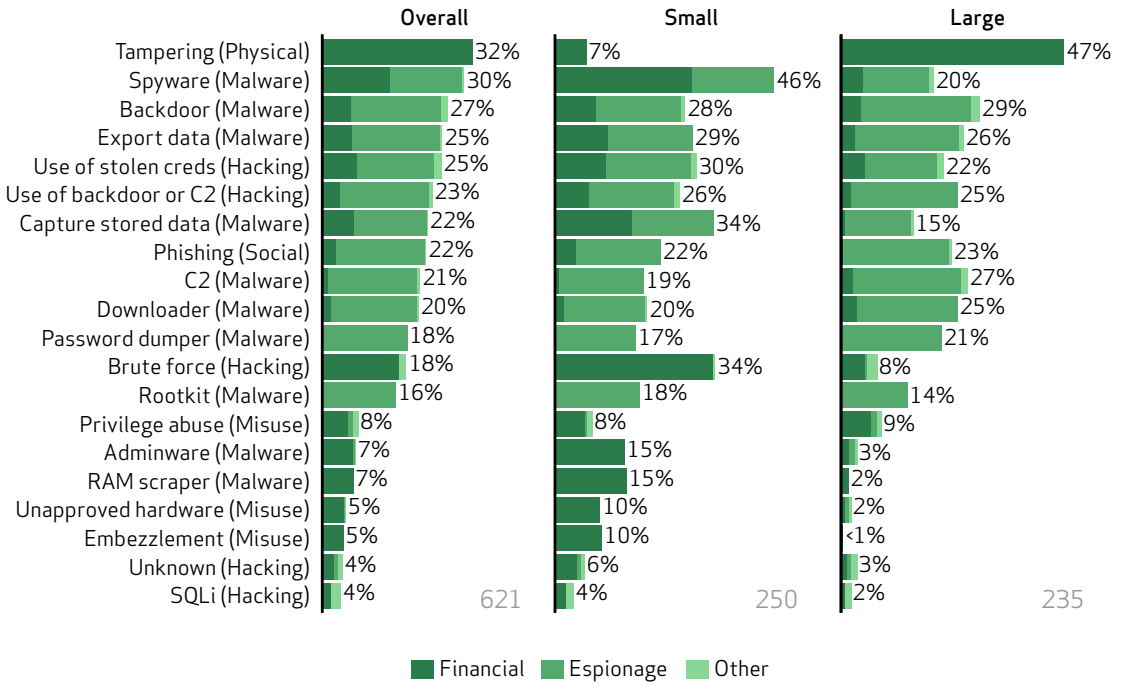

Within these broad attack patterns, however, particular techniques do vary. Figure 17, which ranks some of these, is so jam-packed with information it's probably best to just study it for a while to see what interests you, rather than reading our commentary around it. The action varieties are specific enough that the differences between actor motives and victim sizes really begin to emerge. Once you've gotten your 1,000 words worth, read on, and we'll begin unpacking each action category.

Figure 18: Threat action categories across 47,000+ security incidents

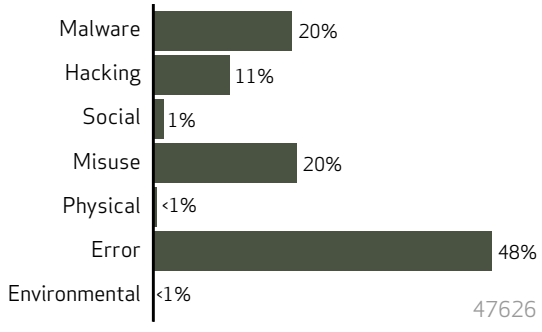

Commensurate with the high proportion of internal actors, error and misuse feature prominently among threat actions within the larger dataset of 47,000+ security incidents. Error consists mainly of lost devices, publishing goof-ups, and mis-delivered e-mails, faxes, and documents that potentially expose sensitive information. Misuse was a mix of malicious privilege abuse and use policy violations. Due to less-detailed reporting, malware specifics usually weren't available, but infections of spyware, botnets, and backdoors were observed most frequently. Use of stolen credentials, backdoor exploits, and SQL injection topped the list in the hacking category.

Regional Threat Trends

We very often receive questions concerning regional trends among the breaches we examine. This topic is challenging to study because so many factors are at play. For instance, breach reporting requirements differ dramatically country to country, and strongly affect data available to the public, law enforcement agencies, and the private sector. Furthermore, the governance, staffing, mission, and reporting methodologies of global CSIRT bodies make their datasets hard to compare on equal footing as well. If incidents reported to IRISS-CERT differ from incidents reported by MyCERT, does this reflect variations in the threat environment of Ireland and Malaysia? Maybe to a certain extent, but numerous confounding factors are likely present as well.

So, while we may not have a high-certainty answer to the question of regional attack trends, we do have some data to share. Figure 19 draws from the full dataset of 47,000+ security incidents shared with us by this year's DBIR partners. To achieve a better baseline for comparison, we filtered out errors and physical threats and also removed any actions for which a specific variety was unknown (e.g., if "unauthorized access" was reported, but no information was given about how that was achieved—which happens a lot—we ignored it). As you can see from the n values in Figure 19, these stipulations whittle the dataset down substantially. The regions depicted correspond to the country of the victim reporting the incident.

Figure 19 is…interesting to say the least. The only real action we recommend based on these results is to consider how we, as a global community, might improve them to produce more complete and reliable data for future research.

Figure 19: Top 20 threat actions by victim region across 47,000+ security incidents (filtered for network intrusions)

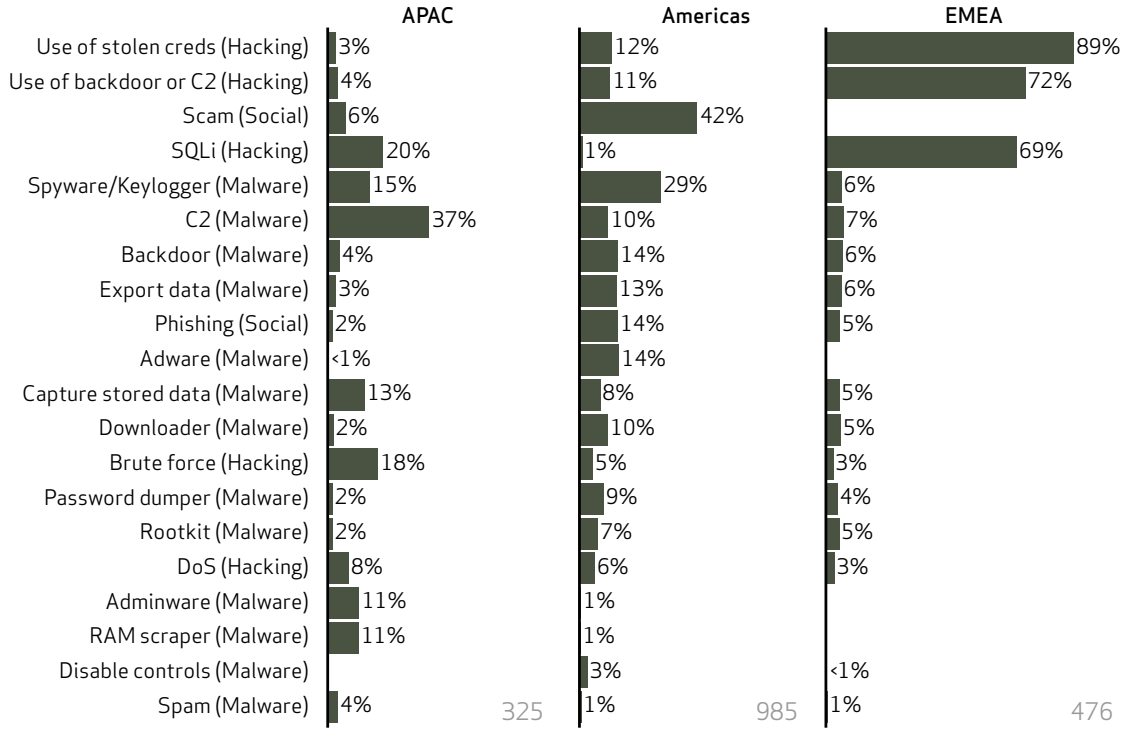

Malware (40% of breaches)

Malware is any **mal**icious soft**ware**, script, or code added to an asset that alters its state or function without permission. The percentage of data breaches involving malware was lower in 2012, but that can be attributed to a relative proportional increase in other categories (social and physical) rather than an actual decline. Malware still ranks in the top two threat actions in our dataset, which speaks to its attractiveness and effectiveness as a tool.

Malware can be classified in many ways, but VERIS uses a two-dimensional approach that records the infection vector and functionality. The theory is that if you know how malware gets on systems and what it tends to do, you're in a good position to make informed decisions about protecting the enterprise.

In a streak that remains unbroken, direct installation of malware by an attacker who has gained access to a system is again the most common vector. And that makes sense; once you own the system, it'll need some fancy accessories. For smaller companies, this often (but not exclusively) plays out as a "POS smash-and-grab" scenario like those mentioned in the previous section.

The primary difference this year is malware distributed via e-mail attachments, and Figure 20 makes it easy to see this is a direct reflection of the increased espionage cases in our dataset. We also see some remnants of the strategic web compromises (aka "watering hole attacks") publicized this year, which utilize drive-by exploits to download malware. This vector is more prevalent among larger organizations in both espionage and financially motivated attacks.

Keep in mind that these vectors are not mutually exclusive. In many cases, an actor may gain initial entry using a malicious e-mail attachment, and then install additional malware on that and other systems throughout the environment.

IN A STREAK THAT REMAINS UNBROKEN, DIRECT INSTALLATION OF MALWARE BY AN ATTACKER WHO HAS GAINED ACCESS TO A SYSTEM IS AGAIN THE MOST COMMON VECTOR. AND THAT MAKES SENSE; ONCE YOU OWN THE SYSTEM, IT'LL NEED SOME FANCY ACCESSORIES.

The second main malware characteristic describes its variety or functionality, and many breaches incorporate quite a few of them (thus, the percentages add up to well over 100). While VERIS defines quite a few specific functionalities, there are three general goals that most malware seeks to achieve within the context of data breaches: 1) grant or prolong access and control, 2) capture data, or 3) weaken the system in some way (usually to enable the first two, so maybe that's two main goals).

Figure 20: Vector for malware actions

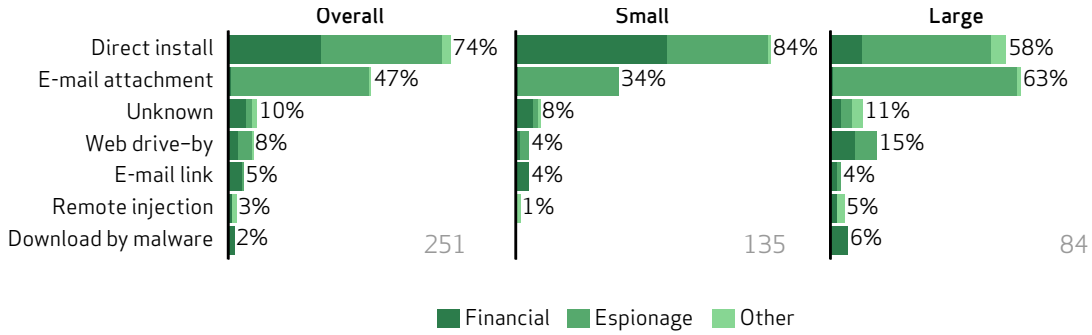

Spyware (which includes keyloggers and form-grabbers) posted the highest percentage, but many other varieties are bunched at or near the top of Figure 21. And that in itself is an interesting observation. We hypothesize that the less scripted, more interactive nature of targeted attacks evens out the playing field a bit. Moreover, the individual pieces of malware used in such attacks are often multi-functional. Figure 21 also illustrates well that malware functionality tends to differ based on victim size and attacker motives.

In the land of financially motivated breaches, spyware is king. Capturing data from payment cards swiped at POS terminals and credentials typed into online bank accounts are two very popular uses of these tools in cybercrime. As an aside, the use of spyware differs in espionage, where it focuses on grabbing screenshots of potentially valuable information and capturing user credentials to further spread the attack. RAM scrapers and network/system utilities ("adminware") are also major players in the financial crime space, and especially so in smaller organizations. The former makes sense, but we suspect the latter might be somewhat under-reported among breaches tied to espionage. That said, state-affiliated actors have the resources to create their own specialized tools rather than repurposing legitimate utilities that might be recognized. As they have been in the past, backdoors, C2, and data capture/export features remain popular with professional criminal groups as a way of setting their hooks in and getting their loot out.

IN THE LAND OF FINANCIALLY MOTIVATED BREACHES, SPYWARE (WHICH INCLUDES KEYLOGGERS) IS KING. ON THE ESPIONAGE SIDE OF THE OCEAN, THERE IS NO KING; THERE IS INSTEAD FAIRLY EQUAL REPRESENTATION ACROSS MANY MALWARE VARIETIES.

On the espionage side of the ocean, there is no king of malware varieties; there is instead fairly equal representation across them all. State-affiliated actors often use the same formula and pieces of multi-functional malware during their campaigns, and this is reflected in the statistics throughout this report. In the last year, we most often saw a variant of AURIGA, which was described in great detail in a February 2013 report from Mandiant on Chinese espionage[18]. After infecting

Figure 21: Variety of malware actions

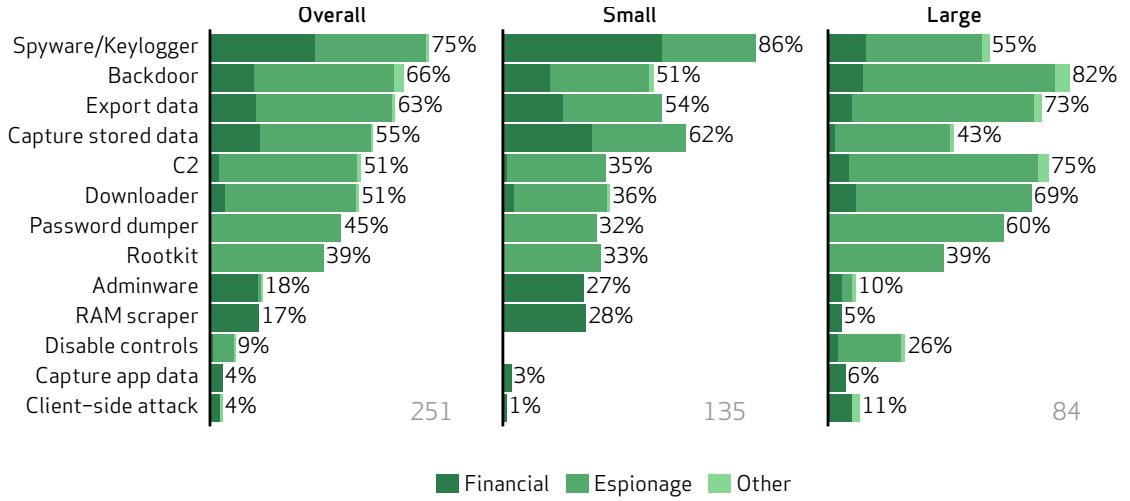

18 http://intelreport.mandiant.com/Mandiant_APT1_Report.pdf

the victim's systems (usually through a phishing e-mail), attackers utilize the backdoor and C2 features to download additional malware and move towards the goal of getting domain-level access. This can be accomplished via keyloggers, capturing credentials stored on end-user systems, or dumping password hashes from the domain controller. Throughout this process, attackers promulgate across the systems within the network, hiding their activities within system processes, searching for and capturing the desired data, and then exporting it out of the victim's environment.

Before we wrap up this section, there are a few miscellaneous points we'd like to mention. First, it's curious that we had no reports of malware disabling security controls in small organizations. Then again, having reviewed "controls" in numerous retailers and restaurants, perhaps that's not so curious after all. Also interesting is that we rarely see password dumpers in financial breaches. This could be under-reported because it falls below the detail threshold tracked by other organizations. It could also be due to a lesser need to "own the environment" and the tighter window of opportunity limited by fraud algorithms that begin to hone in on breach victims once criminals begin profiting from their schemes.

YOUR MONEY JUST RAN SOMEWHERE

We just made you say "ransomware." Ransomware attacks occur when criminals break into the victim's computers and encrypt all data on the system, rendering it inaccessible unless a fee is paid in exchange for the decryption key. Without that key, they're out of luck (and out of data), and this persuades many ransomware victims to pay up. This may be why some sources see ransomware schemes blossoming as an effective tool of choice for online criminals targeting businesses and consumers.

When targeting companies, typically SMBs, the criminals access victim networks via Microsoft's Remote Desktop Protocol (RDP) either via unpatched vulnerabilities or weak passwords. Once they've gained initial access they then proceed to alter the company's backup so that they continue to run each night but no longer actually backup any data.

After a period of weeks, the criminals return to the server and then encrypt it. When the victim tries to access their system, an on-screen message notifies them it is encrypted and that backups are no longer available. Much like the individual scenario, the criminals demand a ransom to supply the key to access the data. If the ransom goes unpaid, the data remains encrypted and may even be deleted from the server. Business owners are then left with the option to either lose their data or pay the ransom demand. In many cases, the business decides to pay the ransom, recover the data, and then implement improved security controls.

The risk of ransomware attacks could be reduced by:

- Ensuring remote access solutions are patched with the latest security software
- Mandating strong passwords for remote access and, where possible, implement two-factor authentication
- Confirming that backups have completed successfully and that the data is available on the backup media
- Keeping all systems up to date with the latest anti-virus software and updates
- Keeping all systems patched with the latest software
- Training users to be aware of the security risks when interacting online

VERIS VEERS DOWN BROAD STREET

And while we're on the subject of malware, let's pause for a moment to consider a different approach to analyzing cases involving this type of threat action. Our colleagues at Microsoft discuss the current state of malware within the Microsoft Security Intelligence Report[19] and have developed a new taxonomy for malware propagation behavior called Broad Street.[20] As researchers with a common interest and a desire to be helpful, we have tried to continue this conversation and extend this effort by identifying areas where our VERIS Framework and the project Broad Street taxonomy overlap, as well as some comparison of findings.

In Figure 22, we have reproduced (with Microsoft's permission) the flow used within the Broad Street taxonomy. Within the boxes are the percentages of VERIS malware vectors within our 2012 breach dataset that map to parts of their taxonomy.

Note that some of our malware vectors do not fit, and for good reason. Vectors such as "Coded into existing program/script," "Installed by other malware," "Installed by remote attacker," and "injected by remote attacker" all need some level of access prior to infection. This means that the taxonomy doesn't apply in those cases because it specifically focuses on the initial entry point onto a system. About 42% of the incidents involving malware in our dataset over the last three years can be classified by the Broad Street taxonomy.

Figure 22: Broad Street flow chart with percentages from DBIR dataset

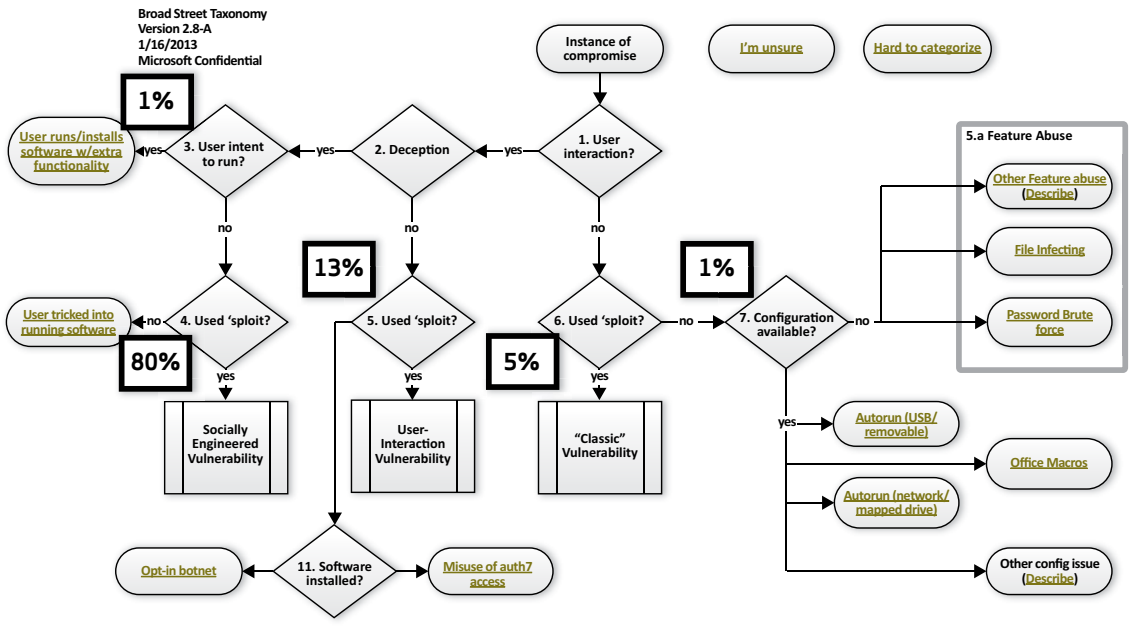

19 http://www.microsoft.com/sir
20 http://download.microsoft.com/download/0/3/3/0331766E-3FC4-44E5-B1CA-2BDEB58211B8/Microsoft_Security_Intelligence_Report_volume_11_English.pdf

Conversely, the VERIS enumerations for malware vectors, by themselves, do not have the same level of granularity to reach some of the classifications in the Broad Street taxonomy. For example, when a user is deceived in Broad Street's taxonomy, we need to look outside of the Malware action to Social actions. So we can only map our vectors onto entire classes of endpoints within their taxonomy (User Interaction, Vulnerability, and Configuration). When we do that and compare data, we see the following proportions:

PROPAGATION METHOD	MICROSOFT	VERIZON
User Interaction	45%	84%
Classic Vulnerability	6%	5%
Configuration	50%	1%

The Microsoft data comes from their latest published Broad Street numbers (in volume 11 of the SIR); the distribution may have changed since then. Our proportions vary from theirs significantly, particularly in terms of malware that requires user interaction and where malware depends on misconfigurations such as Autorun. For user interaction, we believe this is partially due to the significant number of state-affiliated espionage cases that utilize deception to get a user to click on that link or open that attachment. For the misconfigurations, we believe that our numbers are low because most worms that rely on misconfiguration tend to be quite noisy. If you were attacking an organization, siphoning data away from them, you would probably want to remain quiet rather than setting off hundreds of noisy alarms bells. It could also be due to how VERIS classifies incidents; each incident may contain many infections. For example, an incident that involves 200 infected systems gets counted the same in our methodology as another incident that only involves four. Therefore comparisons are rather difficult to make.

Regardless of the issues, we believe that the project Broad Street taxonomy provides a useful tool in examining cases involving malware.

Hacking (52% of breaches)

Hacking includes all attempts to intentionally access or harm information assets without (or in excess of) authorization by circumventing or thwarting logical security mechanisms. The Internet enables many hacking methods to be highly scalable, automated, and conducive to anonymity. This section examines the varieties and vectors of hacking observed in the 2012 dataset.

"It's déjà vu all over again" is a fitting Yogi-ism for Figure 23; it does indeed look familiar. And there's good reason for it—the easiest and least-detectable way to gain *unauthorized* access is to leverage someone's (or something's) *authorized* access. Why reinvent the wheel? So, it really comes as no surprise that authentication-based attacks (guessing, cracking, or reusing valid credentials) factored into about four of every five breaches involving hacking in our 2012 dataset. Nor is it all that surprising that we see this year after year.

If data could start a riot ("Occupy Passwords!"), we could use these statistics to overthrow single-factor passwords: the supreme ruler in the world of authentication. If we could collectively accept a suitable replacement, it would've forced about 80% of these attacks to adapt or die. We've talked about the shortcomings of passwords for years now, and if it were an easy problem (or the pain caused by password problems was greater), it'd be fixed by now.

IF DATA COULD START A RIOT ("OCCUPY PASSWORDS!"), WE COULD USE THESE STATISTICS TO OVERTHROW SINGLE-FACTOR PASSWORDS: THE SUPREME RULER IN THE WORLD OF AUTHENTICATION. IF WE COULD COLLECTIVELY ACCEPT A SUITABLE REPLACEMENT, IT WOULD'VE FORCED ABOUT 80% OF THESE ATTACKS TO ADAPT OR DIE.

Figure 23: Variety of hacking actions

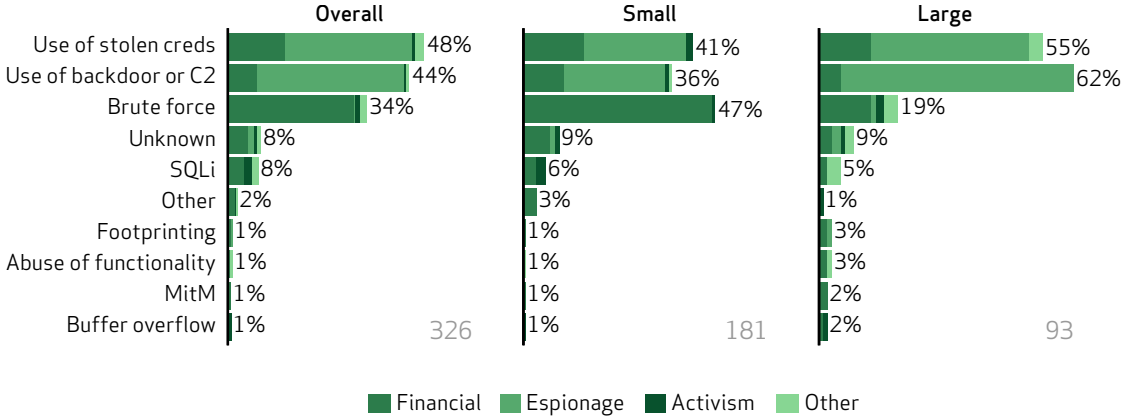

What you can't see from Figure 23 is that VERIS contains more than 40 varieties of hacking. Considering that, the fact that nearly all activity in this threat category is accounted for by a mere five of them is remarkable. Whether this is because attackers don't often use the rest of them or because these are the most successful (in causing breaches) is an interesting research question.

Figure 24: Vector for hacking actions

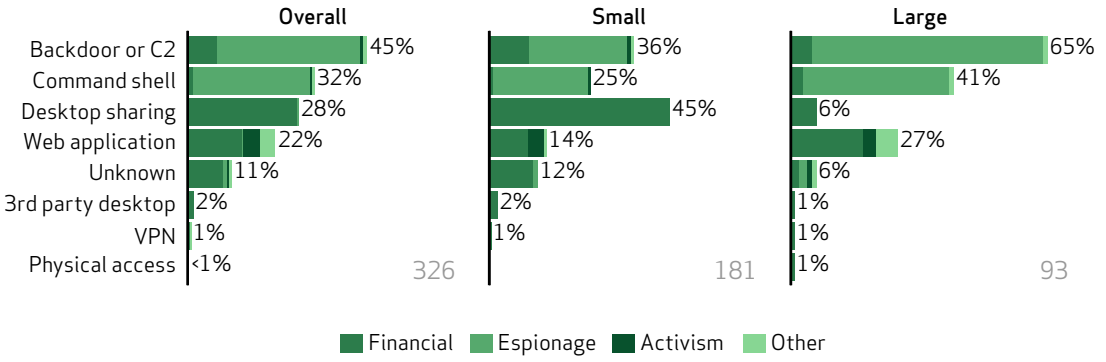

Brute force attacks continue to disproportionately affect smaller organizations. Perhaps larger organizations have more mature processes in place that reduce the opportunities for an attacker to find such an easy method of entry. Of course, that doesn't mean that they do well in this area, just not as poorly. Nearly one-fifth of them still fell victim to the easiest trick in the script-kiddie's toolkit.

Readers will reasonably ask how attackers steal credentials in order to reuse them to gain unauthorized access. Sometimes users are socially engineered to give them up. Sometimes malware captures them from keystrokes, browser cache, or system files. We recommend reviewing the Social and Malware sections for additional discussion and examples on this.

We hit the topic of backdoors in the Malware section, and so will simply clarify here that this relates to threat actors gaining access via a backdoor rather than just malware that includes backdoor functionality. It continues to be an extremely popular and effective way—especially against larger organizations—to circumvent controls and establish many "hooks" into a victim's environment that can be quite difficult to detect.

Figure 24 does an excellent job contrasting techniques with the motive of those who used them. Stolen credentials and backdoors are heavily used in targeted espionage campaigns, while brute force is the tool of choice for financially motivated groups. Activists are a much smaller minority in this dataset, but it is possible to discern their proclivity for brute force and SQL injection. We find this kind of analysis extremely interesting and think it holds great promise for better understanding and countering our adversaries.

What happened to exploitation of default or guessable credentials, the top Hacking variety in the 2012 DBIR? Recent updates to VERIS merged it with brute force and dictionary attacks. From a forensics and incident sharing standpoint, the rather fuzzy line between default/guessable and brute force was causing issues with data consistency. For example, if an attacker uses a tool scripted with the default password for different vendors, should that be recorded as brute force or default credentials or both? Since we weren't getting consistent answers to that and similar questions, we combined these varieties into one. Ultimately, both involve weak credentials, so there is little lost and accuracy gained in combining them.

Changes among hacking vectors largely follow relative shifts in attack genres in this year's dataset. Backdoors and remote shells like Secure Shell (SSH) and Remote Procedure Call (RPC) are up due to their role in targeted attacks. The methodical workflow honed by many state-affiliated actors of setting up a backdoor to gain initial access, and then using shell services to move laterally through the organization, has proven to be successful against victims of all types and sizes. Among financially motivated attacks, desktop-sharing services like Remote Desktop Protocol (RDP) and Virtual Network Computing (VNC)—classic vectors in Point of Sale (POS) hacks—are proportionally down along with those crimes. Web applications are up overall, but are no longer the leading attack vector among larger organizations, as they were in our last report.

Social (29% of breaches)

Since crime began, criminals have consistently taken advantage of human nature to advance their dark enterprises. Stealing information is no different, and threat actors are well aware of the flaws in the carbon layer and the tactics used to exploit them. For example, sending a convincingly crafted malware-laden e-mail to a few key employees could give an attacker the keys to a company's intellectual property kingdom. The 2012 data reveals a big upswing in scenarios like this.

In the last year, phishing jumped bribery and pretexting to become the most widely used social tactic. To what can we attribute this impressive leap? Figure 25 leaves little doubt that state-affiliated espionage is the main driver. In fact, more than 95% of all attacks of this genre employed phishing as a means of establishing a foothold in their intended victims' systems. And we're not the only ones seeing it that way; another recent vendor report put that statistic at 91%[21]. That phishing is similarly prevalent for both small and large organizations is important to note.

MORE THAN 95% OF ALL ATTACKS TIED TO STATE-AFFILIATED ESPIONAGE EMPLOYED PHISHING AS A MEANS OF ESTABLISHING A FOOTHOLD IN THEIR INTENDED VICTIMS' SYSTEMS.

Turning our attention to breaches motivated by financial gain reveals a fairly balanced use of extortion, bribery, and phishing with some pretexting thrown in for larger organizations. Criminals chasing the cash know there's more than one proven way to steal credentials or trick a money handler into committing fraud.

Save for the labels, Figure 26 could be mistaken for Figure 25 below. And there's a good reason for that. The fact that e-mail is the most common vector of social attacks should be no surprise, since phishing is the most common variety. Similarly, criminals wanting to bribe, extort, or con a target often opt for a good old-fashioned phone call or an even older-fashioned face-to-face. Though there's been much discussion regarding the use of SMS communications as a vector of social engineering (aka SMiShing), we're not seeing much activity there.

Figure 25: Variety of social actions

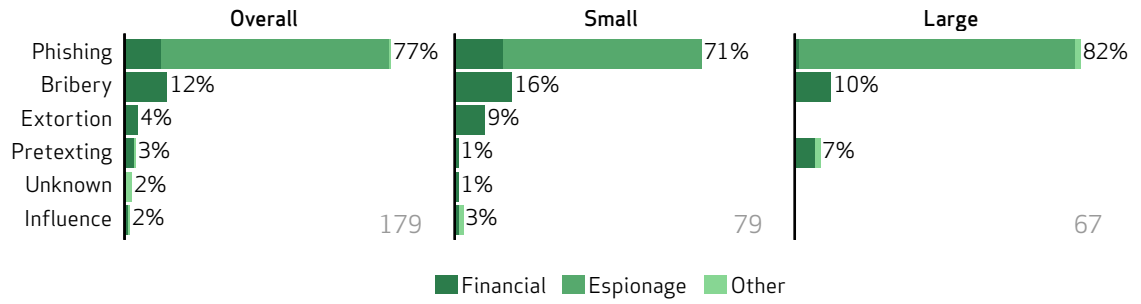

21 Trend Micro - Spear-Phishing Email: Most Favored APT Attack Bait – http://www.trendmicro.com/cloud-content/us/pdfs/security-intelligence/white-papers/wp-spear-phishing-email-most-favored-apt-attack-bait.pdf

Figure 26: Vector for social actions

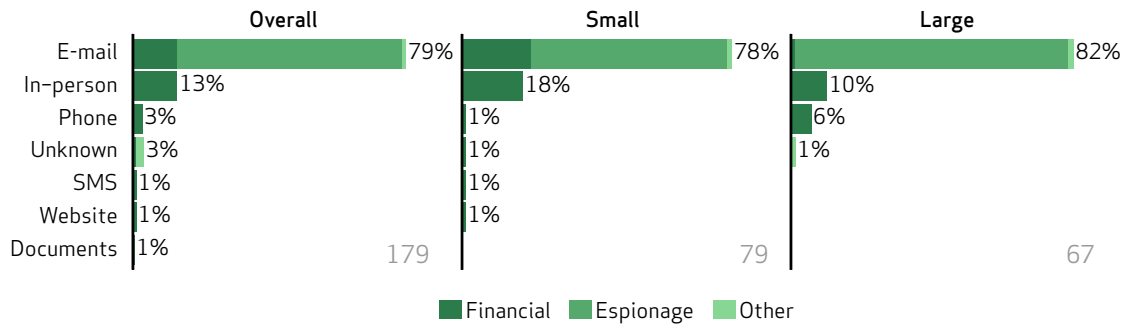

When it comes to the targets of social engineering, the data reveals…not very much. Unfortunately, the positions and titles held by those in the crosshairs of espionage campaigns are unknown to us. There are several reasons for this. Many of these result from network-level intelligence and/or remote investigations. In such cases, we're able to detect malicious communications to compromised systems, and can reconstruct events well enough to determine infections resulted from phishing e-mails, but cannot (or aren't asked to) identify patient zero who clicked the attachment. Exacerbating this is that many victims have been compromised for a long time and relevant logs have long since passed into the ether. Another reason for this high proportion of unknowns is that such details are not contained in the information provided by our contributors.

Beyond "unknown," the next two groups on the list are cause for concern. Executives and managers make sweet targets for criminals looking to gain access to sensitive information via spear phishing campaigns. Not only do they have a higher public profile than the average end user, they're also likely to have greater access to proprietary information. Plus, we all know how much they love .ppt and .pdf attachments.

Figure 27: Targets of social actions

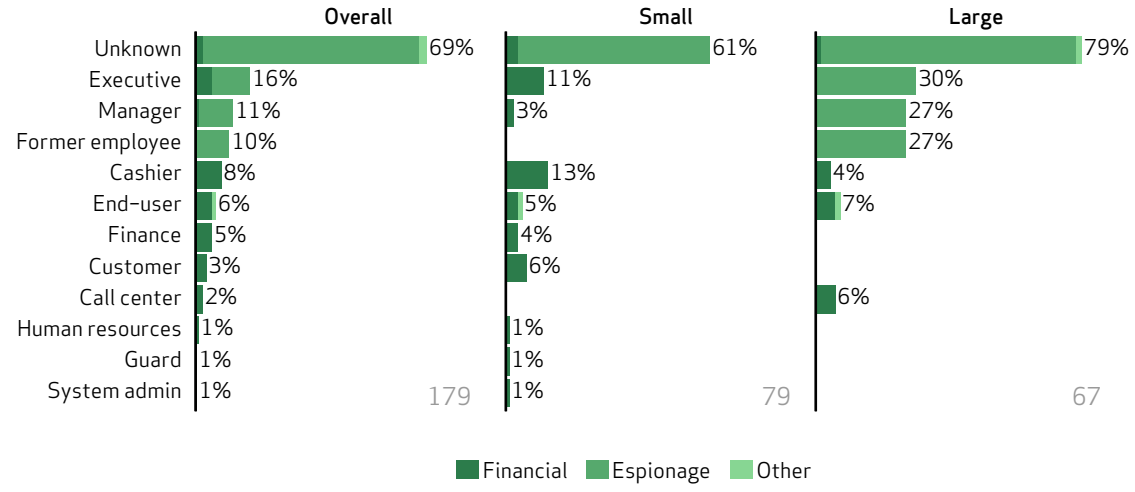

THE INEVITABILITY OF "THE CLICK"

We try to avoid rolling out scary memes like "you will be compromised," but when it comes to phishing attacks, that's exactly what the data tells us.

Phishing e-mails vary in quality, payload, and purpose, but they all share the same initial goal: get the user to take action. Getting the user to click (on a link or attachment) is the first obstacle for all phishing campaigns. So how many e-mails would it take to get one click? Figure 28 attempts to answer that question based on data collected by ThreatSim in their phishing-for-hire campaigns.

It's pretty easy to see why this is a favored attack for espionage campaigns and the answer to our question is "three." Running a campaign with just three e-mails gives the attacker a better than 50% chance of getting at least one click. Run that campaign twice and that probability goes up to 80%, and sending 10 phishing e-mails approaches the point where most attackers would be able to slap a "guaranteed" sticker on getting a click. To add some urgency to this, about half of the clicks occur within 12 hours of the phishing e-mail being sent.

So that's the bad news. The good news is that a user clicking does not automatically lead to a compromise. A successful phishing campaign requires a series of "and" statements for every step in a campaign. With each added step, the probability of a system compromise goes down. For example, a user needs to take action AND there needs to be a vulnerability on the system AND software has to be quietly installed AND there has to be a communication path back to the attacker, and, and, and this is why we have the term "defense in depth."

At each phase in the attack we want to increase the probability of detection and decrease the probability of success.

And while an eventual click might be inevitable, evidence in the Discovery Methods section of this report offers some additional good news. It suggests that equipping employees to recognize and report suspicious occurrences (like phishing e-mails) and monitoring for beacons out to malicious IPs ~~if~~ when they do click on them could be some of the most effective means of discovering a breach.

Figure 28: The inevitability of the click

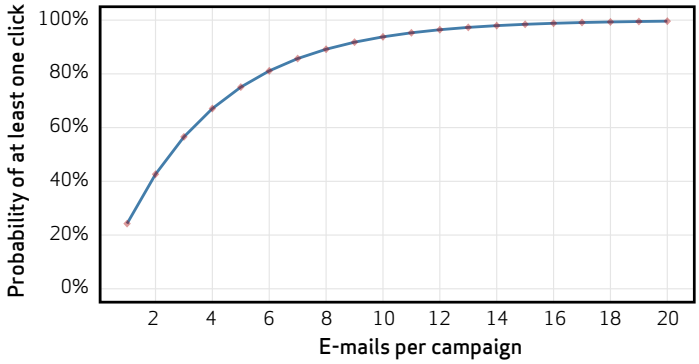

Misuse (13% of breaches)

Organizations expend significant resources trying to hire the best staff. They want someone who is trustworthy, competent, and works well with others. Unfortunately, that's not always what they get, and granting anyone access to confidential information carries the risk of abuse. When privileged parties maliciously or inappropriately use organizational resources in ways they should not, VERIS classifies these actions as misuse.

The top three varieties of misuse from 2011—embezzlement, use of unapproved hardware, and privilege abuse—remain at the top in 2012, but shuffle around as they jockey for the top position. Abuse of system privileges is particularly common in larger organizations, while embezzlement and unapproved hardware (often handheld card skimmers) tend to afflict the smaller ones. All three of these exhibit heavy financial motives. The espionage shading among larger organizations hints at how insiders get the data (abuse their access) and how they get it out (smuggle it on unapproved media devices).

Figure 29: Variety of misuse actions

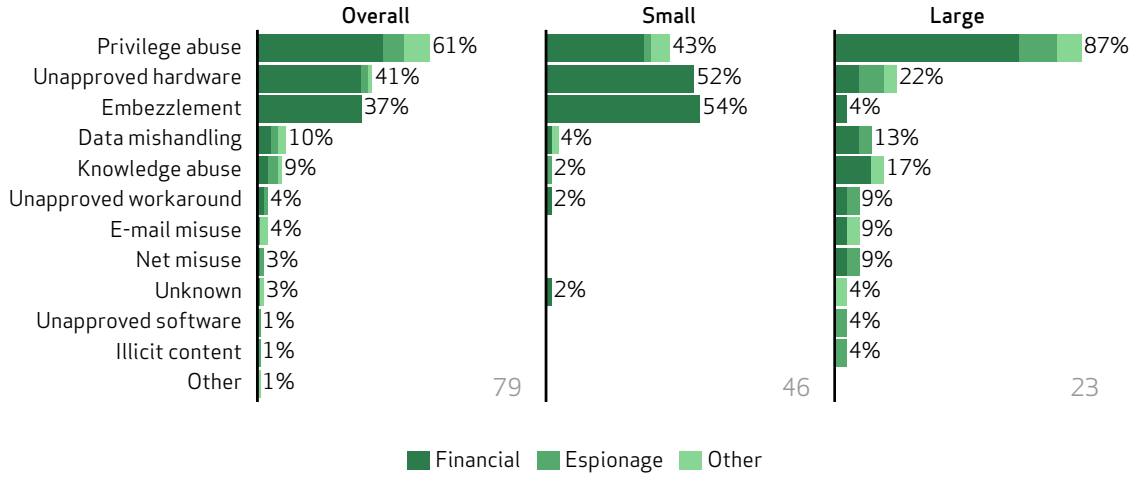

Despite repeated calls to immediately disable access as part of standard employee termination processes, cases of disgruntled ex-employees abusing still-active privileges (treated as external misuse in VERIS) make our dataset every year. We'll say it again—if you take them off your payroll, take them out of your systems too.

Checking in on common vectors of misuse, we once again find breaches carried out through physical access within the corporate facility claim the majority. This vector tilts toward smaller organizations due to the card skimming scenarios discussed above. Employees of larger organizations more often conduct their illicit affairs via the corporate LAN. Given that remote access services post comparatively lower numbers, the concern that working from home is riskier than working inside a corporate facility seems unfounded.

DESPITE REPEATED CALLS TO IMMEDIATELY DISABLE ACCESS AS PART OF STANDARD EMPLOYEE TERMINATION PROCESSES, CASES OF DISGRUNTLED EX-EMPLOYEES ABUSING STILL-ACTIVE PRIVILEGES MAKE OUR DATASET EVERY YEAR. WE'LL SAY IT AGAIN—IF YOU TAKE THEM OFF YOUR PAYROLL, TAKE THEM OUT OF YOUR SYSTEMS TOO.

Figure 30: Vector for misuse actions

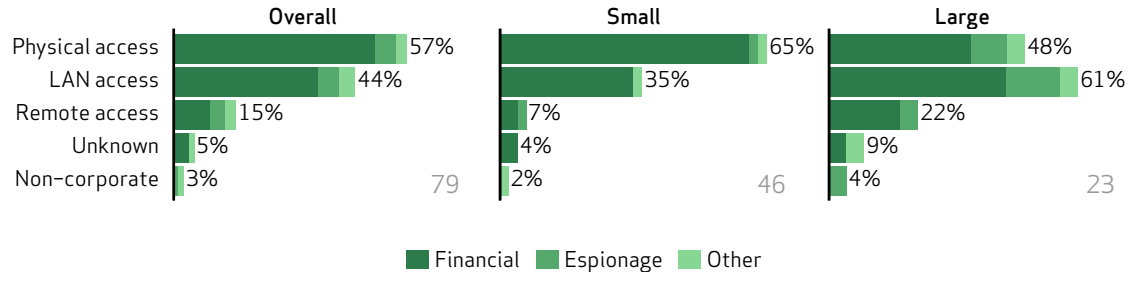

Physical (35% of breaches)

Physical threats encompass deliberate actions that involve proximity, possession, or force. It is important to disclose that physical threats are extremely common but underrepresented in this report. For instance, stolen user devices are less likely to receive a forensic investigation to confirm data compromise or fall under the jurisdiction of our law enforcement contributors.

While not to the extent of some automated network-based attacks, physical threats can be highly organized, and to a certain extent, scalable. This makes them a common modus operandi for financially motived criminal groups. ATM skimming operations (the most common type of physical data breach) frequently target numerous locations of a single large banking institution. The name on the ATM may be of little relevance to the group, but the hardware deployed for the ATMs is likely standardized. This means that skimmer overlays designed for the specific ATM can be built/purchased in numbers and installed across a relatively large geographic region in a single spree. These groups are beginning to leverage 3D printing technology to improve efficiency and adapt to changes in card reader design. As 3D printers become more and more accessible we expect to see more groups utilize this technology in their criminal endeavors.

Similar logic is applied to launching organized attacks involving POS device tampering. Armed with the knowledge that a retail franchise uses a particular card swipe/PIN entry device, groups acquire a set of these devices, install in-line skimmers, and swap them for legitimate ones. These "new" devices are identical in appearance and designed to continue to perform their intended functions, but they are also redesigned to capture payment card data. The stolen devices are then repurposed by again installing hardware, and the cycle continues.

SKIMMERS INSTALLED INSIDE ATM'S, POS DEVICES, AND GAS PUMP TERMINALS COMPRISE ALMOST ALL INCIDENTS IN THE PHYSICAL CATEGORY (SEE FIGURE 31). AND IT'S NOT SURPRISING THEN, THAT TAMPERING IS THE MOST COMMON VARIETY OF PHYSICAL ACTIONS.

The above scenarios, along with skimmers installed inside gas pump terminals, comprise almost all incidents in the physical category (see Figure 31). And it's not surprising then, that tampering is the most common variety of physical actions. ATM skimming operations sometimes combine surveillance in the form of pinhole cameras that capture user PINs as well as mag stripe information. Theft in Figure 31 is largely represented in the POS swap operations. The majority of the attacks are conducted in public areas, which makes perfect sense given the devices targeted are used by the banking and retail customers and must be accessible.

Figure 31: Variety of physical actions

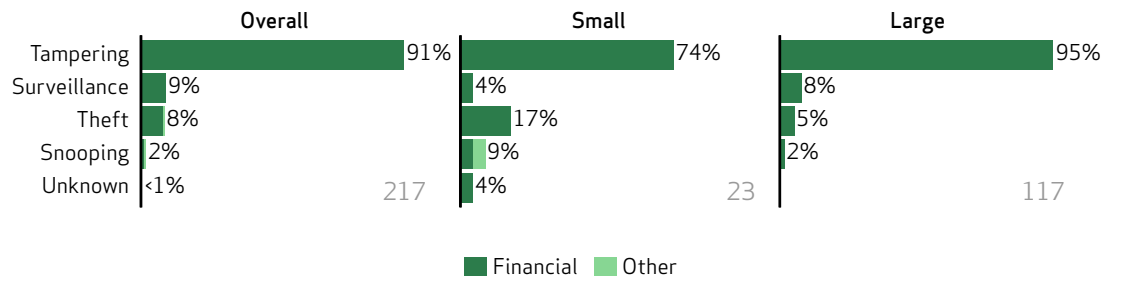

Error (2% of breaches)

Defining thresholds for error as a VERIS threat category has always proven somewhat problematic. If we labeled every poor decision or "oopsie" as a threat action, the resulting metrics would lack meaning and usefulness. Thus, we take a more conservative stance when classifying incidents and compiling statistics. We record an error as a threat action only if it deviates from normal processes within an organization and directly causes or significantly contributes to the incident. It hurts our hearts not to label a blank password as an error, but if the organization doesn't have processes or standards to forbid that and lacks fundamental security as "the norm," it's hard to call it an error. It's not a deviation from the norm. A server misconfiguration that publishes private data to a public website is a different matter, and would be recorded as an error.

Speaking of misconfigurations that directly expose private information, they comprised the majority of the 14 data breaches attributed to error. Not really much else we can unpack there; organizations rarely ask a third party to investigate incidents resulting from mundane mistakes or glitches.

Apparently, however, they very often report them to CSIRTs. As previously mentioned, erroneous delivery of e-mails and documents was the leading threat action among the 47,000+ security incidents we studied from 2012.

Environmental

The environmental category includes natural events such as earthquakes and floods, but also power, water, temperature, and other hazards associated with the immediate environment or infrastructure in which assets are located. Though legitimate threats that must be managed, they rarely directly affect data confidentiality. As one can readily imagine, the environmental section is one of the most anticipated and highly read in our report each year (are you picking up on the sarcasm?). Two great examples of this category from the last few months do come to mind, however. The first is the apocalypse predicted by the Mayans to take place on December 21, 2012. Admittedly, this did not actually occur but, if it had, can you imagine the number of records that would have been lost?

The second event is the meteorite that fell from the sky and exploded above Russia on February 15, 2013. While no serious data outages occurred from this event, and more fortunately, no one was killed, it had the potential to cause widespread damage to information assets. Thus, while rare, we feel it's important to continue including this section in our annual report.

Compromised Assets

Examining details surrounding compromised assets continues to support practical observations about the targets of the actors and actions we've been studying. It shifts us to the closer-to-home question of "what do I have that needs protecting and what happens if I don't?"

After dancing with destiny for several years now (see Figure 32), end-user devices have finally pulled off a prize-winning performance. But the people's champion—"People"—gained the most ground in the standings, thanks to some dazzling, but "phishy" moves by a number of contestants in 2012.

The closer distribution among the categories in Figure 32 hints at more breaches involving different kinds of assets. Single-asset attacks like POS hacks and ATM tampering still occur regularly, but more complex scenarios exist as well. This is the case in the classic spear phishing attack that involves a person (phishing) and their desktop (malware) on the way to rooting a server or two (or 200). These broad categories are worth a quick look, but they're hiding some important details regarding the variety of assets compromised in 2012 breaches. Figure 33 lets in some more light so we can really see what's going on.

Figure 32: Compromised asset categories over time

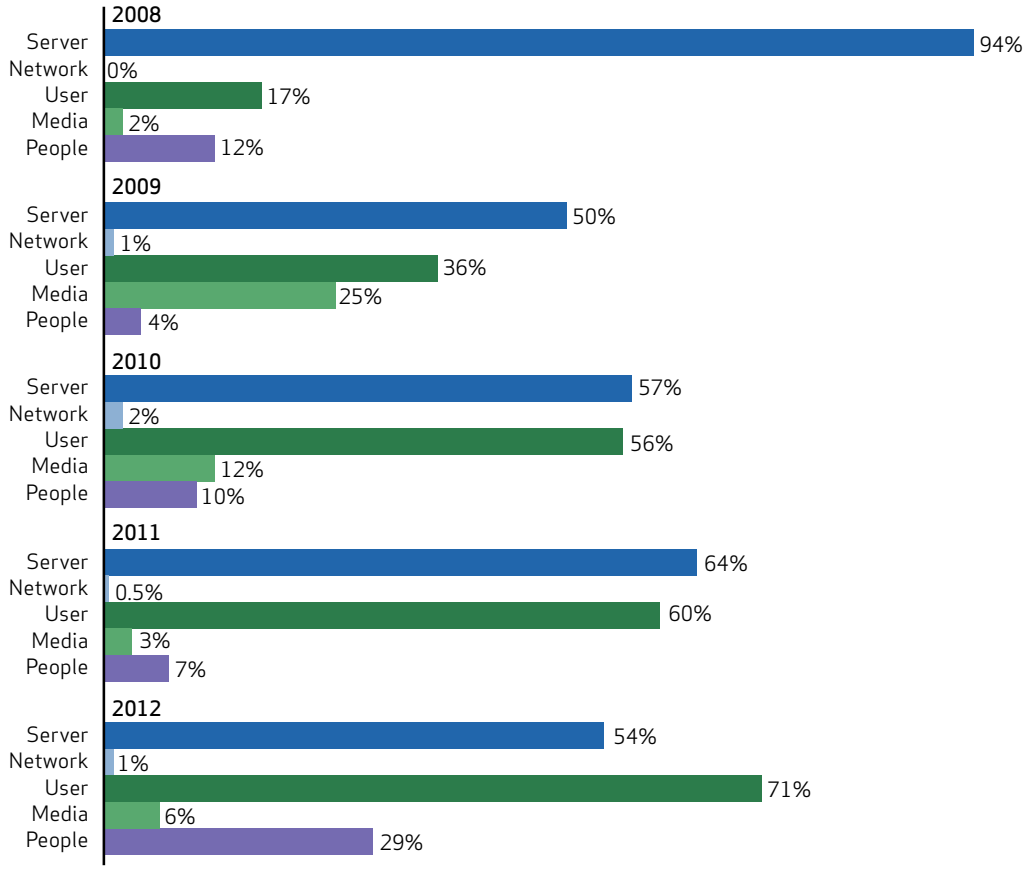

Starting at the top of the list, we see the secret to the end-user devices' success. Three of the four most compromised asset varieties fall within that category, and ATMs lead them all. We discussed ATM skimming in the Physical section, so we won't do that again here. Just note that data thefts involving them are money-driven (shocker), and relevant only to certain kinds of organizations (e.g., larger banks).

Beyond ATMs, the next six asset varieties largely reflect standard targets in espionage campaigns. The standard event chain of phishing (other/unknown people, desktop, laptop), expanding control (directory), and exfiltration of data (database and file servers) is clear. Worth clarifying, though, is "other/unknown server," which reflects the "own the environment" nature of targeted attacks and the fact that we did not do full "boots on the ground" forensics investigations for all of them. Depending on

BEYOND ATM'S, THE NEXT SIX ASSET VARIETIES LARGELY REFLECT STANDARD TARGETS IN ESPIONAGE CAMPAIGNS. THE STANDARD EVENT CHAIN OF PHISHING (OTHER/UNKNOWN PEOPLE, DESKTOP, LAPTOP), EXPANDING CONTROL (DIRECTORY), AND EXFILTRATION OF DATA (DATABASE AND FILE SERVERS) IS CLEAR.

Figure 33: Variety of compromised assets

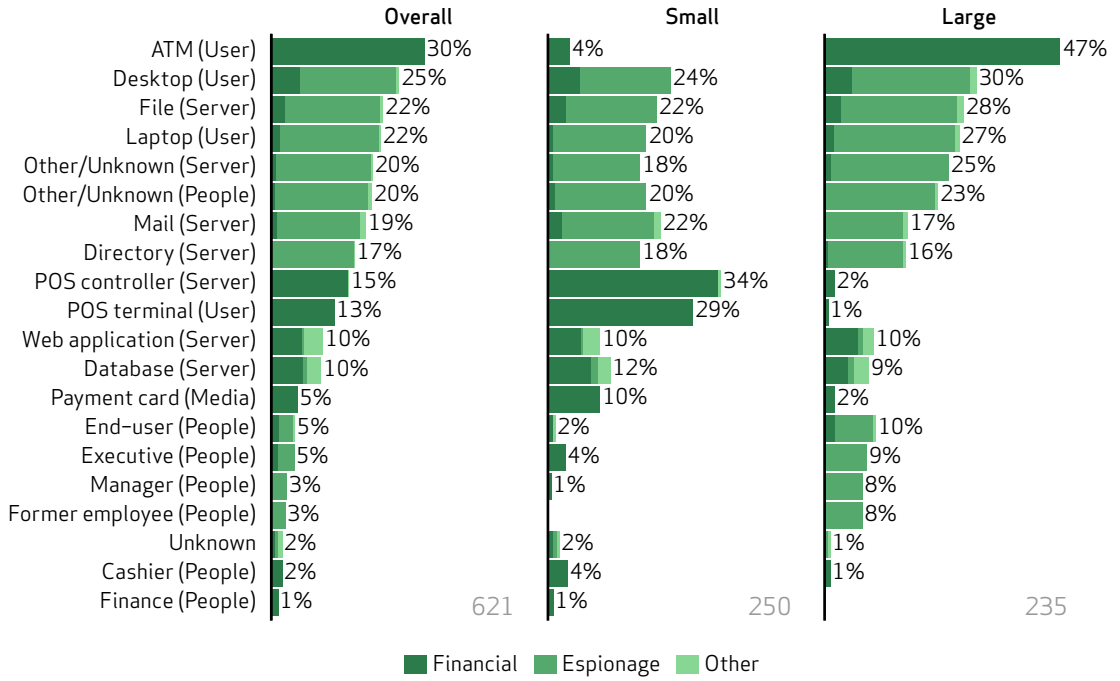

the victim's desire, remediation often involves identifying and redeploying whole network segments rather than enumerating and whack-a-mole-ing individual systems.

Next in line are point of sale controllers and terminals. These are a favorite of financially motivated organized criminal groups looking for a quick score of payment cards from smaller franchises. We've devoted much attention to this in past reports (and in prior sections of this one) and so will save some trees/kilobytes instead of rehashing it all. Rest assured, this kind of crime is alive and well.

Web application and database servers form another logical grouping, and once again account for most of the records breached. That makes sense because, well, those assets store a lot of records. These assets are targeted, almost equally, by financially motivated actors and actors engaged in espionage. It's interesting that they are the most "balanced" asset varieties in terms of actor motivation and organization size.

Rounding out the lower section of Figure 33 are payment cards and a slew of different people varieties. The payment cards are associated with the cashier/waiter embezzlement scenario described in the Misuse section. The people listed represent targets of social engineering schemes.

One might find it curious that no industrial control systems (ICS) appear in Figure 33, since we had two contributors that focus on that space. We did indeed receive incidents involving these types of assets, but they were not cases of confirmed data compromise (most were malware infections only). Our 2012 breach data does include victims that use ICS (quite a few related to state-affiliated espionage), but other varieties of assets were compromised.

MOBILITY, CLOUD, AND BYOD, YIPPEE! We also wanted to acknowledge two subjects that don't feature prominently in our research: mobile devices and ownership, hosting, and management. With respect to mobile devices, obviously mobile malware is a legitimate concern. Nevertheless, data breaches involving mobile devices in the breach event chain are still uncommon in the types of cases Verizon and our DBIR partners investigate. However, we do expect them to make more of an appearance in our data as mobile payment systems continue to become more common.

The "Bring Your Own Device" (BYOD) trend is a current topic of debate and planning in many organizations. Unfortunately, we don't have much hard evidence to offer from our breach data. We saw only one breach involving personally-owned devices in 2011 and a couple more in 2012. We'll keep watching.

In the 2012 dataset we saw many cases that involved devices hosted and/or managed by third parties. However, the fact that these devices were in some form or fashion "in the cloud" was not a significant cause of the data breach, nor did it cause the devices to be more targeted. In other words, attacks against the virtualization technology were not present, but attacks against weakly configured devices that happened to be hosted in an external location were common—but not any more common than among internally-hosted ones. Rest assured, we're still tracking "cloudy things" like hosting and management, and we'll write about them more when we have some useful observations to share.

Figure 34: Compromised asset categories across 47,000+ security incidents

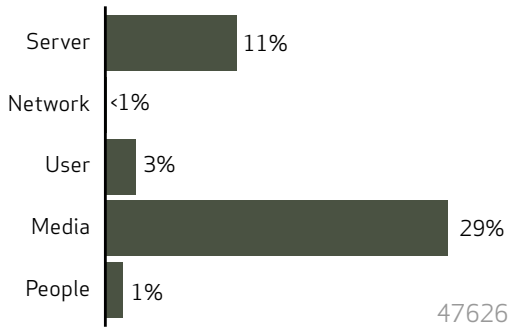

It's uncanny how this dataset of 47,000+ security incidents seems to conflict with our breach-specific findings at every possible point. But actually, these results track with what we've said before: these incidents represent the kind of stuff that happens in offices across the globe every day rather than the "tip of the iceberg" that requires external forensic or law enforcement support. The media category is so high here because of lost, mis-delivered, and stolen documents and faxes. Not the kind of stuff that makes the headlines, but it is the kind of stuff that exposes sensitive corporate data day in, day out.

Compromised Data

In one sense, this section is the focal point of the entire report. After all—if we didn't have data, we wouldn't have data breaches. We might be headed toward a paperless society, but we're sure as heck not headed for a dataless one. And that means we have to figure a way to keep our data while keeping it away from those who would abuse it.

Speaking of those who would abuse your data, we think Figure 35 is the blue ribbon winner in this section, so we're going to lead with it. The intersections, representing the number of breaches for each pairing, show a striking correlation between threat actor motives and the variety of data compromised. The profiteers favor payment and personal information that can easily be converted into cash. Spy types prefer trade secrets (e.g., schematics), internal organizational data (e.g., e-mails and memos), and system information. Hacktivists like the titillating aspect of personal information and internal organizational data. Credentials are fun for the whole family.

Figure 35: Breach count by data variety and actor motive

	Payment	Bank	Credentials	Personal	Medical	Classified	Copyrighted	System	Internal	Secrets	Other	Unknown
Financial	376	37	100	47	1		2	7	10	6	1	13
Espionage	1	1	119	1	1	3	1	113	122	119		21
Activism	2		3	8	1			2	4			
Other	1	1	14	6			1	2	10			8

THERE'S A STRIKING CORRELATION BETWEEN THREAT ACTOR MOTIVES AND THE VARIETY OF DATA COMPROMISED. THE PROFITEERS FAVOR PAYMENT AND PERSONAL INFORMATION THAT CAN EASILY BE CONVERTED INTO CASH. SPY TYPES PREFER TRADE SECRETS, INTERNAL ORGANIZATIONAL DATA, AND SYSTEM INFORMATION. HACKTIVISTS LIKE THE TITILLATING ASPECT OF PERSONAL INFORMATION AND INTERNAL ORGANIZATIONAL DATA. CREDENTIALS ARE FUN FOR THE WHOLE FAMILY.

To determine which varieties of data were compromised most often, you could add up the numbers in Figure 35, but you can also refer to Figure 37, where it's already done.

Payment cards have been a lock as the most oft-stolen data type since this study began, and 2012 was no different. They are the universal currency of the cybercrime marketplace. Theft of trade secrets and sensitive internal data display their highest-ever values (by far), bolstered by the espionage-related breaches in this dataset. Related to this, authentication credentials and system information show similar proportions.

It might be worth referring back to Figure 33 on page 43 of the 2012 DBIR to compare it to Figure 36. The older chart drew a distinction between varieties of data stolen from larger and smaller victims, whereas Figure 36 contains almost mirror images. Our best interpretation of this finding is the size-agnostic nature of many targeted attacks.

By now, you might have noticed that the little red values corresponding to compromised data records are not in any figures in this report. We didn't forget, and we're not hiding data, as some might assume. Each year we struggle with the decision of whether or not to include record counts as well as how much emphasis and credence we should give them.

Figure 36: Variety of compromised data

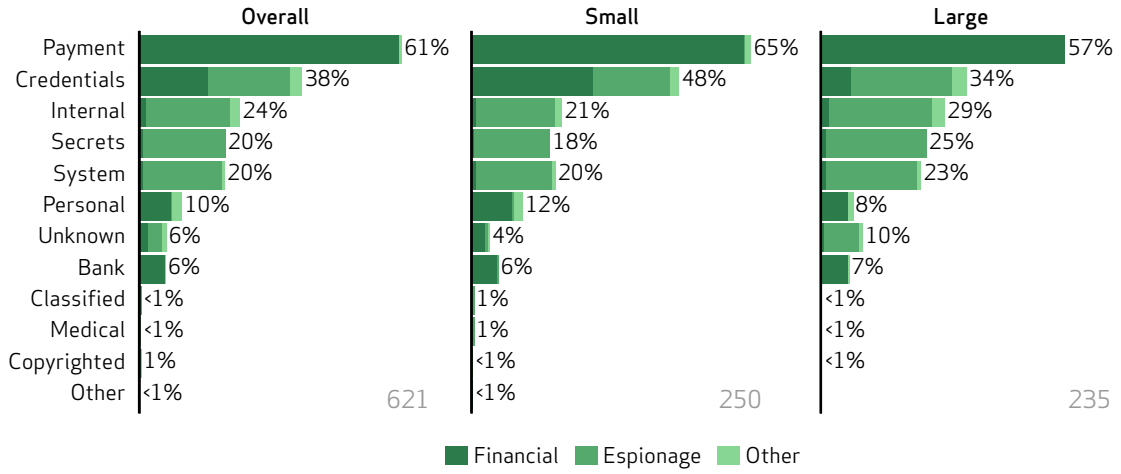

At the end of the day, one statistic persuaded us that any serious analysis based on records counts for this sample set would provide little value and may actually be misleading. Only 15% of breaches had a complete and reliable count of compromised records. That looming shadow of 85% unknown severely limits what we're willing to do with the data, but we can share Figure 37 with a clean conscience. It clearly and honestly shows what is known and what is unknown, and you can make up your own mind about what to do with these numbers.

Figure 37: Breach count by data variety and amount of records

	Payment	Bank	Credentials	Personal	Medical	Classified	Copyrighted	System	Internal	Secrets	Other	Unknown	**Total**
Unknown	115	22	220	28	2	3	3	124	143	124	1	42	**827**
1-100	160	11	11	12					1				**195**
101-1k	49			8					2				**59**
1k-10k	42	2	1	4	1								**50**
10k-100k	8	4	4	7			1			1			**25**
100k-1M	4			1									**5**
1M+	2			2									**4**
Total	**380**	**39**	**236**	**62**	**3**	**3**	**4**	**124**	**146**	**125**	**1**	**42**	

STATE OF THE DATA

In addition to the variety and amount, we track the state in which data existed when compromised—stored, transmitted, or processed. This is only done for Verizon IR cases. Two-thirds of breaches involved data stored or "at rest" on assets like databases and file servers. The other one-third was being processed when compromised. RAM scrapers, skimmers, and keyloggers that grab data in memory or when read/typed into a device are common examples of this. There were no instances in which data was compromised in transit. Owning a backbone router isn't a feasible plan of attack and packet sniffers haven't been common to our caseload recently. Naturally, breach scenarios can involve data compromised in multiple states, and this did indeed occur. More sophisticated espionage cases featured information theft at rest and in process; credential theft via keylogging malware followed by use of the stolen passwords to access a file server is a prime example. Again, the results provided here are for the Verizon caseload only; they would certainly change if we included all payment card skimming incidents contributed by our law enforcement partners.

While we're cautious about them, we don't want to give the impression that record counts are inherently worthless. There are valid reasons to track them and valid uses for them. For instance, there is likely a relationship between records lost and breach impact (at least for certain data varieties), and collecting the tallies enables us to test predictive models. They can also provide insight into the motivations and activities of our adversaries. Furthermore, some third parties (e.g., payment card brands) require total record loss to be reported by victims and/or investigators. So don't get us wrong—we'd love to have a more accurate accounting of compromised records across this dataset. Our main objection is putting too much faith in numbers that are so fraught with uncertainty

Armed with that knowledge, Table 2 can be appreciated for what it is. It tallies the total number of records compromised across all breaches each year. The 44 million posted for 2012 should be considered a lower bound of the true sum (because the full record loss was not known in 85% of those breaches). But it does serve as a wet finger in the wind annual comparative measurement for our dataset. Of note, most of that 44 million traces back to a very few large breaches (as is always the case).

Table 2: Data records compromised over time

YEAR	RECORDS
2004	11,488,000
2005	104,321,000
2006	124,235,000
2007	171,077,984
2008	360,834,871
2009	143,643,022
2010	3,878,370
2011	174,522,698
2012	44,800,841

Attack Targeting and Difficulty

The range of possible breach scenarios is vast, and protecting against them all is challenging. It requires identifying the right resources and where to allocate them to counter the pressures applied by our adversaries. Studying target selection and attack difficulty is a useful way of understanding which points they pick and how hard they press.

Figure 38: Attack targeting

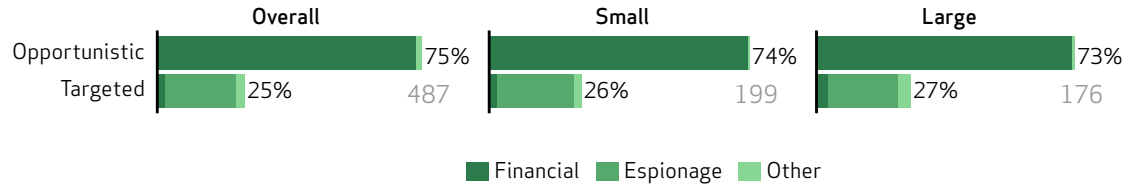

Though it's not found in the VERIS schema, we make an attempt to determine whether an attack leading to a breach was targeted or opportunistic in nature. The differences between these are vast, but let's begin with a description:

- **Opportunistic attacks:** The victim isn't specifically chosen as a target; they were identified and attacked because they exhibited a weakness the attacker knew how to exploit.
- **Targeted attacks:** The victim is specifically chosen as a target; the attacker(s) then determines what weaknesses exist within the target that can be exploited.

From an overall standpoint, the ratio (which is still imbalanced) tilted a little in the direction of targeted attacks in 2012 (a <10% change). But it's still a big world of opportunity out there. Unlike 2011, however, where half of attacks against larger organizations were targeted, these findings look almost identical across organizational size. To what can we attribute this difference? You guessed it—the number and size-agnostic nature of targeted espionage breaches in this dataset evened things out. It pays to follow along.

In one of our previous reports, we pointed out that some organizations will be a target *regardless* of what they do, but most become a target *because* of what they do (or don't do). If your organization is indeed a target of choice, understand as much as you can about what your opponent is likely to do and how far they are willing to go. The rest of us should work to eliminate sloppy configurations, needless services, and exposed vulnerabilities that inevitably bring unwanted attention.

In addition to target selection, examining attack difficulty is also a fruitful endeavor. Though not, perhaps, in the way one might think. Some interpret attack difficulty as synonymous with the skill of the attacker, and while there's some truth to that, it almost certainly reveals much more about the skill and readiness of the defender.

SOME INTERPRET ATTACK DIFFICULTY AS SYNONYMOUS WITH THE SKILL OF THE ATTACKER, AND WHILE THERE'S SOME TRUTH TO THAT, IT ALMOST CERTAINLY REVEALS MUCH MORE ABOUT THE SKILL AND READINESS OF THE DEFENDER.

It must be stated up front that some degree of subjectivity comes into play when rating the relative difficulty of the attacks (and this is why it's not part of the VERIS schema). We provide separate ratings for the method of initial compromise (how the attacker gained access) and for any subsequent actions done after that, including compromise and exfiltration of data. The rating scale is as follows:

- **Very low:** no special skills or resources required. The average user could have done it.
- **Low:** basic methods, no customization, and/or low resources required. Automated tools and scripts.
- **Moderate:** skilled techniques, some customization, and/or significant resources required.
- **High:** advanced skills, significant customizations, and/or extensive resources required.

Figure 39: Difficulty of initial compromise

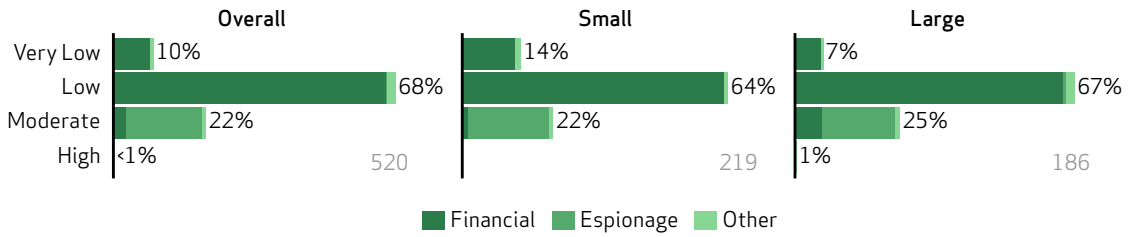

WHEN YOU CONSIDER THE METHODS USED BY ATTACKERS TO GAIN A FOOTHOLD IN ORGANIZATIONS—BRUTE FORCE, STOLEN CREDS, PHISHING, TAMPERING—IT'S REALLY NOT ALL THAT SURPRISING THAT NONE RECEIVE THE HIGHLY DIFFICULT RATING. WOULD YOU FIRE A GUIDED MISSILE AT AN UNLOCKED SCREEN DOOR?

In the overall results from Figure 39, three-quarters of breaches are of low or very low difficulty for initial compromise, and the rest land in the moderate category. This is in keeping with the findings from previous years and those discussed in this report as well. When you consider the methods used by attackers to gain a foothold in organizations—brute force, stolen creds, phishing, tampering—it's really not all that surprising that none receive the highly difficult rating. Would you fire a guided missile at an unlocked screen door?

The same basic trend is visible when one examines the findings based on organizational size. Notice though, that financially motivated attacks fall in the low and very low categories while most moderate attacks tie to espionage. While phishing, the favored method of initial compromise in espionage campaigns, may not seem overly sophisticated, the malware it employs can be quite advanced.

If one were to compare the difficulty ratings in Figure 40 to those in our last report, the difference would be obvious. Only four percent of subsequent actions received the coveted highly difficult rating in 2011, while one in five received that accolade in 2012. At the risk of sounding like a broken record, you can blame this on the larger quantity of espionage cases. Determined threat actors will leverage formidable skills and resources to entrench themselves in the victim's environment and remain hidden until their mission is accomplished.

Figure 40: Difficulty of subsequent actions

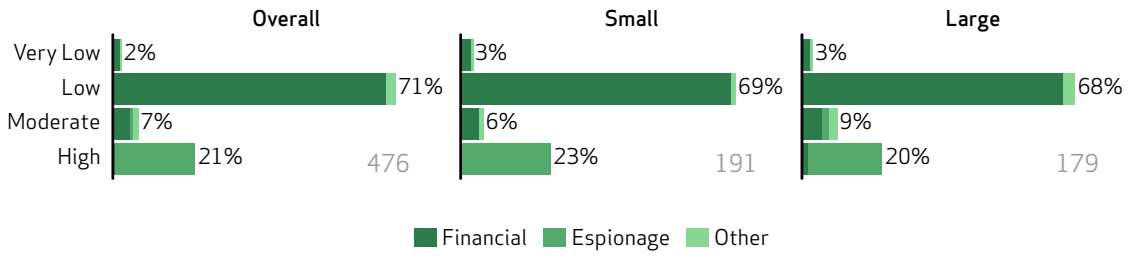

An interesting high-level observation from Figure 40 is the "difficulty gulf" between financial and espionage attacks. This may have something to do with the goals held by the two camps of actors. Those seeking easy money are willing to cut ties and head to the next victim of opportunity when things dry up or turn south. Meanwhile, those focused on longer-term goals set by their employer (or government) have more skin in the game and will protect their investments diligently.

Breach Timeline

As conveyed in previous DBIRs, understanding the time frame of an incident can greatly increase our ability to produce accurate threat models and capability assessments. In the past, we included all types of breaches in our timeline analysis, but we filtered results this year to include only those involving network intrusions. This has the effect of removing things like ATM skimming that are less interesting for timeline analysis (criminals can slap on a faceplate in nothing flat) and tend to throw a wrench into the statistics. Apologies for droning on with a list of caveats, but another thing to note is that we removed any "unknowns" from this analysis. Sometimes it cannot be forensically determined when certain events occurred and sometimes the information is not tracked or provided to us by contributors. Having said all that, let's see what Figure 41 has for us in terms of known timespan details about intrusions leading to data compromise in 2012.

Initial Compromise

This initial phase depicts the time from the first malicious action taken against the victim until the point at which an information asset is negatively affected. In network intrusions, this represents how long it takes the attacker to get his foot in the door.

Current results show a substantial shift toward hours, compared to 2011 where 85% of initial compromises occurred within minutes or less. This is largely attributable to narrowing the scope this year to network intrusions. The lower proportion of smash-and-grab POS hacks in this dataset is another factor, since they typically occur very quickly. Along with SQL injections, they represent a goodly portion of incidents falling into the Seconds and Minutes buckets in Figure 41. After all, how long does it take to crack a default POS vendor password or make a database cough up some records with an invalid SQL query? Interestingly, when comparing organization size, the larger ones were compromised somewhat quicker than were their smaller counterparts, but not to the point that we'd call it a trend or put too much meaning into it.

With regard to motives, most breaches credited to espionage are either not known (and not shown) or fall within the Hours bucket. Due to a lack of logging and a variety of factors, the exact timeframes on such attacks can be a bit fuzzy. We simply may not know the exact amount of time elapsed between the sending of a phishing e-mail and a user clicking on said phish to infect a system. Financially motivated breaches show a more even distribution, likely due to the wider array of actors and actions involved in such crimes.

Figure 41: Timespan of events

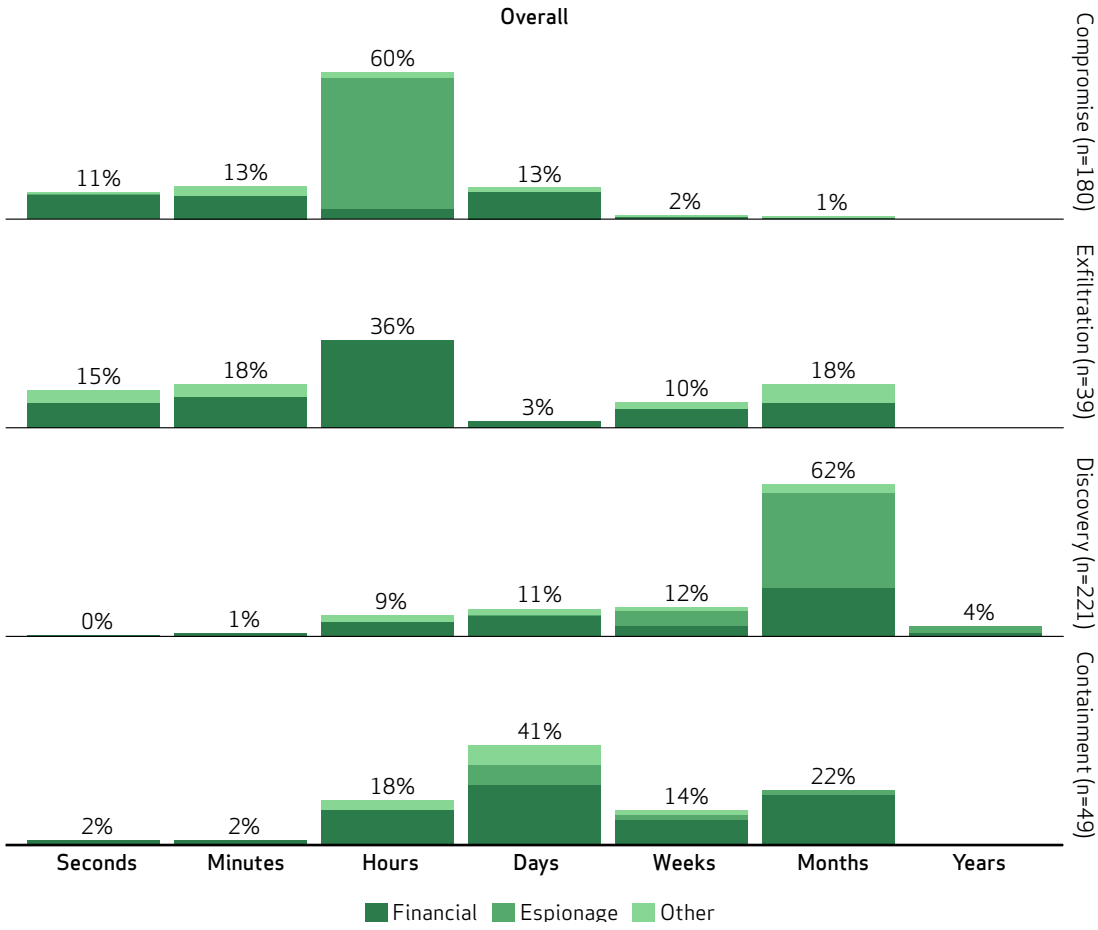

Initial Compromise to Data Exfiltration

This second phase covers the time period from initial compromise to the point when non-public information is first removed from the victim's environment. The number of these having measurable and known timeframes is much lower, especially among espionage victims. Overall, though, times are longer here than in the time-to-compromise phase. This represents the time necessary to explore the network, locate relevant systems, exploit those systems, and then collect and exfiltrate the data. This is akin to physical burglaries where it naturally takes longer to search a house to locate and remove valuables than it does to kick in the door. The shorter timespans in this phase usually represent breaches where data was stored on the point of compromise. POS hacks and SQL injections are good examples of that.

WHILE IT MIGHT BE DIFFICULT TO DETECT, POSITIVELY IDENTIFY, AND RESPOND TO AN INTRUSION WITHIN SECONDS OR MINUTES, OUR ABILITY TO DO SO SHOULD OSTENSIBLY INCREASE THE LONGER THEY POKE AROUND OUR INTERNAL NETWORKS. BUT UNFORTUNATELY, WE'RE NOT REALLY SEEING THAT IMPROVEMENT.

We continue to view this phase in particular as a giant opportunity for improvement in our industry. While it might be difficult to detect, positively identify, and respond to an intrusion within seconds or minutes, our ability to do so should ostensibly increase the longer they poke around our internal networks. But unfortunately, we're not really seeing that improvement.

Initial Compromise to Discovery

This important phase describes the time from initial compromise to when the victim first learns of the incident. Let's just cut to the chase and say the song remains the same here, and while that's a great Led Zeppelin song/album/film, it's not so great when applied to incident discovery. "Same" is not what we want to see here. Unfortunately, the biggest change is that the sizeable proportion representing Months is a stacked bar chart rather than the bubbles of yesteryear. Either way you look at it, the majority of breaches take months or more to discover.

Figure 42: Percent of breaches that remain undiscovered for months or more

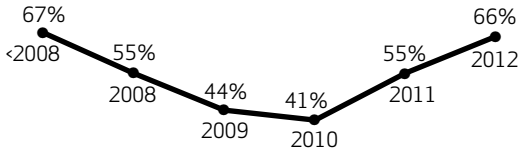

Speaking of rates of change—does Figure 42 give anyone flashbacks to derivatives from high school calculus? While that might not conjure up good memories for some of you, they must be better than the present depressing reality shown here. We've lost any sign of forward progress and are back to where we were when we started this study.

At least the large espionage-shaded region in the Months column in Figure 41 allows for casting off some of the blame for this. That pits the virtually unlimited resources of a nation against the very finite resources of a single company. Nobody can reasonably be expected to withstand THAT, right? Thank goodness for that "get out of jail free" card. For a moment there it was looking like something would actually need to be done about this.

But in all seriousness—something has to be done. If not *the* most, this must be *one of* the most important challenges to the security industry. Prevention is crucial, and we can't lose sight of that goal. But we must accept the fact that no barrier is impenetrable, and detection/response represents an extremely critical line of defense. Let's stop treating it like a backup plan if things go wrong, and start making it a core part of THE plan. With that, we'll draw the curtain of discretion on this act and hope 2013 turns this into a third-degree polynomial.

IF NOT THE MOST, THIS MUST BE ONE OF THE MOST IMPORTANT CHALLENGES TO THE SECURITY INDUSTRY. PREVENTION IS CRUCIAL, AND WE CAN'T LOSE SIGHT OF THAT GOAL. BUT WE MUST ACCEPT THE FACT THAT NO BARRIER IS IMPENETRABLE, AND DETECTION/RESPONSE REPRESENTS AN EXTREMELY CRITICAL LINE OF DEFENSE. LET'S STOP TREATING IT LIKE A BACKUP PLAN IF THINGS GO WRONG, AND START MAKING IT A CORE PART OF THE PLAN.

Discovery to Containment

This last phase measures time between the discovery of a breach to when it is successfully contained (when "the bleeding has stopped"). Depending on the circumstances, this can be challenging to measure. It's more likely in a forensics capacity where there is high touch with the victim and containment is a central goal. But CSIRTs and other authorities are less likely to have the longer-term and containment-focused relationship with the victim in order to obtain this data.

Based on the data that we do possess, there is no discernible difference from our last report. The majority of breaches are contained within days, but almost a quarter

of them took months. Due to the sparseness of the results, it's hard to glean much from a motive perspective. Containment timeframes run the gamut for financially motivated crimes, and show mostly Days for the small proportion of espionage cases with known timeframes. The latter is a rather counterintuitive finding due to a few special cases and not representative of the norm.

Discovery Methods

Having examined the (really long) time it takes to discover a breach in the previous section, we now turn our attention to how those discoveries are made. Similar to the Timeline section, we've filtered the results to focus on network intrusions. Approximately 70% of breaches were discovered by external parties who then notified the victim. This is admittedly better than the 92% observed in our last report, but well within the range of prior years. In fact, the imaginary regression line in Figure 43 looks pretty flat to our eyes, and that suggests internal detection capability is lacking, not widespread, or both.

Figure 43: Percent of breaches discovered external to victim

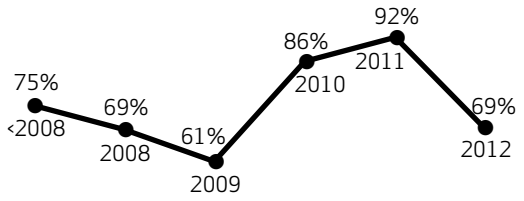

Moving away from the high-level comparison of External to Internal discovery, Figure 44 offers a more itemized view of discovery methods in 2012. In a real-life "Cinderella story, out of nowhere," breaches reported to the victim by unrelated external parties had a huge turnaround year, capturing their first major tour win. (Didn't catch the Caddyshack reference? Well, you should. It's funny; trust us.) "Unrelated" in this sense refers to external parties with whom the victim has no business relationship specific to detection services and are also not law enforcement, card brands, etc., whose official mission is to notify victims. Common examples are ISPs,

APPROXIMATELY 70% OF BREACHES WERE DISCOVERED BY EXTERNAL PARTIES WHO THEN NOTIFIED THE VICTIM. THIS IS ADMITTEDLY BETTER THAN THE 92% OBSERVED IN OUR LAST REPORT, BUT WELL WITHIN THE RANGE OF PRIOR YEARS. IN FACT, THE IMAGINARY REGRESSION LINE IN FIGURE 43 LOOKS PRETTY FLAT TO OUR EYES, AND THAT SUGGESTS INTERNAL DETECTION CAPABILITY IS LACKING, NOT WIDESPREAD, OR BOTH.

ISACs, and intelligence organizations that track threat actors and, when appropriate, inform potential victims of suspicious activity. As for the suspicious activity detected, that varies from situation to situation, but often involves communication to and from malicious IPs and domains associated with known threat groups. Due to the effectiveness of monitoring IOCs for state-affiliated groups, this method accounts for the discovery of many of the espionage-related breaches in this dataset.

Far from a Cinderella story, it's not surprising that third-party fraud detection (e.g., Common Point of Purchase, or CPP) tops Figure 44 as the leading way to discover financially motivated attacks (and would be even higher if ATM and cashier skimming were included). This method is especially prevalent for smaller retail or food services establishments, which have fewer human and technical resources to deter and detect attacks. The advantage of CPP lies in its ability to correlate suspicious patterns across many disparate organizations, thereby providing a more comprehensive vantage point than any single organization could achieve. The main disadvantage is that it only kicks in after fraud has begun, using stolen payment card data.

Once again, end users represent the most effective means of detecting a breach internally (and it would be even higher if ATM skimmers spotted by employees were included). Typically, this involves a regular employee who, in the course of their daily responsibilities, notices something strange (e.g., slower system performance or an e-mail that looks suspicious) and alerts IT or management. Let that fact and all its ramifications sink in.

Figure 44: Discovery methods

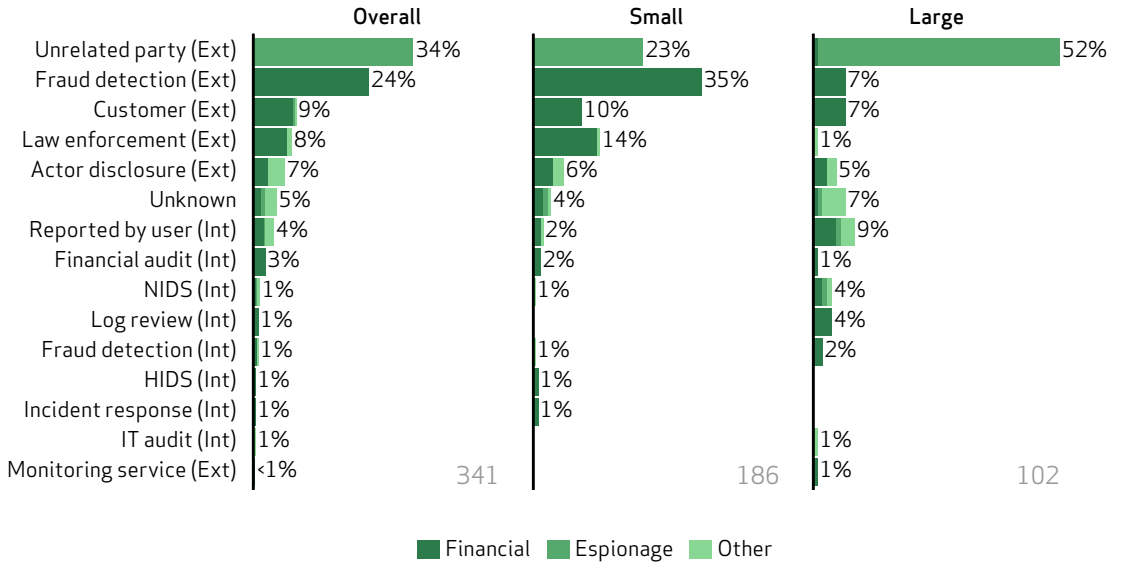

USING IOC'S AND NETFLOW TO GUIDE INVESTIGATION

In any investigation, one must rely on evidence as a guide down the appropriate path. Some of the most important signposts on this path are indicators of compromise (IOCs): identifiable events and artifacts that suggest a security incident occurred. Consistently collecting and maintaining the right data sources provides an organization with a resource from which to mine for IOCs, and a basic foundation for a stronger investigation. One should ask such questions as: *What level of logging does my host-based security solution have? What data am I logging on the network? Am I tracking DNS requests?* Retention time is also a crucial factor since most data breaches aren't discovered for weeks or months.

Although the victim's own data sources are an important element of an investigation, the data provided by external parties can also be of great value. For instance, the RISK Team aggregates IOCs from publicly-available feeds, information sharing groups, law enforcement relationships, and our own investigations. Matching this IOC library with victim-side evidence kick starts an investigation and allows for much quicker and more effective progress.

Another useful investigative resource that complements IOCs nicely is netflow data. Netflow consists of basic routing information for sequences of packets (e.g., source and destination IP, ports, etc.). If IOCs describe "what to look for," netflow often provides a broad lens of "where to look," and aids investigators in determining the nature and extent of a particular case or larger campaign.

Thanks to our role as an ISP, we can—with client permission—bring netflow into the investigative process to augment client evidence and even partially compensate for missing event logs. On-site investigators, while following the chain of evidence, also pass possible IOCs and data to intelligence analysts on the team. This second group works in tandem with investigators and compares this data to previous incidents, known IOCs, open-source data, etc. In one recent case, this process yielded links from four separate cases across three continents, identified more potential victims, and assisted the law enforcement entities involved. As the field continues to improve cooperation and intelligence sharing, look for IOCs to be one of the primary currencies of information.

We suspect organizations spend a lot more time and money on things that fall below the one percent mark in Figure 44, and do very little to hone and support the detection capability of their human resources. Maybe larger organizations—which discovered about a quarter of all breaches in this manner!—realize this and actually train employees to keep their eyes open and empower them to *act* on what they observe. We can't prove that connection exists from this data, but if you're looking to support a business case to management for IR training for end users, Figure 44 might help.

And for the closing data visualization of the 2013 DBIR, we offer this bit of lagniappe. Figure 45 plots discovery methods against timeframes observed across 2012 breaches, and yields some intriguing results. There's plenty of food for thought and discussion here, but we're going to bow out here and let that occur outside the pages of the DBIR.

Figure 45: Breach count by discovery method and time to discovery

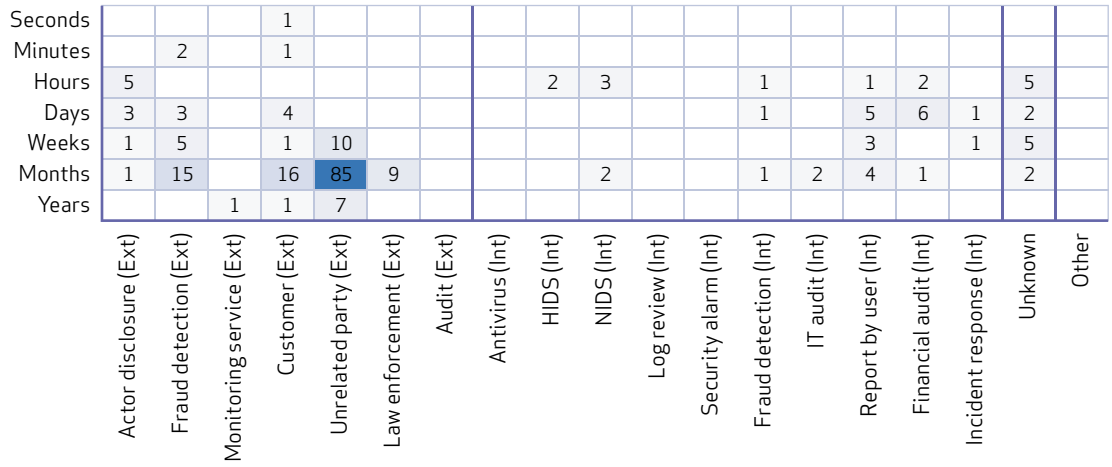

Questions? Comments? Brilliant ideas?
We want to hear 'em. Drop us a line at dbir@verizon.com, find us on LinkedIn and Facebook, or post to Twitter with the hashtag #dbir.

CONCLUSIONS AND RECOMMENDATIONS

We have lamented in years past the difficulty in creating a list of recommendations and providing our readership prescriptive marching orders. The most common threat actions have realized some shifts over the years, but we have failed to see any cutting-edge methods introduced. Better reconnaissance to craft better spear-phishing campaigns? Sure. Automation and scalability improvements incorporated by financially motivated groups? Absolutely. But at the end of the day, phishing (even targeted phishing campaigns) and attacks against weak passwords are not new, and there are no new whiz-bang controls to solve the world's problems. So we will continue to provide solid recommendations based on the story the data tells us.

To that end, we worked with the recently formed Consortium for Cybersecurity Action (CCA) and mapped the most common threat action varieties to their Critical Security Controls for Effective Cyber Defense[22]. This control set is widely vetted and adopted, and we appreciate the opportunity to collaborate with the CCA to provide improved recommendations in the DBIR. All of the high-level control categories are broken down into sub-controls as well as implementation and testing guidance. If you haven't already, the first recommendation of this section is to familiarize yourself with the content and structure of the 20 Critical Security Controls (CSC)[23].

Even with a well-regarded control set established by a consortium of security professionals to back up our recommendations, there is no one-size-fits-all solution. The feasible level of implementation across all CSC controls (or any set of controls) will differ among organizations depending on size, budget, business need, etc. Plus, the order of effective implementation will also differ, based on the need to address the most critical threats first. And if we have demonstrated anything in this report, we hope it is evident that list of threats can differ dramatically from one organization to another.

22 The Critical Security Controls were originally developed under the leadership of the Center for Strategic and International Studies (CSIS) and The SANS Institute, under the name of the Consensus Audit Guidelines (CAG).
23 http://www.sans.org/critical-security-controls/

Below are the 20 Critical Security Controls along with some examples of their areas of focus.

1. **Inventory of Authorized and Unauthorized Devices:** Asset tracking
2. **Inventory of Authorized and Unauthorized Software:** Software inventories, monitoring and notifications regarding unapproved software, application whitelisting, and software identification tagging
3. **Secure Configurations for Hardware and Software on Laptops, Workstations, and Servers:** Configuration monitoring and management, standard system images, software currency, and file integrity checks
4. **Continuous Vulnerability Assessment and Remediation:** Automated vulnerability scanning, port checking, and patch management solutions
5. **Malware Defenses:** Anti-virus tools, disabling auto-run, traffic analysis, secure e-mail usage, and sandboxing
6. **Application Software Security:** Application testing and code review
7. **Wireless Device Control:** Wireless device identifiers, network access control
8. **Data Recovery Capability:** No sub-controls were primary mitigators of top threat actions
9. **Security Skills Assessment and Appropriate Training to Fill Gaps:** Security awareness training, security policies, and awareness testing
10. **Secure Configurations for Network Devices such as Firewalls, Routers, and Switches:** Strong authentication for network infrastructure
11. **Limitations and Control of Network Ports, Protocols, and Services:** Conservative device configuration, default-deny stance
12. **Controlled Use of Administrative Privileges:** Identification and monitoring of administrative accounts, restriction of access to administrative accounts, and securing administrative accounts with strong authentication
13. **Boundary Defense:** Ingress and egress filtering based on blacklists, and default deny principle, DMZ traffic monitoring, IDS technologies, application proxies
14. **Maintenance, Monitoring, and Analysis of Security Audit Logs:** Audit log settings, storage, retention, and review
15. **Controlled Access Based on the Need to Know:** Network segmentation, logical access control
16. **Account Monitoring and Control:** Account auditing, password parameters, account lockout settings, monitoring attempts to access disabled accounts and atypical account usage
17. **Data Loss Prevention:** Mobile hard drive encryption, DLP software
18. **Incident Response and Management:** No sub-controls were primary mitigators of top threat actions
19. **Secure Network Engineering:** Network segmentation, establishment of security zones
20. **Penetration Tests and Red Team Exercises:** Inclusion of social attacks in sanctioned penetration testing

In an effort to show the mappings in a concise manner, and save innocent trees (along with the .pdf sub-species of tree) in the process, we have only included a mapping of the high-level controls to each of the most prevalent threat actions from this year's overall dataset. The full mappings, including detailed sub-controls, are available on the CSC site[24]. Please use this page for the descriptions and intersections that are not listed in their entirety here.

24 http://www.sans.org/critical-security-controls/

Figure 46: CCA's Critical Security Controls mapped to common VERIS threat actions

Top VERIS Threat Actions	20 Critical Security Controls																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Tampering	•		•				•		•											
Spyware		•	•	•	•							•								
Backdoor**		•	•	•	•	•				•	•	•	•							
Export data		•	•	•	•					•		•	•				•		•	
Use of stolen creds							•		•	•		•								
Capture stored data		•	•	•	•							•			•		•		•	
Phishing		•	•	•	•				•			•	•							•
C2		•	•	•	•	•				•	•	•	•							
Downloader		•	•	•	•							•	•							
Brute force				•		•	•		•	•		•	•	•	•	•				

** BKDOOR includes the Malware threat actions of backdoor and command and control, along with the Hacking action that represents the use of backdoor and command and control channels.

With seven of the top 10 threat actions belonging to the Malware threat action category, it's not surprising that there's a great deal of overlap in the Critical Security Controls across Figure 46. It should be noted again that this is showing only the controls that were primary recommendations. Practices like security awareness campaigns (CSC9) and secure network architecture (CSC19) will likely result in some risk reduction across the board, but they are only highlighted where they will act as primary security controls for particular threat actions. Incident Response Capabilities (CSC18) are extremely important in identifying and containing issues and ultimately lessening their impact. Data Recovery Capabilities (CSC8) are essential to protect business functionalities and to recover from availability issues. Neither of these—in the restricted scope of the top 10 threat actions leading to data breaches—provided primary risk mitigation. In general, well-designed controls do not represent a one-to-one defense against individual types of attack, but are instead measures that provide value against multiple classes of attack.

Most organizations should implement all 20 of the Critical Security Controls to some level. In this report and others we have produced, you can find lists of the top threat actions for various industries and sizes of organizations. And because the full threat-to-control mappings are publicly available, anyone has the ability to produce their own set of "top-of-the-Top-20" controls to evaluate and use as they see fit.

Recommendations for Mitigating Highly Targeted Attacks

As you may be aware, we support evidence-based risk management. We believe that using the data at your disposal is a great way to shape security programs that fit your organization. In addition to the SANS Critical Security Control mapping project, we wanted to focus on strategies for managing risks around determined adversaries and targeted attacks. Conventional controls such as updated patches and antivirus signatures certainly have their place, but focusing solely on controls such as these will not meet with success. So where should you focus? That is a complex question, but we think our findings, combined with others' research, do support a few recommendations.

First, focus on the kill-chain approach. This approach is effective against highly targeted attacks largely because these attacks typically have long event chains that involve several assets. Hutchins and others from Lockheed Martin first published about the methodology in 2010,[25] but we don't have room here for a full description. Essentially, it recognizes that "the adversary must progress successfully through each stage of the chain before it can achieve its desired objective; just one mitigation disrupts the chain and the adversary." By analyzing the entire chain, the approach essentially shifts security from pure reaction to proactive anticipation. It also raises costs for highly targeted attackers. Public discussions of success are rare, but they exist.[26] And we also hear from others in private that they're using the approach successfully. We would be remiss if we didn't bring up the natural connection between the VERIS framework and kill chain analysis. VERIS allows for dissection of an incident into a series of independent events, allowing you to view the entire chain and determine the most efficient and feasible stage at which to focus additional security controls. All this to say, don't just focus on preventative controls. Detection is equally as important, and response could be even more important in today's threat landscape.

Secondly, actions that evade signature detection require a more preventive approach to protecting assets; e.g., Microsoft's EMET (Enhanced Mitigation Experience Toolkit).[27] As history has shown, focusing on finding specific vulnerabilities and blocking specific exploits is a losing battle. EMET functions by blocking entire classes of exploits instead of only the specifics. In doing this, EMET helps shift security from being reactive to proactive and raises costs for attackers. There's a growing body of evidence indicating the approach is effective: we're seeing more reports about new zero-day attacks that simply do not work when EMET is in place.[28] So assets are protected before a patch or signature is released—they're even protected before the zero-day is discovered. In contrast to kill chain analysis, the cost of deploying EMET arguably favors small organizations. It's free, but larger organizations will likely have some configuration management cost. Also, EMET is specifically referenced in the Critical Security Controls (sub-control 5.7) and is part of the recommended control set for the Malware threat actions called out in the preceding section.

Finally, targeted attacks frequently rely on social methods to compromise people, not just computers (see our section on compromised assets). We've seen examples of great corporate security programs that are bypassed as a result of this. Likewise, we've also seen actors go after high-value targets in their personal lives, using social tactics like phishing, doxing, and watering hole attacks to compromise personal e-mail accounts and computing devices. In the grand scheme of espionage, targeting specific key personnel isn't anything new (think about the Cold War era). In the past, fewer organizations may have thought about extending corporate security into the living rooms of their CEOs. More organizations are now at least considering that option, and depending on your posture, it may be something you want to consider as well.

25 https://www.vita.virginia.gov/uploadedfiles/VITA_Main_Public/Security/Meetings/ISOAG/2012/Sept_ISOAG_NetworkDefense.pdf
26 http://www.darkreading.com/authentication/167901072/security/attacks-breaches/240148399/how-lockheed-martin-s-kill-chain-stopped-securid-attack.html
27 http://www.microsoft.com/en-us/download/details.aspx?id=30424
28 https://isc.sans.edu/diary/14797

APPENDIX A: A PERSPECTIVE FROM THE NEW EUROPEAN CYBERCRIME CENTER (EC3)

We spend most of our time in the DBIR discussing who the bad guys are, what they tend to do, how victims respond, etc. We spend comparatively little time discussing who the good guys are, what they're doing, and how they're responding to the changing global nature of data breaches and related threats. Thus, we like to include perspectives from some of the good guys willing to share that perspective with us. This year, that perspective comes from Troels Oerting, Assistant Director at the newly-established European Cybercrime Center (EC3).

The European Union is comprised of 500 million citizens in 27 (soon to be 28) member states. We have 23 different languages, 27 different legal systems, and 2.9 million law enforcement officers. We have no internal borders and a single market. And we have a new borderless crime with no link between crime and perpetrator.

We are 72% wired in the EU and heavily dependent on our digital infrastructure. Our way out of these times of economic austerity is based on our ability to invent, brand, innovate, produce and deliver in a global economy with increasing competition from emerging economic areas like the BRIC states.

Our citizens shop on the Internet. They operate digital signatures in exchanging sensitive information with their banks, doctors, municipalities, libraries, governments, and this development will continue. We will see more opportunities and possibilities—more transparency and democracy.

In the future we will likely ***always*** be online, even without direct access to a PC.

Even today when we, as citizens, societies, governments, retail, academia and industry go online, multitudes of criminals lie in wait. Adversaries try to steal our identities, our information, and our money. Cyberspace is often misused to facilitate various types of crime such as fraud, sexual abuse of children, sale of illegal commodities and drugs—the list is endless.

But a free and transparent Internet is of no value if it is not safe. And a safe and secure Internet is not attractive if it does not protect freedom of speech. We have to balance these two important principles.

Based on the increasing threat in Cyberspace, the EU Commission, with the support of the EU Parliament and the EU Council, decided to establish a new European Cybercrime Centre (EC3). The Centre opened its doors on January 1st 2013 in Europol's headquarters in The Hague.

EC3 has two years to reach "cruising speed" and deliver in all fields. In the beginning we will focus on the following areas:

- Develop a fusion center to create an overview on cybercrime in the EU and coordinate Member State (MS) investigations.
- Support operational cases in MS in three prime areas: Intrusions, Fraud and Child Sexual Abuse online.

- Establish a Cyber Lab to assist in complicated cases.
- Establish a Cyber Innovation Room including a Large File Exchange facility and a malware sandbox to support Joint Investigation Teams in MS including private partners.
- Develop an outreach strategy for including key private and public partners more directly in EC3's work.
- Initiate Research and Development to develop forensic and investigative tools to help MS investigators and use financial support from EU.
- Enhance our ability to draft threat assessments, scan notices and reports on emerging trends and mitigation suggestions.
- Initiate capacity building in MS and abroad.
- Keep an eye on attacks on EU critical infrastructure.

The above tasks cannot be delivered by EC3 alone. Our approach is very inclusive and our program board covers representatives from ENISA, EEAS, EU Commission, EU Cybercrime Task Force, CEPOL, CERT-EU, and EMPACT. It will soon open for advisory positions to public and private partners in academia, industry and companies.

Our work is closely coordinated with Interpol and a representative has a seat in the program board. Eurojust and Interpol have posted liaison officers in EC3; more will follow.

Our ambition is to reach out to all stakeholders with respect for their tasks and responsibilities and to focus on the criminal—not the crime. We need to make it very unattractive to be a cybercriminal, and today it is almost a free ride.

Adding value to the front line work within Member States is key. The future will show if we can deliver as expected.

I am not in doubt.

Troels Oerting
Assistant Director,
European Cybercrime Centre (EC3)
Europol.

APPENDIX B: FULL LIST OF 2013 DBIR CONTRIBUTORS

1. Australian Federal Police (AFP)
   www.afp.gov.au/policing/cybercrime
2. CERT Insider Threat Center at the Carnegie Mellon University Software Engineering Institute (CERT)
   www.cert.org/insider_threat/index.html
3. Consortium for Cybersecurity Action (CSIS control mapping)
   www.sans.org/critical-security-controls/
4. Danish Ministry of Defence, Center for Cybersecurity
   www.fmn.dk/Eng/Pages/Frontpage.aspx
5. Danish National Police, NITES (National IT Investigation Section)
   www.politi.dk/en/servicemenu/home/
6. Deloitte
   www.deloitte.com
7. Dutch Police: National High Tech Crime Unit (NHTCU)
   www.politie.nl
8. Electricity Sector Information Sharing and Analysis Center (ES-ISAC)
   www.esisac.com/SitePages/Home.aspx
9. European Cyber Crime Center (EC3)
   www.europol.europa.eu/ec3
10. G-C Partners, LLC
    www.g-cpartners.com/
11. Guardia Civil (Cybercrime Central Unit)
    www.gdt.guardiacivil.es
12. Industrial Control Systems Cyber Emergency Response Team (ICS-CERT)
    ics-cert.us-cert.gov/ics-cert/
13. Irish Reporting and Information Security Service (IRISS-CERT)
    www.iriss.ie/iriss/RFC_2350.htm
14. Malaysia Computer Emergency Response Team (MyCERT), CyberSecurity Malaysia
    www.mycert.org.my/en/
15. National Cybersecurity and Communications Integration Center (NCCIC)
    www.us-cert.gov/nccic/
16. ThreatSim
    threatsim.com
17. U.S. Computer Emergency Readiness Team (US-CERT)
    www.us-cert.gov/
18. U.S. Secret Service
    www.secretservice.gov
19. Verizon
    www.verizonenterprise.com

For additional information on the DBIR and access to related content, please visit
www.verizonenterprise.com/DBIR/2013

Questions? Comments? Brilliant ideas?
We want to hear 'em. Drop us a line at dbir@verizon.com, find us on LinkedIn and Facebook, or post to Twitter with the hashtag #dbir.

verizonenterprise.com