2012 DATA BREACH INVESTIGATIONS REPORT

A study conducted by the Verizon RISK Team with cooperation from the Australian Federal Police, Dutch National High Tech Crime Unit, Irish Reporting & Information Security Service, Police Central e-Crime Unit, and United States Secret Service.

verizon

2012 DATA BREACH INVESTIGATIONS REPORT

A study conducted by the Verizon RISK Team with cooperation from the Australian Federal Police, Dutch National High Tech Crime Unit, Irish Reporting and Information Security Service, Police Central e-Crime Unit, and United States Secret Service.

2012 DATA BREACH INVESTIGATIONS REPORT

TABLE OF CONTENTS

Executive Summary 2
Methodology 5
- Classifying Incidents Using VERIS 6
- A Word on Sample Bias 8

Results and Analysis 9
- Demographics 10
- 2011 DBIR: Threat Event Overview 13
- Threat Agents 16
  - Breach Size by Threat Agents 18
  - External Agents (98% of breaches, 99+% of records) 19
  - Internal Agents (4% of breaches, <1% of records) 21
  - Partner Agents (<1% of breaches, <1% of records) 22
- Threat Actions 23
  - Malware (69% of breaches, 95% of records) 26
  - Hacking (81% of breaches, 99% of records) 30
  - Social (7% of breaches, 37% of records) 33
  - Misuse (5% of breaches, <1% of records) 35
  - Physical (10% of breaches, <1% of records) 36
  - Error (<1% of breaches, <1% of records) 37
  - Environmental (0% of breaches, 0% of records) 38
- Compromised Assets 38
- Compromised Data 41
- Attack Difficulty 45
- Attack Targeting 47
- Timespan of Events 48
- Breach Discovery Methods 51
- Anti-Forensics 55
- PCI DSS 56
- The Impact of Data Breaches 58

2012 DBIR: Conclusions and Recommendations 61
Appendix A: Examining relationships among threat actions 67
Appendix B: A USSS case study of large-scale "industrialized" cybercrime 72
About the 2012 DBIR Contributors 74
- Verizon RISK Team 74
- Australian Federal Police 74
- Dutch National High Tech Crime Unit 74
- Irish Reporting & Information Security Service 75
- Police Central e-Crime Unit 75
- United States Secret Service 76

For additional updates and commentary, please visit verizon.com/enterprise/securityblog

EXECUTIVE SUMMARY

2011 will almost certainly go down as a year of civil and cultural uprising. Citizens revolted, challenged, and even overthrew their governments in a domino effect that has since been coined the "Arab Spring," though it stretched beyond a single season. Those disgruntled by what they perceived as the wealth-mongering "1%" occupied Wall Street along with other cities and venues across the globe. There is no shortage of other examples.

This unrest that so typified 2011 was not, however, constrained to the physical world. The online world was rife with the clashing of ideals, taking the form of activism, protests, retaliation, and pranks. While these activities encompassed more than data breaches (e.g., DDoS attacks), the theft of corporate and personal information was certainly a core tactic. This re-imagined and re-invigorated specter of "hacktivism" rose to haunt organizations around the world. Many, troubled by the shadowy nature of its origins and proclivity to embarrass victims, found this trend more frightening than other threats, whether real or imagined. Doubly concerning for many organizations and executives was that target selection by these groups didn't follow the logical lines of who has money and/or valuable information. Enemies are even scarier when you can't predict their behavior.

This re-imagined and re-invigorated specter of "hacktivism" rose to haunt organizations around the world.

It wasn't all protest and lulz, however. Mainline cybercriminals continued to automate and streamline their method du jour of high-volume, low-risk attacks against weaker targets. Much less frequent, but arguably more damaging, were continued attacks targeting trade secrets, classified information, and other intellectual property. We certainly encountered many faces, varied tactics, and diverse motives in the past year, and in many ways, the 2012 Data Breach Investigations Report (DBIR) is a recounting of the many facets of corporate data theft.

855 incidents, 174 million compromised records.

This year our DBIR includes more incidents, derived from more contributors, and represents a broader and more diverse geographical scope. The number of compromised records across these incidents skyrocketed back up to 174 million after reaching an all-time low (or high, depending on your point of view) in last year's report of four million. In fact, 2011 boasts the second-highest data loss total since we started keeping track in 2004.

Once again, we are proud to announce that the United States Secret Service (USSS) and the Dutch National High Tech Crime Unit (NHTCU) have joined us for this year's report. We also welcome the Australian Federal Police (AFP), the Irish Reporting & Information Security Service (IRISSCERT), and the Police Central e-Crime Unit (PCeU) of the London Metropolitan Police. **These organizations have broadened the scope of the DBIR tremendously with regard to data breaches around the globe. We heartily thank them all for their spirit of cooperation, and sincerely hope this report serves to increase awareness of cybercrime, as well as our collective ability to fight it.**

It wasn't all protest and lulz, however. Mainline cybercriminals continued to automate and streamline their method du jour of high-volume, low-risk attacks against weaker targets.

With the addition of Verizon's 2011 caseload and data contributed from the organizations listed above, the DBIR series now spans eight years, well over 2000 breaches, and greater than one billion compromised records. It's been a fascinating and informative journey, and we are grateful that many of you have chosen to come along for the ride. As always, our goal is that the data and analysis presented in this report prove helpful to the planning and security efforts of our readers. We begin with a few highlights below.

WHO IS BEHIND DATA BREACHES?

98% stemmed from external agents (+6%)

4% implicated internal employees (-13%)

<1% committed by business partners (<>)

58% of all data theft tied to activist groups

No big surprise here; outsiders are still dominating the scene of corporate data theft. Organized criminals were up to their typical misdeeds and were behind the majority of breaches in 2011. Activist groups created their fair share of misery and mayhem last year as well—and they stole more data than any other group. Their entrance onto the stage also served to change the landscape somewhat with regard to the motivations behind breaches. While good old-fashioned greed and avarice were still the prime movers, ideological dissent and schadenfreude took a more prominent role across the caseload. As one might expect with such a rise in external attackers, the proportion of insider incidents declined yet again this year to a comparatively scant 4%.

HOW DO BREACHES OCCUR?

Incidents involving hacking and malware were both up considerably last year, with hacking linked to almost all compromised records. This makes sense, as these threat actions remain the favored tools of external agents, who, as described above, were behind most breaches. Many attacks continue to thwart or circumvent authentication by combining stolen or guessed credentials (to gain access) with backdoors (to retain access). Fewer ATM and gas pump skimming cases this year served to lower the ratio of physical attacks in this report. Given the drop in internal agents, the misuse category had no choice but to go down as well. Social tactics fell a little, but were responsible for a large amount of data loss.

81% utilized some form of hacking (+31%)

69% incorporated malware (+20%)

10% involved physical attacks (-19%)

7% employed social tactics (-4%)

5% resulted from privilege misuse (-12%)

WHAT COMMONALITIES EXIST?

79% of victims were targets of opportunity (-4%)

96% of attacks were not highly difficult (+4%)

94% of all data compromised involved servers (+18%)

85% of breaches took weeks or more to discover (+6%)

92% of incidents were discovered by a third party (+6%)

97% of breaches were avoidable through simple or intermediate controls (+1%)

96% of victims subject to PCI DSS had not achieved compliance (+7%)

Findings from the past year continue to show that target selection is based more on opportunity than on choice. Most victims fell prey because they were found to possess an (often easily) exploitable weakness rather than because they were pre-identified for attack.

Whether targeted or not, the great majority of victims succumbed to attacks that cannot be described as highly difficult. Those that were on the more sophisticated side usually exhibited this trait in later stages of the attack after initial access was gained.

Given this, it's not surprising that most breaches were avoidable (at least in hindsight) without difficult or expensive countermeasures. Low levels of PCI DSS adherence highlight a plethora of issues across the board for related organizations.

While at least some evidence of breaches often exists, victims don't usually discover their own incidents. Third parties usually clue them in, and, unfortunately, that typically happens weeks or months down the road.

Did you notice how most of these got worse in 2011?

Once again, this study reminds us that our profession has the necessary tools to get the job done. The challenge for the good guys lies in selecting the right tools for the job at hand and then not letting them get dull and rusty over time. Evidence shows when that happens, the bad guys are quick to take advantage of it.

As you'll soon see, we contrast findings for smaller and larger organizations throughout this report. You will get a sense for how very different (and in some cases how very similar) their problems tend to be. Because of this, it makes sense that the solutions to these problems are different as well. Thus, most of the recommendations given at the end of this report relate to larger organizations. It's not that we're ignoring the smaller guys—it's just that while modern cybercrime is a plague upon their house, the antidote is fairly simple and almost universal.

Larger organizations exhibit a more diverse set of issues that must be addressed through an equally diverse set of corrective actions. We hope the findings in this report help to prioritize those efforts, but truly tailoring a treatment strategy to your needs requires an informed and introspective assessment of your unique threat landscape.

WHERE SHOULD MITIGATION EFFORTS BE FOCUSED?

Smaller organizations

- ✔ Implement a firewall or ACL on remote access services
- ✔ Change default credentials of POS systems and other Internet-facing devices
- ✔ If a third party vendor is handling the two items above, make sure they've actually done them

Larger organizations

- ✔ Eliminate unnecessary data; keep tabs on what's left
- ✔ Ensure essential controls are met; regularly check that they remain so
- ✔ Monitor and mine event logs
- ✔ Evaluate your threat landscape to prioritize your treatment strategy
- ✔ Refer to the conclusion of this report for indicators and mitigators for the most common threats

Got a question or a comment about the DBIR?
Drop us a line at dbir@verizon.com, find us on Facebook, or post to Twitter with the hashtag **#dbir**.

METHODOLOGY

Based on the feedback we receive about this report, one of the things readers value most is the level of rigor and honesty we employ when collecting, analyzing, and presenting data. That's important to us, and we appreciate your appreciation. Putting this report together is, quite frankly, no walk in the park (855 incidents to examine isn't exactly a light load). If nobody knew or cared, we might be tempted to shave off some time and effort by cutting some corners, but the fact that you do know and do care helps keep us honest. And that's what this section is all about.

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.

Verizon 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 conducted by Verizon from 2004 to 2011. The 2011 caseload is the primary analytical focus of the report, but the entire range of data is referenced extensively throughout. Though the RISK team works a variety of engagements (over 250 last year), only those involving confirmed data compromise are represented in this report. There were 90 of these in 2011 that were completed within the timeframe of this report. To help ensure reliable and consistent input, we use the Verizon Enterprise Risk and Incident Sharing (VERIS) framework to record case data and other relevant details (fuller explanation of this to follow). VERIS data points are collected by analysts throughout the investigation lifecycle and completed after the case closes. Input is then reviewed and validated by other members of the RISK team. During the aggregation process, information regarding the identity of breach victims is removed from the repository of case data.

Data Collection Methodology for other contributors

The USSS, NHTCU, AFP, IRISSCERT, and PCeU differed in precisely how they collected data contributed for this report, but they shared the same basic approach. All leveraged VERIS as the common denominator but used varying mechanisms for data entry. For instance, agents of the USSS used a VERIS-based internal application to record pertinent case details. For the AFP, we interviewed lead agents on each case, recorded the required data points, and requested follow-up information as necessary. The particular mechanism of data collection is less important than understanding that all data is based on real incidents and, most importantly, real facts about those incidents. These organizations used investigative notes, reports provided by the victim or other forensic firms, and their own experience gained in handling the case. The collected data was purged of any information that might identify organizations or individuals involved and then provided to Verizon's RISK Team for aggregation and analysis.

From the numerous investigations worked by these organizations in 2011, in alignment with the focus of the DBIR, the scope was narrowed to only those involving confirmed organizational data breaches.[1] The scope was further narrowed to include only cases for which Verizon did not conduct the forensic investigation.[2] All in all, these agencies contributed a combined 765 breaches for this report. Some may raise an eyebrow at the fact that Verizon's caseload represents a relatively small proportion of the overall dataset discussed in this report, but we couldn't be happier with this outcome. We firmly believe that more information creates a more complete and accurate understanding of the problem we all collectively face. If that means our data takes a backseat in a Verizon-authored publication, so be it; we'll trade share of voice for shared data any day of the week.

1 "Organizational data breach" refers to incidents involving the compromise (unauthorized access, theft, disclosure, etc.) of non-public information while it was stored, processed, used, or transmitted by an organization.

2 We often work, in one manner or another, with these agencies during an investigation. To eliminate redundancy, Verizon-contributed data were used when both Verizon and another agency worked the same case.

While we're on that topic, if your organization investigates or handles data breaches and might be interested in contributing to future DBIRs, let us know. The DBIR family continues to grow, and we welcome new members.

A BRIEF PRIMER ON VERIS

VERIS is a framework 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 presented in this report. Because many readers asked about the methodology behind the DBIR and because we hope to facilitate more information sharing on security incidents, we have released VERIS for free public use. A brief overview of VERIS is available on our **website**[3] and the complete framework can be obtained from the **VERIS community wiki**.[4] Both are good companion references to this report for understanding terminology and context.

Classifying Incidents Using VERIS

The Incident Classification section of the VERIS Framework 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 A^4 Threat Model developed by Verizon's RISK team. In the A^4 model, a security incident is viewed as a series of events that adversely affects the information assets of an organization. Every event is comprised of the following elements (the four A's):

- **Agent:** Whose actions affected the asset
- **Action:** What actions affected the asset
- **Asset:** Which assets were affected
- **Attribute:** 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.

If we calculate all the combinations of the A^4 model's highest-level elements, (three Agents, seven Actions, five Assets, and three Attributes), 315[5] distinct threat events emerge. The grid in Figure 1 graphically represents these and designates a Threat Event Number (hereafter referenced by TE#) to each. TE1, for instance, coincides with External Malware that affects the Confidentiality of a Server. Note that not all 315 A^4 combinations are feasible. For instance, malware does not, insofar as we know, infect people…though it does make for intriguing sci-fi plots.

It is our position that the four A's represent the minimum information necessary to adequately describe any incident or threat scenario.

Turning the Incident Narrative into Metrics

As stated above, incidents often involve multiple threat events. Identifying which are in play, and using them to reconstruct the chain of events is how we model an incident to generate the statistics in this report. By way of example, we describe below a simplified hypothetical incident where a "spear phishing" attack is used to exfiltrate sensitive data and intellectual property (IP) from an organization.

The flowchart representing the incident includes four primary threat events and one conditional event.[6] A brief description of each event is given along with the corresponding TE#s and A^4 categories from the matrix exhibited earlier.

3 http://www.verizonbusiness.com/resources/whitepapers/wp_verizon-incident-sharing-metrics-framework_en_xg.pdf
4 https://verisframework.wiki.zoho.com/
5 Some will remember that this grid showed 630 intersections as presented in the 2011 DBIR. The difference is a result of the number of security attributes depicted. While we still recognize the six attributes of the "Parkerian Hexad," we (with input from others) have decided to use and present them in paired format (e.g., "confidentiality and possession losses"). Thus, the notions of confidentiality versus possession are preserved, but data analysis and visualization is simplified (a common request from VERIS users). More discussion around this change can be found on the Attributes section of the VERIS wiki.
6 See the Error section under Threat Actions for an explanation of conditional events.

Figure 1. VERIS A[4] Grid depicting the 315 high-level threat events

		Malware			Hacking			Social			Misuse			Physical			Error			Environmental		
		Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt
Servers	Confidentiality & Possession	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
	Integrity & Authenticity	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42
	Availability & Utility	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
Networks	Confidentiality & Possession	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79	80	81	82	83	84
	Integrity & Authenticity	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	100	101	102	103	104	105
	Availability & Utility	106	107	108	109	110	111	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126
User Devices	Confidentiality & Possession	127	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143	144	145	146	147
	Integrity & Authenticity	148	149	150	151	152	153	154	155	156	157	158	159	160	161	162	163	164	165	166	167	168
	Availability & Utility	169	170	171	172	173	174	175	176	177	178	179	180	181	182	183	184	185	186	187	188	189
Offline Data	Confidentiality & Possession	190	191	192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207	208	209	210
	Integrity & Authenticity	211	212	213	214	215	216	217	218	219	220	221	222	223	224	225	226	227	228	229	230	231
	Availability & Utility	232	233	234	235	236	237	238	239	240	241	242	243	244	245	246	247	248	249	250	251	252
People	Confidentiality & Possession	253	254	255	256	257	258	259	260	261	262	263	264	265	266	267	268	269	270	271	272	273
	Integrity & Authenticity	274	275	276	277	278	279	280	281	282	283	284	285	286	287	288	289	290	291	292	293	294
	Availability & Utility	295	296	297	298	299	300	301	302	303	304	305	306	307	308	309	310	311	312	313	314	315

Once the construction of the main event chain is complete, additional classification can add more specificity around the elements comprising each event (i.e., the particular type of External agent or exact Social tactics used, etc.). The incident is now "VERIS-ized" and useful metrics are available for reporting and further analysis.

The process described above has value beyond just describing the incident itself; it also helps identify what might have been done (or not done) to prevent it. The goal is straightforward: break the chain of events and you stop the incident from proceeding.

One final note before we conclude this sub-section. The process described above has value beyond just describing the incident itself; it also helps identify what might have been done (or not done) to prevent it. The goal is straightforward: break the chain of events and you stop the incident from proceeding. For instance, security awareness training and e-mail filtering could help keep E1 from occurring. If not, anti-virus and a least-privilege implementation on the laptop might prevent E2. Stopping progression between E2 and E3 may be accomplished through egress filtering or netflow analysis to detect and prevent backdoor access. Training and change control procedures could help avoid the administrator's misconfiguration described in the conditional event and preclude the compromise of intellectual property in E4. These are just a few examples of potential controls for each event, but the ability to visualize a layered approach to deterring, preventing, and detecting the incident should be apparent.

Figure 2. Sample VERIS incident scenario

External agent sends a phishing e-mail that successfully lures an executive to open the attachment.	Malware infects the exec's laptop, creating a backdoor.	External agent accesses the exec's laptop via the backdoor, viewing e-mail and other sensitive data.	System administrator misconfigures access controls when building a new file server.	External agent accesses a mapped file server from the exec's laptop and steals intellectual property.
TE#280 External Social People Integrity	**TE#148** External Malware User Devices Integrity	**TE#130** External Hacking User Devices Confidentiality	**TE# 38** Internal Error Servers Integrity	**TE#4** External Hacking Servers Confidentiality

A Word on Sample Bias

Allow us to reiterate: we make no claim that the findings of this report are representative of all data breaches in all organizations at all times. Even though the merged dataset (presumably) more closely reflect reality than they might in isolation, it is still a sample. 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. 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 2011. 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 knowledge grows along with what we are able to study and that grew more than ever in 2011. 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.

Got a question or a comment about the DBIR?
Drop us a line at dbir@verizon.com, find us on Facebook, or post to Twitter with the hashtag **#dbir**.

RESULTS AND ANALYSIS

The 2011 combined dataset represents the largest we have ever covered in any single year, spanning 855 incidents and over 174 million compromised records (the second-highest total, if you're keeping track). These next few paragraphs should help make some sense of it all.

In several places throughout the text, we present and discuss the entire range of data from 2004 to 2011. As you study these findings, keep in mind that the sample dataset is anything but static. The number, nature, and sources of cases 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. As in previous reports, the chosen approach is to present the combined dataset intact and highlight interesting differences (or similarities) within the text where appropriate. There are, however, certain data points that were only collected for Verizon cases; these are identified in the text and figures.

The figures in this report utilize a consistent format. **Values shown in dark gray pertain to breaches** while **values in red pertain to data records.** The "breach" is the incident under investigation in a case and "records" refer to the amount of data units (files, card numbers, etc.) compromised in the breach. In some figures, we do not provide a specific number of records, but use a red "#" to denote a high proportion of data loss. If one of these values represents a substantial change from prior years, this is marked with an orange "+" or "–" symbol (denoting an increase or decrease). Many figures and tables in this report add up to over 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.

Table 1. Key for translating percents to numbers for the 2012 DBIR dataset

855 BREACHES	
%	#
1%	9
5%	43
10%	86
25%	214
33%	282
50%	428

Values shown in dark gray pertain to breaches while **values in red pertain to data records.** The "breach" is the incident under investigation in a case and "records" refer to the amount of data units (files, card numbers, etc.) compromised in the breach. In some figures, we do not provide a specific number of records, but use a red "#" to denote a high proportion of data loss. If one of these values represents a substantial change from prior years, this is marked with an orange "+" or "–" symbol (denoting an increase or decrease).

Because the number of breaches in this report is so high, the use of percentages is a bit deceiving in some places (5 percent may not seem like much, but it represents over 40 incidents). Where appropriate, we show the raw number of breaches instead of or in addition to the percentages. A handy percent-to-number conversion table is shown in Table 1. 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.

Some constructive criticism we received about the 2011 report suggested the dataset was so rife with small breach victims that it didn't apply as strongly to larger organizations as it had in years past. (The nerve—can you believe those people?)

We're kidding, of course; this critique is both understandable and helpful. 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. This is unavoidable. We've made the conscious decision to study all types of data breaches as they affect all types of organizations, and if small businesses are dropping like flies, we're not going to exclude them because they infest our data. What we can do, however, is to present the results in such a way that they are more readily applicable to certain groups.

We could split the dataset a myriad of ways, but we've chosen (partially due to the initial criticism mentioned above) to highlight differences (and similarities) between smaller and larger organizations (the latter having at least 1000 employees).

We could split the dataset a myriad of ways, but we've chosen (partially due to the initial criticism mentioned above) to highlight differences (and similarities) between smaller and larger organizations (the latter having at least 1000 employees). We hope this alleviates these concerns and makes the findings in this report both generally informative and particularly useful.

Oh—and though we don't exactly condone schadenfreude, we do hope you'll find it enjoyable.

Demographics

Every year we begin with the demographics from the previous years' breach victims because it sets the context for the rest of the information presented in the report. Establishing how the breaches break down across industries, company size, and geographic location should help you put some perspective around all the juicy bits presented in the following sections.

This year we altered how we collect some of the demographic data. We decided to stop using our own list of industries and adopt the North American Industry Classification System (which is cross-referenced to other common classifications). As a result, some of the trending and comparisons from the industry breakdown in previous years lose some consistency, but for the most part the classifications map closely enough that comparisons are not without value.

As Figure 3 shows, the top three spots carry over from our last report. The most-afflicted industry, once again, is Accommodation and Food Services, consisting of restaurants (around 95%) and hotels (about 5%). The Financial and Insurance industry dropped from 22% in 2010 to approximately 10% last year. While we derived a range of plausible (and not-so-plausible) explanations for the widening gap between Financial and Food Services, we will reserve most of those for more applicable sections in the report. Suffice it to say that it appears the cybercrime "industrialization" trend that so heavily influenced findings in our last report (and has been echoed by other reports in the industry[7]), is still in full swing.

"The North American Industry Classification System (NAICS) is the standard used by Federal statistical agencies in classifying business establishments for the purpose of collecting, analyzing, and publishing statistical data related to the U.S. business economy.

NAICS was developed under the auspices of the Office of Management and Budget (OMB), and adopted in 1997 to replace the Standard Industrial Classification (SIC) system. It was developed jointly by the U.S. Economic Classification Policy Committee (ECPC), Statistics Canada , and Mexico's Instituto Nacional de Estadistica y Geografia , to allow for a high level of comparability in business statistics among the North American countries."

Source: http://www.census.gov/eos/www/naics/

When looking at the breakdown of records lost per industry in Figure 4, however, we find a very different result. The chart is overwhelmed by two industries that barely make a showing in Figure 3 and have not previously contributed to a large share of data loss—Information and Manufacturing. We'll touch more on this throughout the report, but this surprising shift is mainly the result of a few very large breaches that hit organizations in these industries in 2011. We suspect the attacks affecting these organizations were directed against their brand and for their data rather than towards their industry.

7 For instance, see Trustwave's 2012 Global Security Report discussing growing attacks against franchises.

Figure 3. Industry groups represented by percent of breaches

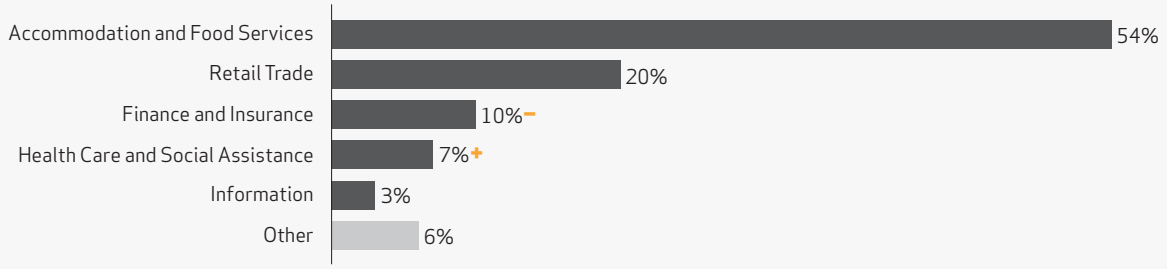

Redrawing Figure 5 with these outliers removed reveals what is perhaps a more representative or typical account of compromised records across industries. Figure 4 is a bit more in line with historical data and also bears some resemblance to Figure 3 above.

Figure 4. Compromised records by industry group

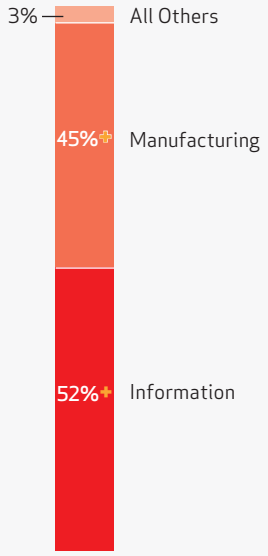

Figure 5: Compromised records by industry group with breaches >1M records removed

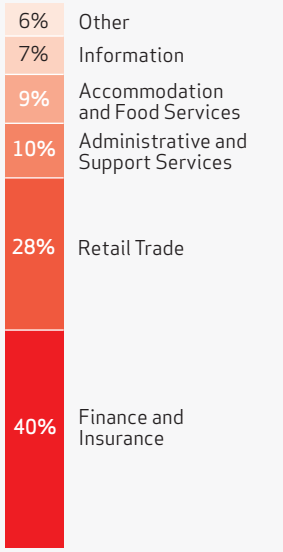

Once again, organizations of all sizes are included among the 855 incidents in our dataset. Smaller organizations represent the majority of these victims, as they did in the last DBIR. Like some of the industry patterns, this relates to the breed of "industrialized" attacks mentioned above; they can be carried out against large numbers in a surprisingly short timeframe with little to no resistance (from the victim, that is; law enforcement is watching and resisting. See the "Discovery Methods" section as well as Appendix B.). Smaller businesses are the ideal target for such raids, and money-driven, risk-averse cybercriminals understand this very well. Thus, the number of victims in this category continues to swell.

The rather large number of breaches tied to organizations of "unknown" size requires a quick clarification. While we ask DBIR contributors for demographic data, sometimes this information is not known or not relayed to us. There are valid situations where one can know details about attack methods and other characteristics, but little about victim demographics. This isn't ideal, but it happens. Rather than brushing these aside as useless data, we're using what can be validated and simply labeling what can't as "unknown." (See Table 2.)

As mentioned in the Methodology section, we will be breaking out findings where appropriate for larger organizations. **By "larger" we're referring to those in our sample with at least 1000 employees.** Remember that as you read this report. So that you have a better idea of the makeup of this subset, Figure 6 shows the industries of the 60 organizations meeting this criterion.

Table 2. Organizational size by number of breaches (number of employees)

1 to 10	42
11 to 100	570
101 to 1,000	48
1,001 to 10,000	27
10,001 to 100,000	23
Over 100,000	10
Unknown	135

Figure 6. Industry groups represented by percent of breaches – LARGER ORGS

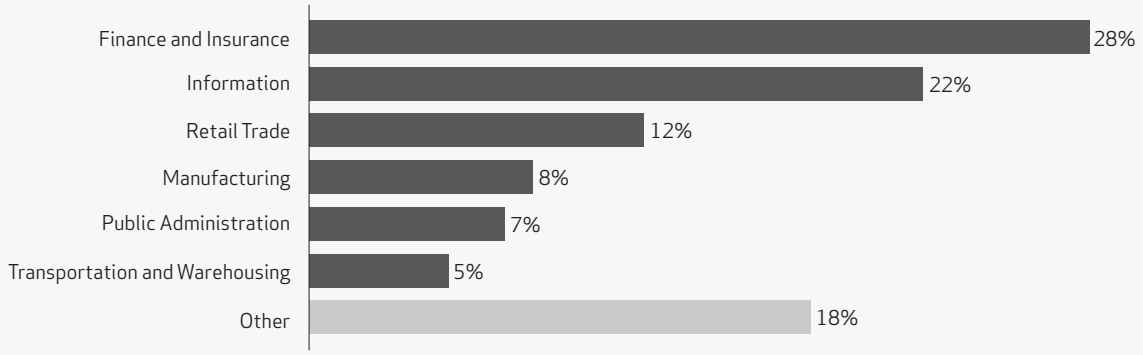

As usual, it's hard to pull meaning from where victims base their operations, since most breaches do not require the attacker to be physically present in order to claim their prize. We set a high mark in 2010 with 22 countries represented, but smashed that record in 2011 with a whopping 36 countries hosting organizations that fell victim to a data compromise. This is an area where the contributions of our global law enforcement partners really highlight the fact that data breaches are not an isolated regional problem.

Figure 7. Countries represented in combined caseload

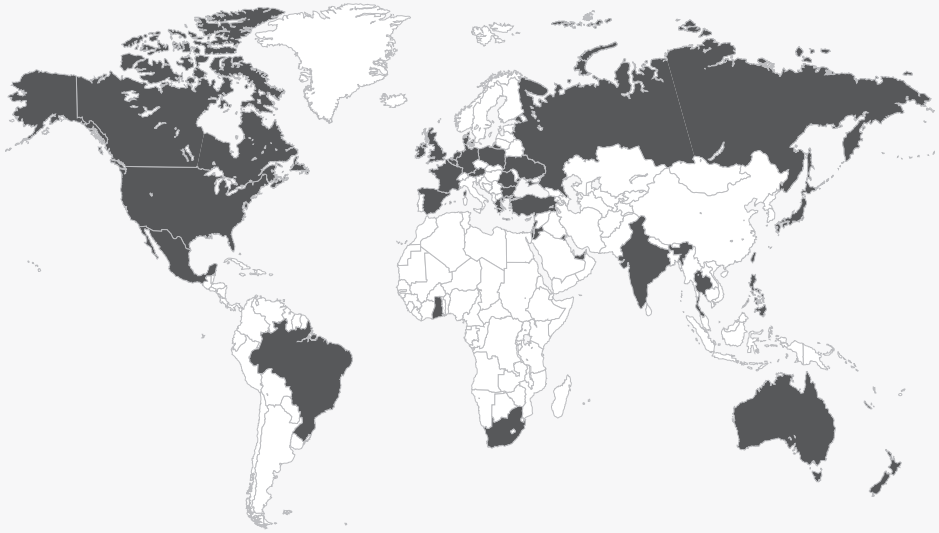

Countries in which a breach was confirmed

Australia	France	Jordan	Poland	United Arab Emirates
Austria	Germany	Kuwait	Romania	Ukraine
Bahamas	Ghana	Lebanon	Russian Federation	United Kingdom
Belgium	Greece	Luxembourg	South Africa	United States
Brazil	India	Mexico	Spain	
Bulgaria	Ireland	Netherlands	Taiwan	
Canada	Israel	New Zealand	Thailand	
Denmark	Japan	Philippines	Turkey	

We set a high mark in 2010 with 22 countries represented, but smashed that record in 2011 with a whopping 36 countries hosting organizations that fell victim to a data compromise.

2011 DBIR: Threat Event Overview

In last year's DBIR, we presented the VERIS threat event grid populated with frequency counts for the first time. Other than new data sharing partners, it was one of the most well received features of the report. The statistics throughout this report provide separate analysis of the Agents, Actions, Assets, and Attributes observed, but the grid presented here ties it all together to show intersections between the four A's. It gives a single big-picture view of the threat events associated with data breaches in 2011. Figure 8 (overall dataset) and Figure 9 (larger orgs) use the structure of Figure 1 from the Methodology section, but replace TE#s with the total number of breaches in which each threat event was part of the incident scenario.[8] This is our most consolidated view of the 855 data breaches analyzed this year, and there are several things worth noting.

Figure 8. VERIS A⁴ Grid depicting the frequency of high-level threat events

		Malware			Hacking			Social			Misuse			Physical			Error			Environmental		
		Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt
Servers	Confidentiality & Possession	381			518		1				9	8	1					2	1			
	Integrity & Authenticity	397			422		1				6	1	1									
	Availability & Utility	2			6						5											
Networks	Confidentiality & Possession										1											
	Integrity & Authenticity	1									1											
	Availability & Utility	1			1						1											
User Devices	Confidentiality & Possession	356			419						1			86								
	Integrity & Authenticity	355			355						1	1		86								
	Availability & Utility										1			3								
Offline Data	Confidentiality & Possession											23							1			
	Integrity & Authenticity																					
	Availability & Utility																					
People	Confidentiality & Possession							30	1													
	Integrity & Authenticity							59	2													
	Availability & Utility																					

When we observe the overall dataset from a threat management perspective, only 40 of the 315 possible threat events have values greater than zero (13%). Before going further, we need to restate that not all intersections in the grid are feasible. Readers should also remember that this report focuses solely on data breaches. During engagements where we have worked with organizations to "VERIS-ize" all their security incidents over the course of a year, it's quite interesting to see how different these grids look when compared to DBIR datasets. As one might theorize, Error and Misuse as well as Availability losses prove much more common.

8 In other words, 381 of the 855 breaches in 2011 involved external malware that affected the confidentiality of a server (the top left threat event).

Figure 9. VERIS A⁴ Grid depicting the frequency of high-level threat events – LARGER ORGS

		Malware			Hacking			Social			Misuse			Physical			Error			Environmental		
		Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt	Ext	Int	Prt
Servers	Confidentiality & Possession	7			33						3							2	1			
	Integrity & Authenticity	10			18						1											
	Availability & Utility				1																	
Networks	Confidentiality & Possession																					
	Integrity & Authenticity																					
	Availability & Utility	1			1																	
User Devices	Confidentiality & Possession	3			6									10								
	Integrity & Authenticity	4			2									10								
	Availability & Utility													1								
Offline Data	Confidentiality & Possession											1							1			
	Integrity & Authenticity																					
	Availability & Utility																					
People	Confidentiality & Possession							7														
	Integrity & Authenticity							11														
	Availability & Utility																					

USING VERIS FOR EVIDENCE-BASED RISK MANAGEMENT

This may sound like an advertisement, but it's not—you can do this using VERIS (which is free!). Imagine, as a risk manager, having access to all security incidents within your organization classified using VERIS (if you really want to let your imagination run wild, think about also having similar data from other organizations like your own). Over time, a historical dataset is created, giving you detailed information on what's happened, how often it's happened, and what hasn't happened within your organization. Unknowns and uncertainties begin to recede. You give it to your data visualization guy who cranks out a grid for your various business groups similar to Figure 9. Hotspots on the grid focus your attention on critical problem areas and help to properly diagnose underlying ailments. From there, treatment strategies to deter, prevent, detect, or help recover from recurring (or damaging) threat events can be identified and prioritized. But you don't stop there; you actually measure the effectiveness of your prescriptions to track whether incidents and losses decrease after these treatments are administered. Thus, you achieve a state where better measurement enables better management. Colleagues start referring to you as the "Risk Doctor" and suddenly your opinion matters in security spending discussions. This could be you.

Obviously, this is meant to be tongue in cheek, but we truly do believe in the merit of an approach like this. We like to refer to this approach as "Evidence-Based Risk Management" (EBRM), borrowing from the concept of evidence-based medicine. Essentially, EBRM aims to apply the best available evidence gained from empirical research to measure and manage information risk. Security incidents, whether large or small, are a huge part of that "best available evidence." This is why we assert that meticulously analyzing them is a highly beneficial practice.

Now back to the grids, where the results for the overall dataset share many similarities with our last report. The biggest changes are that hotspots in the Misuse and Physical areas are a little cooler, while Malware and Hacking against Servers and User Devices are burning brighter than ever. Similarly, the list of top threat events in Table 3 feels eerily familiar.

The results for the overall dataset share many similarities with our last report. The biggest changes are that hotspots in the Misuse and Physical areas are a little cooler, while Malware and Hacking against Servers and User Devices are burning brighter than ever.

Table 3. Top 10 VERIS threat events

	Threat Event	Threat Event #	Counts
1	External.Hacking.Server.Confidentiality	4	518
2	External.Hacking.Server.Integrity	28	422
3	External.Hacking.UserDevice.Confidentiality	130	419
4	External.Malware.Server.Integrity	22	397
5	External.Malware.Server.Confidentiality	1	381
6	External.Malware.UserDevice.Confidentiality	127	356
7	External.Malware.UserDevice.Integrity	148	355
8	External.Hacking.UserDevice.Integrity	151	355
9	External.Physical.UserDevice.Confidentiality	139	86
10	External.Physical.UserDevice.Integrity	160	86

Table 4. Top 10 VERIS threat events – LARGER ORGS

	Threat Event	Threat Event #	Counts
1	External.Hacking.Server.Confidentiality	4	33
2	External.Hacking.Server.Integrity	28	18
3	External.Social.People.Integrity	280	11
4	External.Malware.Server.Integrity	22	10
5	External.Physical.UserDevice.Confidentiality	139	10
6	External.Physical.UserDevice.Integrity	160	10
7	External.Malware.Server.Confidentiality	1	7
8	External.Social.People.Confidentiality	259	7
9	External.Hacking.UserDevice.Confidentiality	130	6
10	External.Malware.UserDevice.Integrity	148	4

Separating the threat events for larger organizations in Figure 9 yields a few additional talking points. Some might be surprised that this version of the grid is less "covered" than Figure 8 (22 of the 315 events—7%—were seen at least once). One would expect that the bigger attack surface and stronger controls associated with larger organizations would spread attacks over a greater portion of the grid. This may be true, and our results shouldn't be used to contradict that point. We believe the lower density of Figure 9 compared to Figure 8 is mostly a result of size differences in the datasets (855 versus 60 breaches). With respect to threat diversity, it's interesting that the grid for larger organizations shows a comparatively more even distribution across in-scope threat events (i.e., less extreme clumping around Malware and Hacking). Related to this, Social and Physical events make the top 10 list in Table 4. Based on descriptions in the press of prominent attacks leveraging forms of social engineering, this isn't a shocker.

Naturally, we'll expound on all of this throughout the following sections.

Threat Agents

Entities that cause or contribute to an incident are known as threat agents. There can, of course, be more than one agent 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 agent-specific sections). Identification of the agents associated with an incident is critical to taking specific corrective actions as well as informing decisions regarding future defensive strategies. VERIS specifies three primary categories of threat agents—External, Internal, and Partner.

VERIS Classification Note: If the agent's role in the breach is limited to a contributory error, the agent would not be included here. For example, if an insider's unintentional misconfiguration of an application left it vulnerable to attack, the insider would not be considered a threat agent if the application were successfully breached by another agent. An insider who deliberately steals data or whose inappropriate behavior (e.g., policy violations) facilitated the breach would be considered a threat agent in the breach.

- **External:** External threats originate from sources outside of the organization and its network of partners. Examples include former employees, lone hackers, organized criminal groups, and government entities. External agents also include environmental events such as floods, earthquakes, and power disruptions. Typically, no trust or privilege is implied for external entities.
- **Internal:** Internal threats are those originating from within the organization. This encompasses company executives, employees, independent contractors, interns, etc., as well as internal infrastructure. Insiders are trusted and privileged (some more than others).
- **Partners:** Partners include any third party sharing a business relationship with the organization. This includes suppliers, vendors, hosting providers, outsourced IT support, etc. Some level of trust and privilege is usually implied between business partners.

Figure 10 displays the distribution of threat agents by percentage of breaches in this year's dataset, along with all previous years of this study. It's important to keep in mind that we're not looking at a consistent sample. The first few years were based only on Verizon cases, then the USSS (2007-2011), NHTCU (2006-2011), AFP (2011), IRISSCERT (2011), and PCeU (2011) joined at various points in the years that followed. Thus, trends are the combination of changes in the threat environment and changes within the sample dataset.

Figure 10. Threat agents over time by percent of breaches

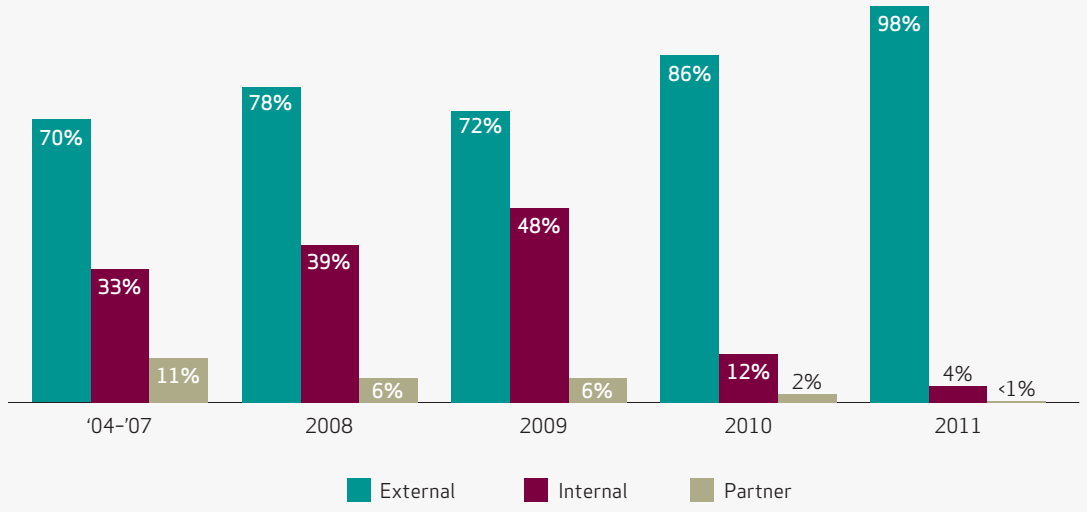

2011 continued the shift towards external agents' involvement in a high percentage of data breaches. Though we have always seen an external majority, never before has any year been so one-sided. 2009 was the closest to an exception to that rule, but the rise in internal agents was mostly the by-product of incorporating the insider-heavy USSS caseload (see the 2010 DBIR[9] for more detail). Since then, it's been primarily outsiders in the caseloads we've examined.

2011 continued the shift towards external agents' involvement in a high percentage of data breaches. Though we have always seen an external majority, never before has any year been so one-sided.

Apart from yearly sample variations, there are several factors contributing to the escalating percentage of external agents vs. insiders and partners in this report. The primary factor, which was addressed at length in the 2011 DBIR[10], is the continued effect of "industrialized" attacks on these ratios. Organized criminal groups targeting payment card information from Internet-facing POS systems or physically-exposed ATMs and gas pumps can launch a sting against hundreds of victims during the same operation. From a percentage standpoint, the resulting effect that these commoditized yet highly-scalable attacks have on threat agent trends makes perfect sense. Insiders, by definition, have a smaller number of potential targets.

Another contributor to the continued rise of external agents in 2011 was the reinvigorated conducts of activist groups. Commonly known as "hacktivism," these attacks are inherently external in nature. They are not nearly as frequent (one might even say "constant") as mainline cybercrime, but as will be seen below, they can be quite damaging.

We would be remiss if we did not point out that in 2011, there were several investigations involving internal agents that did not meet the definition of a data breach. When insiders misuse access or information provided for their job duties, but did not disclose information to an unauthorized party, then no loss of confidentiality has occurred.[11] Such incidents are not included in this report.

Another interesting observation about 2011 is the much lower percentage of multi-agent breaches. Back in 2009, over one-quarter of all incidents was the work of more than one category of threat agent. Such incidents sometimes involve overt collusion, but more often outsiders solicit insiders to participate in some aspect of the crime. In 2011, that figure was just 2%. The decline here can also be attributed to the "industrialization" trend discussed above.

Partner threat agents have realized a steady decrease over the last few years, and this dataset is no exception.[12] With less than 1% of breaches caused by a partner, it will be hard to go anywhere but up in the next report. Similar to insiders, the dramatic increase in external agents helps to explain this decline, but there are other factors as well. Notice that the downward trend began in 2008, which precedes the major shift towards highly-scalable attacks by outsiders. We have given several hypotheses in past reports, including increased awareness, regulation, and technology advancements. More significant is how we define causal and contributory agents. Partners that did not have a causal role in the incident are not included in these percentages. More discussion on such scenarios can be found in the Partner and Error sections of this report.

It is also entirely possible that malicious insiders and/or partners are flying under the radar and thus avoiding discovery. We have lamented in previous reports (and will lament in later sections) that a high percentage of breaches are identified by fraud detection. However, compromises of non-financial data do not have these mechanisms to trigger awareness, and are therefore more difficult to discover. Our data consistently shows that trusted parties are

9 http://www.verizonbusiness.com/resources/reports/rp_2010-data-breach-report_en_xg.pdf
10 http://www.verizonbusiness.com/go/2011dbir/us/
11 A frequent example of this is a bank employee who uses system privileges to make an unauthorized withdrawal or transfer of funds. This is certainly a security violation, but it is not a data breach.
12 Some may rightly remember that the percentage tied to partners was substantially higher in prior reports. Keep in mind that those reports showed Verizon data separately, whereas this is the combined data from all participating organizations "retrofitted" to historical data. It definitely changes the results.

Figure 11. Threat agents (exclusive) by percent of breaches

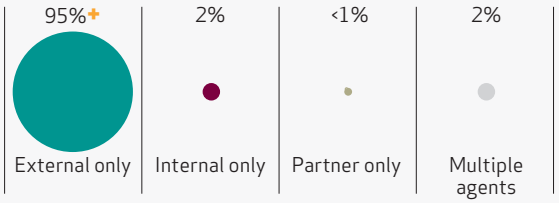

Figure 12. Threat agents by percent of breaches – LARGER ORGS

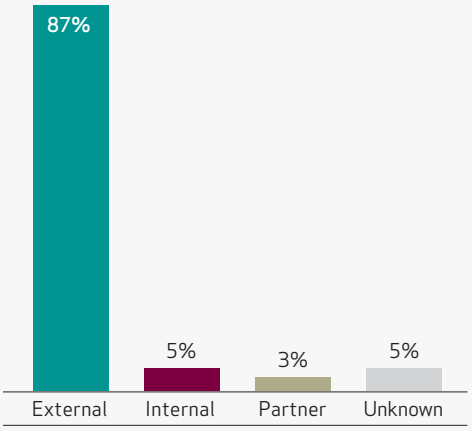

considerably more likely to steal intellectual property and other sensitive (non-financial) data, and there's a good chance these activities would never be detected. This is not included to "apologize" for bias or to spread FUD, but to raise a valid point that insiders and partners are probably under-represented in Figure 10 (though, in the grand scheme of things, we still don't think they're anywhere close to outsiders).

In keeping with our promise to give findings specific to larger organizations, we present Figure 12. Those hoping to see a significantly different result here are bound for disappointment. (Don't you hate it when data gets in the way of a good theory?) We had an incredibly insightful and rational explanation ready to explain why insiders and partners were more likely to attack larger organizations, but alas, it's gone to waste.

Breach Size by Threat Agents

Data compromise, as measured by number of records lost, is not indicative of the full impact of the breach, but is a useful and measurable indicator of it. We agree that it would be optimal to include more information on losses associated with response, brand damage, business disruption, legal penalties, etc. As a small step in this direction, we have added a short section to this report discussing some of these consequences. Here, we focus exclusively on the amount of data loss.

Figure 13 shows the distribution among threat agents of the approximately 174 million records compromised across the merged 2011 dataset. No, we didn't forget to include bubbles for insiders and partners; it's just that outsiders stole virtually all of it. When compared to the entire dataset encompassing all years of this study (Figure 14), the effect isn't much different (but we can at least see colors other than greenish-blue). Mega-breaches, involving millions of records in a single incident, have consistently skewed data loss numbers toward external agents. The high-volume, low-yield attacks also mount up in their favor over time.

It's important to recognize the various types of data compromised and their influence on this metric. Payment card data and personal information are frequently stored and stolen in bulk, whereas intellectual property or classified data theft often involve only a single "record." As mentioned previously, insiders are more likely to target the latter.

Figure 13. Compromised records by threat agent, 2011

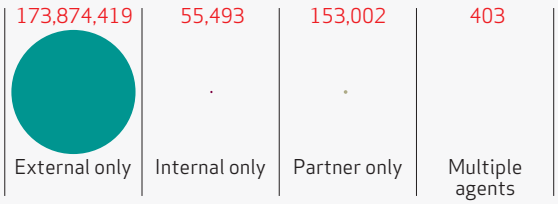

Figure 14. Compromised records by threat agent, 2004-2011

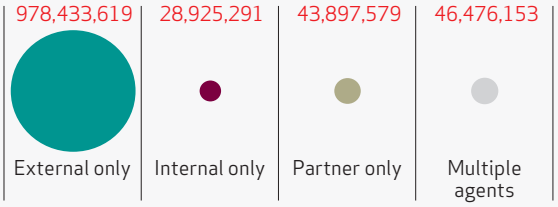

External Agents (98% of breaches, 99+% of records)

As with all of our previous DBIRs, this version continues to reinforce the finding that external parties are responsible for far more data breaches than insiders and partners. This go-around, they were tied to 98% of all incidents. At a quick glance, much about the roles, varieties, and motives of external agents in 2011 appears to be just a continuation of the same ol' story.

Bottom line: most data thieves are professional criminals deliberately trying to steal information they can turn into cash. Like we said—same ol' story.

Outsiders almost always engaged in direct, intentional, and malicious actions. Only a scant 2% of cases featured external agents in indirect roles, where they solicited or aided someone else to act against the victim. Organized criminal groups were once again behind the lion's share (83%) of all breaches. One may wonder why it is they do what they do (we surely do, and that's why we started tracking more about motives last year), the answer is pretty straightforward—they do it for the money (96%). Bottom line: most data thieves are professional criminals deliberately trying to steal information they can turn into cash. Like we said—same ol' story.

Figure 15. Motive of external agents by percent of breaches within external

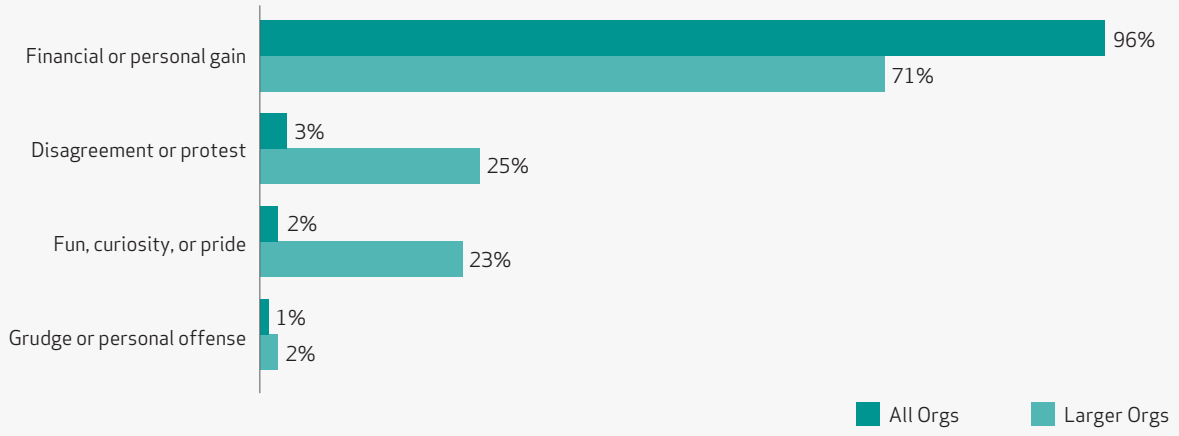

It's not the whole story, however. Nor is it the most important one. The most significant change we saw in 2011 was the rise of "hacktivism" against larger organizations worldwide. The frequency and regularity of cases tied to activist groups that came through our doors in 2011 exceeded the number worked in all previous years combined.

It's not the whole story, however. Nor is it the most important one. The most significant change we saw in 2011 was the rise of "hacktivism" against larger organizations worldwide.

But this was not restricted to our caseload alone; the other organizations participating in this report also spent a great deal of effort responding to, investigating, and prosecuting hacktivist exploits. It was extremely interesting to piece these different perspectives together to form a global view of investigations into activist groups and their victims. 3% of all external attacks may not seem like much (though remember we're dealing with over 850 incidents here, and notice related motives are higher than that; plus we suspect some "unknown" agents are actually activists), but this trend is probably the biggest and single most important change factor in this year's DBIR.

That is not to say that hacktivsm is new; the term has been standard lexicon since it was coined by the Cult of the Dead Cow hacker collective in the late 90's.[13] Back then, it mostly consisted of website defacements, coordinated denial of service attacks, and other antics to express disagreement, obtain bragging rights, or "just because." The major shift that occurred in 2011 was that activist groups added data breaches to their repertoire with much-heightened intensity and publicity. In other words, 2011 saw a merger between those classic misdeeds and a new *"oh by the way, we're gonna steal all your data too"* twist.

Table 5. Varieties of external agents by percent of breaches within External and percent of records

	All Orgs		Larger Orgs	
Organized criminal group	83%	35%−	33%	36%
Unknown	10%	1%	31%	0%
Unaffiliated person(s)	4%	0%	10%	0%
Activist group	2%	58%+	21%	61%
Former employee (no longer had access)	1%	0%	6%	0%
Relative or acquaintance of employee	0%	0%	2%	0%

But even that's not the whole story. Although activist groups accounted for a relatively small proportion of the 2011 caseload, they stole over 100 million records. That's almost twice the amount pinched by all those financially-motivated professionals we discussed earlier. So, although ideological attacks were less frequent, they sure took a heavy toll.

Why the disparity between the total records stolen by professional cybercriminals versus activist groups? Looking through the case data, it is apparent that money-driven crooks continue to focus more on opportunistic attacks against weaker targets. This may be at least partly because a good number of their brethren are enjoying jail time. Instead of major (and risky) heists, they pilfer smaller hauls of data from a multitude of smaller organizations that present a lower risk to the attacker. Think of it as a way to streamline business processes. Find an easy way to prey on the unsuspecting, the weak, and the lame, and then simply repeat on a large scale. This high-volume, low-yield business model has become the standard M.O. for organized criminal groups.

An important observation before we close this discussion is that nearly all data stolen by activist groups were taken from larger organizations. Furthermore, the proportion of breaches tied to hacktivism-related motives rises to 25 percent. This stands to reason, since a low-profile brand is less likely to draw the ire of these groups.

Just like the security professionals with whom they contend, criminals are constantly assessing risk—the risk of apprehension. One of the greatest challenges for law enforcement in the fight against cybercrime is merging a criminal's real world identity with their online identity. Unfortunately, across 10% of the 2011 caseload, investigators were unable to identify a specific variety of external agent. There are several valid reasons for this. First and foremost, many clients do not maintain sufficient log data that would enable attribution. In many cases, the determination cannot be made through disk forensics alone. Many victims (for various reasons) do not wish to expand the investigation to include this line of inquiry once the breach has been successfully contained. Sometimes the perpetrator is able to erase his tracks or hide them among a host of intermediary systems. Every now and then, just as we think we've correctly identified the assailant—nope! Chuck Testa *(just look it up—it's worth the break).*

Origin of External Agents

As is always the case, determining the geographic origin of external attackers based solely on IP address can be problematic. Even if the country of the source IP addresses can be pinpointed, this may not be where the attacker actually resides. It's quite likely that it's just a host in a botnet or another "hop" used by the agent. In some cases, various types of additional data, such as those provided by law enforcement and/or netflow analysis, can help to

13 http://www.wired.com/techbiz/it/news/2004/07/64193

determine the attacker's true origin. Either way, examining the geographic origin of attacks is valuable for a number of reasons.

2011 findings look similar to previous years, with threat agents hailing from Eastern Europe accounting for two-thirds of all external breaches (see Figure 16). However, if examining only large organizations, this number drops to 27%. This statistic falls in line with the increasing tendency of organized criminal groups (that often hail from Eastern Europe) to target smaller, lowest-hanging-fruit victims. Attacks against larger organizations originated from a comparatively more diverse set of regions around the world.

Figure 16. Origin of external agents by percent of breaches within External

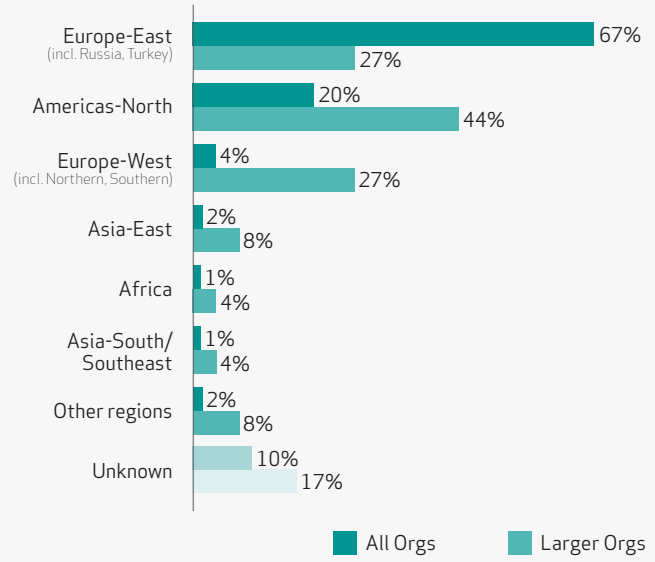

Internal Agents (4% of breaches, <1% of records)

As discussed in the Threat Agent Overview section, the decline of internal agents as a percentage of our dataset is due more to the continued rise of industrialized attacks than the demise of all insider crime. We hypothesize that many insider crimes go unreported because the organization is unaware of them, or because they decide for political reasons to handle it internally in lieu of calling for a third-party forensic investigation or referring it to law enforcement.

Nevertheless, when insiders do directly cause or contribute to a data breach, they do so in multiple ways. For our purposes, we classify them according to three main roles. Insiders either acted deliberately and maliciously, inappropriately (but not maliciously), or acted unintentionally. For the third year in a row, nearly all the internal breaches were a result of deliberate and malicious actions (each year ~90%). It should be noted, however, that there were a handful of unintentional errors made by insiders in our caseload that directly led to data loss. In these instances, it was due to an employee accidentally publishing information to the web that shouldn't have been made public.

We hypothesize that many insider crimes go unreported because the organization is unaware of them, or because they decide for political reasons to handle it internally.

There are many ways that an insider may indirectly contribute to an incident,[14] but they are not considered a threat agent in such circumstances, and thus are not the focus of this section. What we're dealing with here are scenarios where insiders were the direct or primary cause of data loss within their organizations.

This year we also separated out "cashiers/tellers/waiters" from the "regular employee/end users" category. We found that the sorts of actions involved with these money handlers were quite different than those of traditional end users within corporations. By so doing, we are able to get a more accurate picture of who's behind the historically-large percentage of incidents attributed to regular employees. The money handlers mentioned above account for 65% of all internal incidents. These individuals, often solicited by external organized gangs, regularly

14 See the Partner Agents and Error sections for discussion and examples of how an agent can contribute indirectly to an incident, but not be considered a threat agent.

Table 6. Types of internal agents by percent of breaches within Internal

Cashier/Teller/Waiter	65%
Manager/Supervisor	15%
Regular employee/end-user	12%
Finance/Accounting staff	6%
System/network administrator	6%
Auditor	3%
Executive/Upper Management	3%
Internal system or site	3%
Unknown	3%

skim customer payment cards on handheld devices designed to capture magnetic stripe data. The data is then passed up the chain to criminals who use magnetic stripe encoders to fabricate duplicate cards. Not surprisingly, such incidents are almost entirely associated with smaller businesses or independent local franchises of large brands.

On the other hand, when regular corporate end users are involved (12%), their actions are quite different. In most instances, these employees abuse system access or other privileges in order to steal sensitive information. Almost all of the scenarios listed above are motivated by financial or personal gain.

Outside of the varieties mentioned above, we observed a mixture of executives, managers, and supervisors (totaling 18%). Like the regular employees and end users, these individuals are also exploiting system access and privileges for personal gain. For three years running, we have seen a decline in finance and accounting staff. Still, the daily responsibilities of these folks, which involve the oversight of and/or direct access to valuable assets and information, put them in a position to engage in a multitude of misdeeds. One can't help but wonder what the data would show if we were to track these types of insiders through the ever-changing regulatory landscape, from before Glass-Stegall, to Graham-Leach-Bliley, and now to Dodd-Frank. The ebb and flow of these numbers would have been very interesting to witness.

Finally, it might be negligent of us if we didn't provide some mention of system or network administrators. These trusty technological warriors help make IT organizations around the world hum with productivity, and they oftentimes possess the proverbial "keys-to-the-kingdom." Though we have seen cases in which they were responsible for data breaches, they have barely registered more than a blip on the radar in the last couple of years. We mentioned in an earlier section that we have analyzed the incidents for a single organization over the course of a year. In these datasets, admin-related incidents occur frequently, but they are mostly of the availability and downtime variety.

Partner Agents (<1% of breaches, <1% of records)

Continuing the trend we observed in 2010, breaches caused by partners were few and far between. There were exactly three (that is correct—three—the same number as our last report) partner breaches in the entire combined 2011 caseload. In two of those, a publishing error was identified as the primary cause; the partner accidently posted sensitive data to a public-facing website. The third partner-sourced breach involved deliberate and malicious misuse motivated by financial gain. A third-party database developer identified a SQL vulnerability while performing contract work and then abused this knowledge in order to compromise the victim.

Note that the statistic above refers only to partners identified as a threat agent (the direct cause/contributor); it does not include the many other ways a partner can indirectly factor into the circumstances surrounding the breach. We realize this is a bit confusing. Indeed, we have received a number of inquiries from DBIR readers and VERIS users about whether various scenarios should be attributable to partners (and insiders too, for that matter). Having nothing further to say about the three incidents above, we will switch gears and try to give some clarification on how we classify the role of partners in an incident. If you never plan to use VERIS and/or just don't care, skip it.

A few examples should help:

1. If the partner's actions are the direct cause of the incident, they ARE a threat agent.
2. If the partner's actions create a circumstance or condition that—if/when acted upon by another agent—allows the primary chain of threat events to proceed, the partner is NOT a threat agent. We consider this to be a conditional event, and the partner can be viewed as a contributing agent. Their actions had more to do with the victim's security or vulnerability than the threat itself.
3. If the partner owns, hosts, or manages the victim's asset involved in an incident, it does NOT necessarily follow that they are a threat agent. They may be (if their actions led to the incident), but they are not guilty simply by this association.

Example #2 seems to be a sticking point for most people. To further illustrate what we mean, let us consider the following scenario. Suppose a third party remotely administers a customer's devices over the Internet via some kind of remote access or desktop service. Further suppose this partner forgot to enable or misconfigured a security setting (let's pick something no admin would ever do, like neglecting to change default credentials). Then, lo and behold, that device gets popped within 30 seconds of being identified when an organized criminal group operating out of Eastern Europe guesses the username/password. All of this, of course, is purely figurative; this would never actually happen in the real world (wink, wink). In such circumstances, the criminal group would be the only threat agent. One could capture the partner's [indirect] contribution using the VERIS-specified role of "contributed to conditional event(s)" along with a suitable "Error" threat action. This essentially notes that the partner created a vulnerability (the conditional event) that was exploited by the external threat agent.

A partner's lax security practices and poor governance—often outside the victim's control or expertise—are frequently catalysts in security incidents.

All in all, the assertion made for the last two years remains true: organizations that outsource their IT management and support also outsource a great deal of trust to their chosen partners. A partner's lax security practices and poor governance—often outside the victim's control or expertise—are frequently catalysts in security incidents. Nevertheless, outsourcing can have many benefits, and the best way to counteract the associated risk is through third-party policies, contracts, controls, and assessments. One caveat of outsourcing is that you can outsource business functions, but you cannot outsource the risk and responsibility to a third party. These must be borne by the organization that asks the population to trust they will do the right thing with their data.

Threat Actions

Threat actions describe what the threat agent did to cause or to contribute to the breach. Every data breach contains one or more of them, causing percentages to add up to more than 100%. Within VERIS, actions are classified into seven high-level categories (each of which will be covered in detail in the following sections).

Hacking and malware have traditionally led the pack, but this year they've pulled away from the group even further while waving "Hi Mom!" to the camera. Out of the 855 incidents this year, 81% leveraged hacking, 69% included malware, and an impressive 61% of all breaches featured a combination of hacking techniques and malware. Out of the 602 incidents with two or more events, hacking and malware were used in 86% of the attacks (more on the relationships of threat actions can be found in Appendix A).

Hacking and malware have traditionally led the pack, but this year they've pulled away from the group even further while waving "Hi Mom!" to the camera.

Figure 17. Threat action categories over time by percent of breaches and percent of records

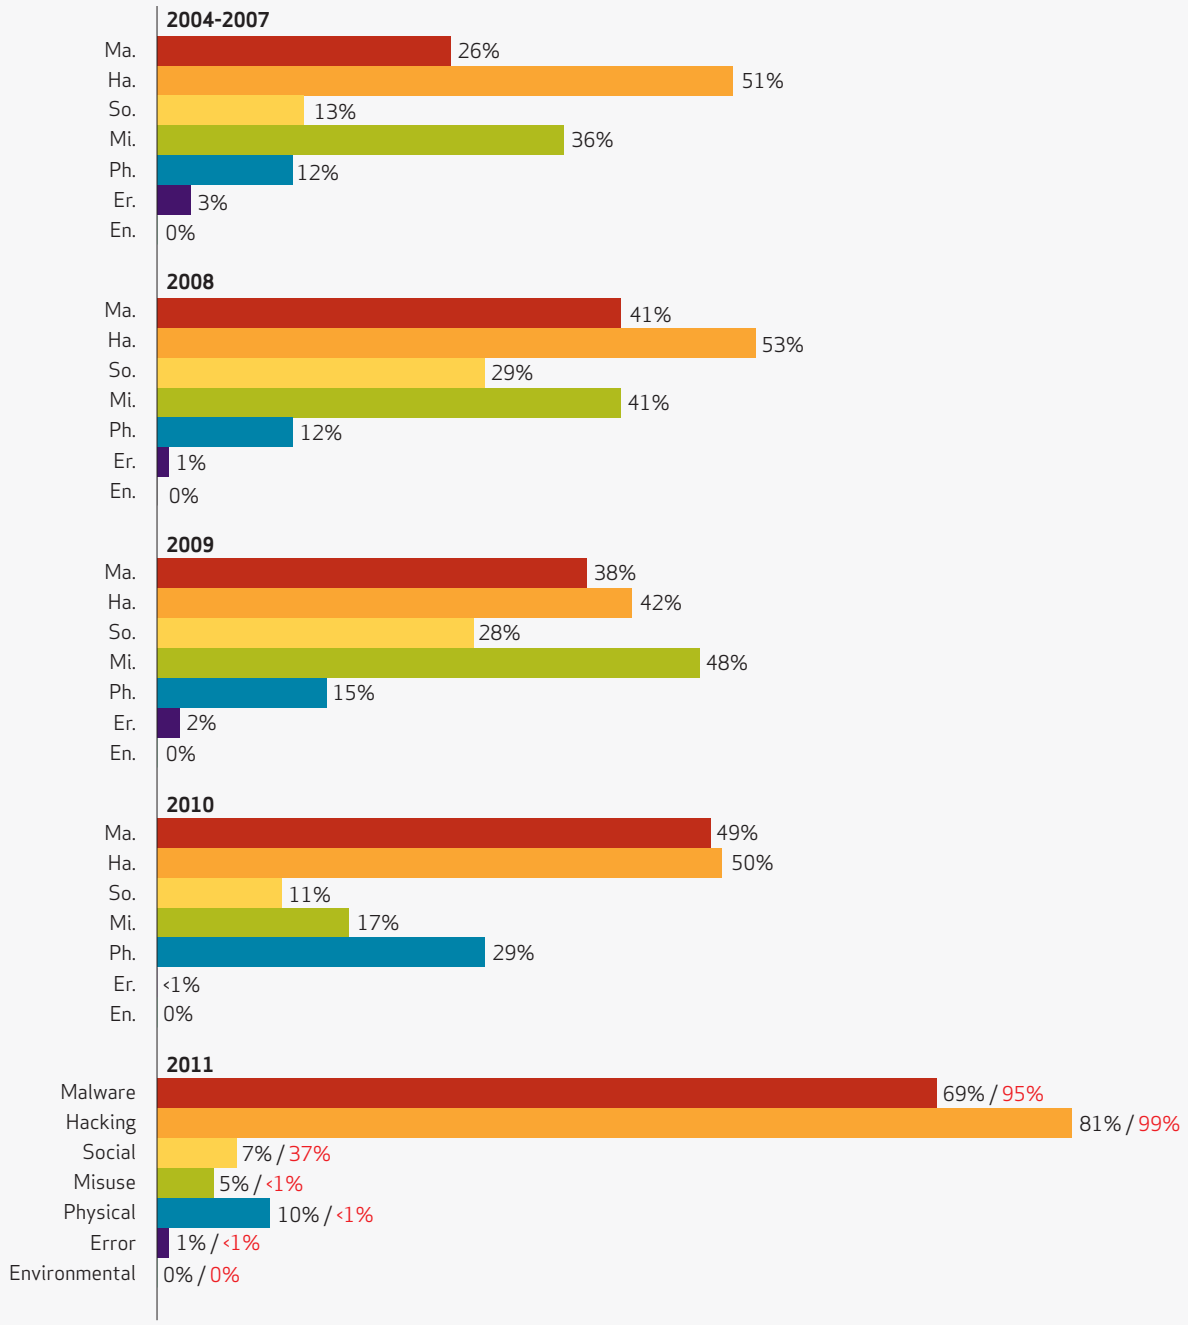

Overall, we've seen the categories bounce around a bit over the years. Misuse and social tactics stepped up their game in 2009 while physical techniques made a respectable appearance the year after that. The rather sharp drop in physical attacks this past year may be due to global law enforcement agencies successfully flipping the freedom bit on those involved with skimming incidents. They focused heavily on the criminal rings behind these skimming activities rather than individual incidents themselves, and we may be starting to see the fruits of those efforts.

Figure 18. Threat action categories by percent of breaches and percent of records – LARGER ORGS

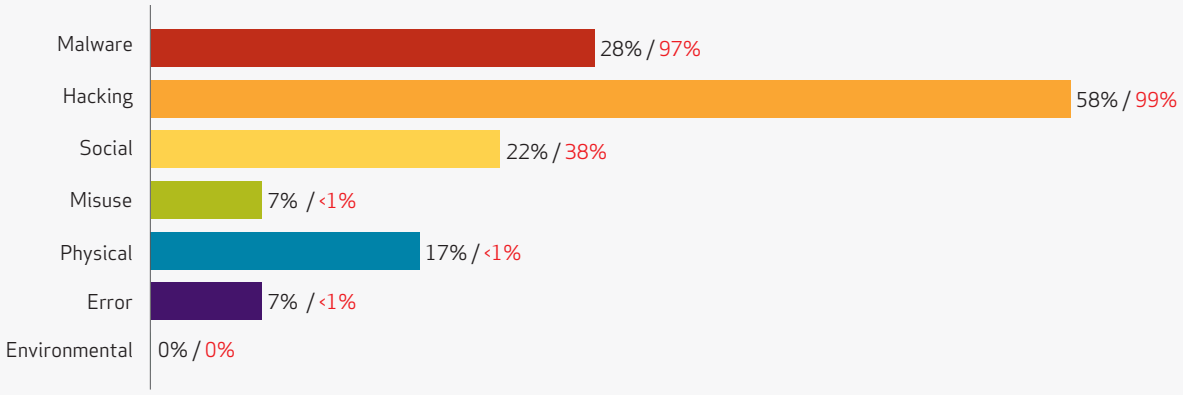

If we look at bigger organizations, however, we find a slightly different picture. Figure 18 hints at an obvious and simple truth worth mentioning: large company problems are different than small company problems.

Whatever the explanation, one thing is absolutely clear: we see a definite pattern emerging over the years with respect to threat actions across the full dataset.

If we look at bigger organizations, however, we find a slightly different picture. Figure 18 hints at an obvious and simple truth worth mentioning: large company problems are different than small company problems. Perhaps it's because enterprises have the IT staff to address some of the low-hanging fruit (or, what is often more apropos, the fallen fruit rotting in the yard). However, to get at the actionable items for large versus small organizations, we must look at the breakdown of threat actions beyond these high-level categories (see Table 7).

Table 7. Top 10 Threat Action Types by number of breaches and records

Rank	Variety	Category	Breaches	Records
1	Keylogger/Form-grabber/Spyware (capture data from user activity)	Malware	48%	35%
2	Exploitation of default or guessable credentials	Hacking	44%	1%
3	Use of stolen login credentials	Hacking	32%	82%
4	Send data to external site/entity	Malware	30%	<1%
5	Brute force and dictionary attacks	Hacking	23%	<1%
6	Backdoor (allows remote access/control)	Malware	20%	49%
7	Exploitation of backdoor or command and control channel	Hacking	20%	49%
8	Disable or interfere with security controls	Malware	18%	<1%
9	Tampering	Physical	10%	<1%
10	Exploitation of insufficient authentication (e.g., no login required)	Hacking	5%	<1%

Companies, big and small, saw a fair amount of malicious code designed to capture user inputs, commonly called keyloggers—they were present in almost half of all breaches (48%). This most likely contributed to the use of stolen credentials in roughly one out of three incidents. Another consistent threat action for large and small companies was the installation (and exploitation) of backdoors; these were leveraged in one out of every five attacks. We can get a feel for the differing threat landscapes of big and small companies by comparing Table 8, which lists top threat actions used against larger enterprises.

Table 8. Top 10 Threat Action Types by number of breaches and records – LARGER ORGS

Rank	Overall Rank	Variety	Category	Breaches	Records
1	3	Use of stolen login credentials	Hacking	30%	84%
2	6	Backdoor (allows remote access/control)	Malware	18%	51%
3	7	Exploitation of backdoor or command and control channel	Hacking	17%	51%
4	9	Tampering	Physical	17%	<1%
5	1	Keylogger/Form-grabber/Spyware (capture data from user activity)	Malware	13%	36%
6	11	Pretexting (classic Social Engineering)	Social	12%	<1%
7	5	Brute force and dictionary attacks	Hacking	8%	<1%
8	15	SQL Injection	Hacking	8%	1%
9	20	Phishing (or any type of *ishing)	Social	8%	38%
10	22	Command and Control (listens for and executes commands)	Malware	8%	36%

Pulling information from Table 8 is a little problematic since the numbers are smaller (smaller datasets have larger swings in sampling error), but we can see some interesting trends. The first thing we notice is the increased presence of social tactics; a disproportionate 22% of incidents incorporated these within larger organizations. This could be because they have better perimeter defenses (forcing attackers to target humans instead of systems) or that employees of larger companies have a more complex social web (they are less likely to know all the co-workers they should (or should not) trust).

Another interesting take-away from Table 8 is the lack of exploitation of default credentials. It dropped off the radar and a few of the 60 large company breaches included that threat action. Again, this could be because larger organizations have the talent and resources to tackle some of the menial tasks or it could be that larger companies likely have more than a single default password between the attacker and the crown jewels. This reinforces the need for the bad guys to steal login credentials to breach larger organizations. In the pages that follow, we dig deeper into each of these categories to see what else we can learn about the actions leading to data breaches in 2011.

Malware (69% of breaches, 95% of records)

Malware is any **mal**icious soft**ware**, script, or code developed or used for the purpose of compromising or harming information assets without the owner's informed consent. Malware factored in over two-thirds of the 2011 caseload and 95% of all stolen data. Upon identification of malware during an investigation, the Verizon RISK team conducts an objective analysis to classify and ascertain its capabilities with regard to the compromise at hand. The RISK team uses the analysis to assist the victim with containment, removal, and recovery of the infection. Malware can be classified in many ways, but we utilize a two-dimensional approach within the VERIS framework that identifies the infection vector and the functionality used to breach data. These two dimensions are directly relevant to identifying appropriate detective and preventive measures for malware.

Malware Infection Vectors

Much as it has in the past, the most common malware infection vector continues to be installation or injection by a remote attacker. This covers scenarios in which an attacker breaches a system via remote access and then deploys malware or injects code via web application vulnerabilities. Over the past few years, the data shows that this infection vector continues on an upward trend. Attackers utilized this vector in slightly more than half of malware-related cases in 2009, about 80% in 2010, and a staggering 95% in the past year. Its popularity as an infection vector likely stems both from the attacker's desire to remain in control after gaining access to a system, and its use in high-volume automated attacks against remote access services. This is most evident in the broader financially-motivated crimes (such as payment card breaches) where malware is not typically the initial vector of intrusion, but rather is installed by the attacker after gaining access. This is not always true for other genres of attacks. With IP theft scenarios, malware often provides the entry point after a successful social attack such as a phishing e-mail. In both cases, good defense-in-depth controls, not just antivirus software, could aid in keeping the attacker out in the first place.

Much as it has in the past, the most common malware infection vector continues to be installation or injection by a remote attacker. This covers scenarios in which an attacker breaches a system via remote access and then deploys malware or injects code via web application vulnerabilities.

Figure 19. Malware infection vectors by percent of breaches within Malware

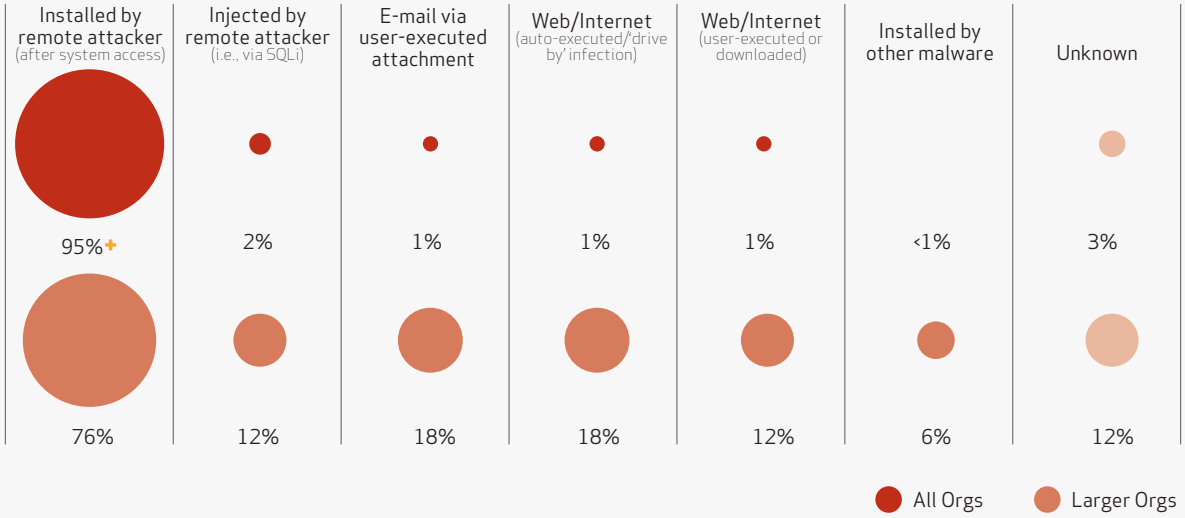

When focusing on data compromise situations, e-mail is less common as an infection vector. Many organizations employ antivirus products and other filtering mechanisms to successfully block or quarantine millions of malware strains floating around the Internet. It is highly likely that e-mail would be a much larger vector if these controls were revoked.

Infections via the web decreased again this past year in proportion to other vectors. We divide web-based malware into two subcategories: code that is auto-executed (a.k.a. drive-by downloads) and software that the user needs to execute (clicking on a malicious hyperlink). We realize that web-based malware results in countless infected systems, but only a portion of those lead to confirmed data thefts.

For many web-based types of malware, a user is required to visit a certain infected website. This certainly works for some scenarios, such as password-stealing Zeus malware, but not for large-scale compromises of payment systems. Most of the infected systems appear simply to join the thousands of botnets used for DDoS and other types of attacks.

For larger organizations, the distribution of malware infection vectors is less one-sided; the data shows higher frequencies of web and e-mail infection vectors and lower frequencies of malware installed directly by attackers.

For larger organizations, the distribution of malware infection vectors is less one-sided; the data shows higher frequencies of web and e-mail infection vectors and lower frequencies of malware installed directly by attackers.

For larger organizations, the distribution of malware infection vectors is less one-sided; the data shows higher frequencies of web and e-mail infection vectors and lower frequencies of malware installed directly by attackers. Our leading theory for this shift is that attackers may find it easier to get users to install malware rather than breach the perimeter defense of larger organizations through a direct attack. The amount of "unknown" infection vectors is attributable to many different factors. Most often it is due to a lack of evidence (no log data, anti-forensics by the attacker, and/or premature clean-up) on the system. In these cases, it is known that malware was present, but the infection vector cannot be conclusively determined.

Malware Functionality

Of equal importance to the pathway of malware infection are the functions exhibited once deployed within a victim's environment. We mostly focus on malware that directly relates to the breach, but we often find all sorts of extraneous malicious or unwanted files during the course of an investigation. This serves as an additional indication of inadequately managed systems and a lack of security processes. Although malware frequently utilizes several methods to harm a system, it still serves one or more of three basic purposes in data breach scenarios: enable or prolong access, capture data, or further the attack in some other manner.

Per Figure 20, the three most commonly found functions of malware continue to be logging keystrokes (and other forms of user input), sending data to external locations, and backdoors. It is important to note that none of these functionalities are mutually exclusive and it's common for a single piece of malware to feature several components.

As mentioned, keyloggers appeared in over two-thirds of malware-related cases, slightly more than the previous year. These tools include commercially available software packages, which are freely available on the web, and for which fully functional pirated versions can be found on peer-to-peer (P2P) networks and torrent sites. Some of these keyloggers also allow the attacker to build a pre-configured remote installation package that can be deployed on a target system. Their availability, ease of use, and configuration, as well as their anti-forensic capabilities, such as hiding from a list or running processes, make them attractive for attackers to use.

The next most common functions relate to data exfiltration. In general, there are two ways for an attacker to leverage malware to accomplish this. The first is programming the malware to send the data out of the victim's environment. This method was more prominent in last year's report (79% compared to this one (43%). This can be accomplished both in real-time (as soon as the data is captured), or it can be done in batches at certain intervals or after certain actions (such as starting a program). It's quite common to see this functionality bundled with keyloggers, as shown in Appendix A.

The second method of exfiltrating data calls for the attacker to re-enter the network and retrieve it. If the attacker doesn't return via their original vector of intrusion, backdoors are a favored tool to remotely retrieve captured data. In addition, this type of malware also allows full control of compromised systems, which can be used to install additional malware, maintain persistence in the victim's environment, use the system to launch further attacks, and so on. Backdoors have consistently been one of the most common malware functions within our dataset for removing data from the victim's environments.

Figure 20. Malware functionality by percent of breaches within Malware

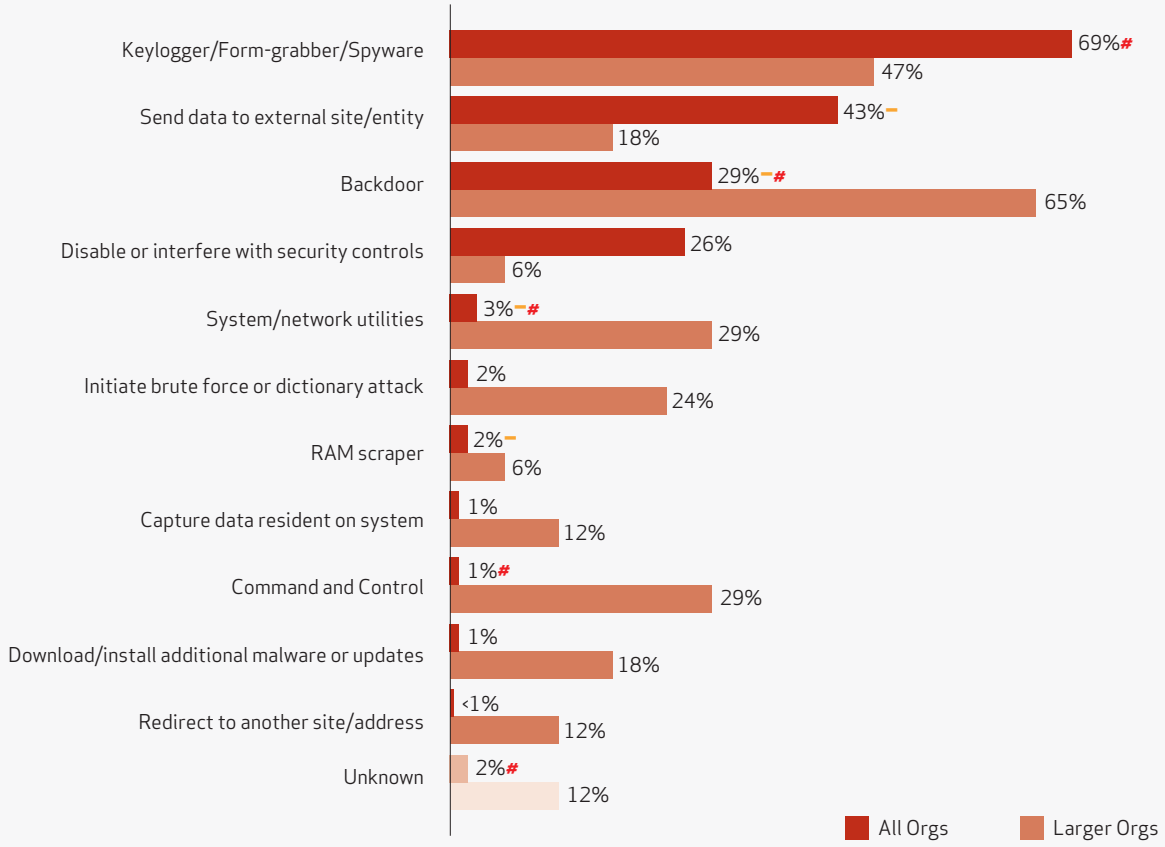

2011 data shows that the above-mentioned methodologies for exfiltration are differentiated based on organizational size. Sending data to an external site is more likely to be seen in malware that affects smaller organizations. This is because these functions are typically bundled with malware seen in large-scale automated attacks to which small and medium businesses are more prone. Large organizations, however, are more likely to see attackers utilize backdoors to exfiltrate data. Larger organizations typically have stronger authentication and access control in place for externally-facing remote services. Therefore, a popular way for an attacker to gain access into a larger organization's network is to get a foothold on an inside system through a backdoor and when they have a foothold the attacker installs multiple backdoors to maintain access over a long period of time.

Another difference seen within breaches of large organizations is the use of certain non-malicious network utilities, such as the SysInternals tools. Sometimes, system/network administrators use these tools for conducting normal maintenance on a system. However, if an attacker places them on a victim's system, we categorize them as malware. Many times antivirus products do not rate these as malicious.

As mentioned in previous reports, we highly encourage organizations to run systems in a least-privilege mode, monitor both ingress and egress points of their networks, and look for unauthorized changes to systems. Employing these practices can provide evidence of foul play and could potentially limit the impact of a data breach or any security incident. Additionally, look within your environment for indicators of compromise; seek out PSTools, growing archive files such as ZIPs and RARs, and new files with hidden and read-only attributes.

Although it's become passing fancy to kick dirt on the antivirus vendors, it's also important to investigate and follow-up on antivirus alerts. In many instances, the proverb "where there's smoke, there's fire" holds true. Detecting and responding to malware quickly may be the difference between a malware infection and a full-blown data breach. Multiple investigations have shown that incidents could have been stopped early if organizations followed up more promptly on antivirus alerts.

Malware Customization

This year about one-third of the malware investigated in the Verizon caseload was customized, which is quite a change from our last few reports. Malware customization is much more in line with statistics seen in our 2005-2007 caseload. Agents used customized malware slightly more against larger organizations, but not significantly so. When customization is used, in large organizations as well as in the overall dataset, it typically involves malware written from scratch by the attacker or involves the modification of existing code. We suspect that one of the main reasons for the change in malware customization is the usage of commercial applications within "industrialized" attacks. In these large-scale, multiple victim compromises, attackers simply don't need to bother with customizing malware since they can successfully use "canned" attacks against thousands of victims.

Hacking (81% of breaches, 99% of records)

Hacking is defined within the VERIS framework as all attempts to intentionally access or harm information assets without authorization or in excess of authorization by thwarting logical security mechanisms. Hacking is advantageous as an attack method for several reasons. It is usually conducted remotely, allowing attackers the benefits of anonymity and scalability. Automated tools and basic scripting, often written by someone else and made available to attackers, make many varieties of hacking extremely easy to conduct, and allow for attacks against multitudes of potential victims.

Classification note: There is an action category for Hacking and for Misuse. Both can utilize similar vectors and achieve similar results; in Misuse, the agent was granted access (and used it inappropriately), whereas with Hacking, access was obtained illegitimately.

This section will examine the hacking methods utilized in the combined 2011 dataset, as well as the vectors targeted in hacking actions.

Hacking Varieties

The graphic on the next page depicts the most commonly used varieties of hacking in our 2011 caseload. Those who follow this series may recognize some of these from previous reports. Like every year, a handful of techniques dominate the charts. Generally, the hit parade can be subdivided into the authentication attacks (stealing, brute forcing, or guessing of credentials), and technical attacks that bypass or break authentication altogether (e.g., SQL injection or backdoors).

Like every year, a handful of techniques dominate the charts. Generally, the hit parade can be subdivided into the authentication attacks, and technical attacks that bypass or break authentication altogether.

Also, the re-appearance of "mega-breaches" this year has shattered any correlation of breach and record loss percentages that may have existed. Some techniques, varieties such as abuse of functionality, were responsible for significant amounts of compromised records in one or two incidents. Interestingly, over half of all record loss had an unknown hacking method somewhere along the chain of events. This can be explained either by the lack of available logs, anti-forensic measures taken by the attackers, or the limitation in scope of the investigation based on client requirements.

Figure 21. Hacking methods by percent of breaches within Hacking

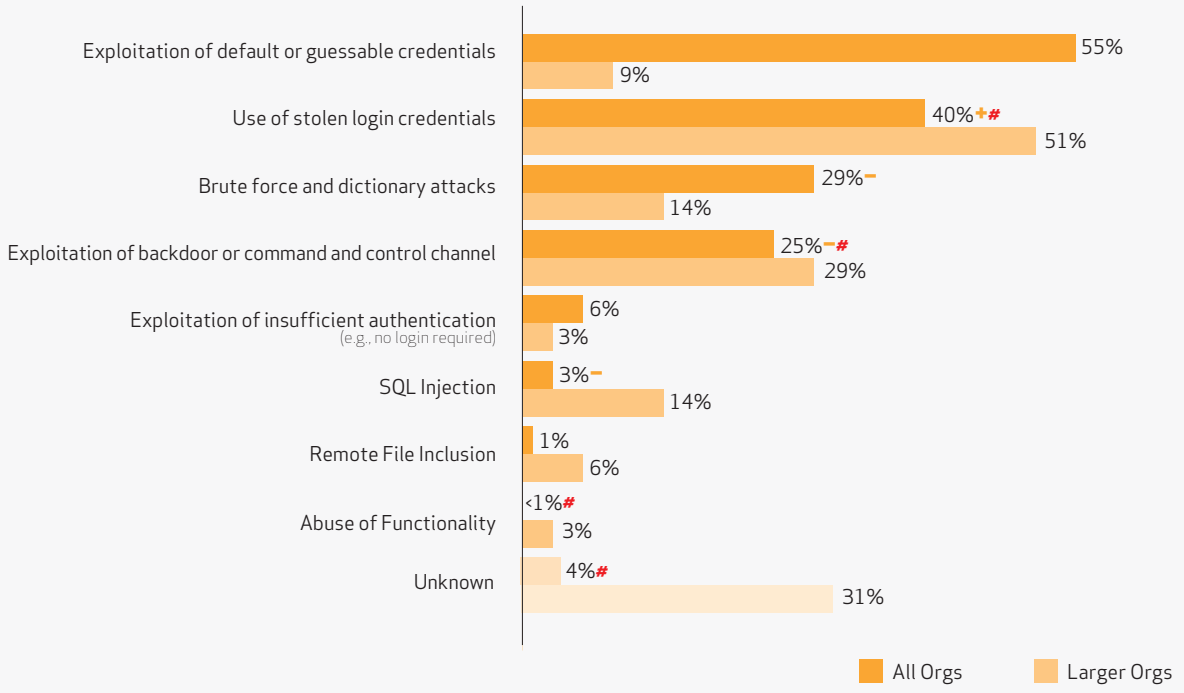

The low-level and highly automated attacks are still used on a wide scale but yield a relatively low number of records per attack. Typically these attacks involve smaller businesses mainly in the retail trade and hospitality sectors. Also, there are few obvious distinctions between the methods used for hacking into small companies compared to those utilized to compromise large ones. Larger companies do seem to be more adept at warding off the easier-to-prevent attacks; however, approximately 98% of all records breached via stolen credentials occurred in larger organizations.

As there's typically not a great deal of data to harvest within smaller companies, hackers revert to what could be considered "army ant" tactics. In other words, they strip the server clean of whatever data happens to reside there, but do not bother to make a home for themselves. Larger companies offer a lusher environment, and from the attacker's point of view, warrant a bit more investment. Like sedentary ants, attackers build tunnels and backdoors to easily and safely get to their pastures. As expected, such attacks require more technical proficiency—but typically don't fall into the "advanced" category of attacks—and yield more data in the long run.

Larger companies offer a lusher environment, and from the attacker's point of view, warrant a bit more investment. Like sedentary ants, attackers build tunnels and backdoors to easily and safely get to their pastures.

Both brute force attacks and exploitation of default or easily guessable credentials were down from last year's report, but still endemic in the retail and hospitality industries—typically smaller businesses. Unfortunately it's still possible to go to a vendor's site, get the client list and just hit those with the default or guessable username-password combination. These are relatively easy attacks that require little in-depth knowledge or creativity. They are usually scripted, aimed at many targets, and, if unsuccessful, exhibit little persistence.

In fact, the thief often doesn't even know what he's stolen until checking the remote server to which his scripts have been sending the captured data. The targets simply are not worth much effort to the attacker, since few records are stolen in such incidents. Scale of targets is what matters.

One particular case illustrates the lack of individual attention that goes along with most of these attacks. In this scenario, an online FTP server that had been misconfigured to allow anonymous FTP access was under constant attack by brute-forcing tools (like most online systems). What was noticeable for these scans was that many exotic username-password combinations were attempted, but very few tried to use anonymous access, which would have gotten them in.

Luckily, larger enterprises typically have basic authentication and access control configurations in place and enforce these policies on both their personnel and vendors. This becomes apparent when comparing the statistics for larger enterprises, where brute forcing is found in only 14% of cases (compared to 29% of all cases), and exploitation of default and guessable credentials is present in less than 10% of cases (<1% of data loss). We still see very few hacks exploiting patchable or 0-day vulnerabilities. As long as authentication is easy to fake or bypass, we don't expect this to change.

Hacking Vectors

The list of hacking-related pathways in Figure 22 tells a very similar story to years past. Remote access services (e.g., VNC, RDP) continue their rise in prevalence, accounting for 88% of all breaches leveraging hacking techniques—more than any other vector. Remote services accessible from the entire Internet, combined with default, weak, or stolen credentials continue to plague smaller retail and hospitality organizations. Often these victims share the same support and/or software vendor. Scripted attacks seeking victims with known remote access ports (TCP 3389, RDP or VNC), followed with issuance of known default vendor credentials, allow for targets of opportunity to be discovered and compromised in an automated and efficient manner.

Backdoors must be brought up yet again (we have reviewed the installation and utilization of backdoors earlier) as an attack vector. A quarter of all hacking cases feature a device that has a backdoor installed at some point in the event chain, typically post-compromise. Over 90 percent of known record loss is associated with attacks that exploit backdoors.

Web applications remain the third most common vector overall, and while they realized a decrease in the percentage of breaches, they were associated with over a third of total data loss. The inherent need for many web applications to be Internet-visible makes them a logical target; the potential to use them as an entry point into a corporate database makes them an attractive one.

When focusing on larger companies, there are some obvious shifts in the attack vector landscape. Most notable is that only 20% of the cases involve remote access services. Large organizations with more mature security practices have done a

Figure 22. Hacking vectors by percent of breaches within Hacking

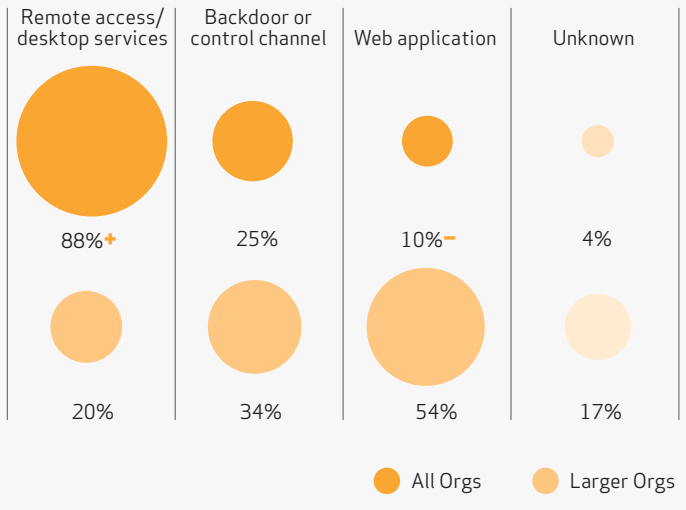

better job of limiting the accessibility of these services and have implemented better authentication controls. Backdoors are slightly more prominent in large organization breaches and almost all data loss can be associated with cases involving them. Web applications abound in many larger companies, and remain a popular (54% of breaches) and successful (39% of records) attack vector. As stated earlier, large companies have done a good job of limiting Internet visibility so attackers will either go after what is available (web applications) or attempt to bypass perimeter security with malware (backdoors, command and control).

Web applications abound in many larger companies, and remain a popular (54% of breaches) and successful (39% of records) attack vector.

Social (7% of breaches, 37% of records)

The centuries-old tradecraft of exploiting people in order to further criminal aims is alive and well in the information security field. The "carbon layer" of information assets (the user) is notoriously susceptible to social tactics such as deception, manipulation, and intimidation, and savvy threat agents know how to use this to their advantage. Even though humans are the most complex creatures on earth, criminals have consistently been able to outwit or otherwise leverage them in the process of stealing data. Although a little less common in 2011 (down from 11% in the previous year), the amount of data stolen during incidents employing social tactics rose dramatically from less than 1% to 37%—the highest in the DBIR's history (though it must be acknowledged that most of this ties to one of those "mega-breaches" we spoke of earlier).

Figure 23. Social tactics by percent of breaches within Social

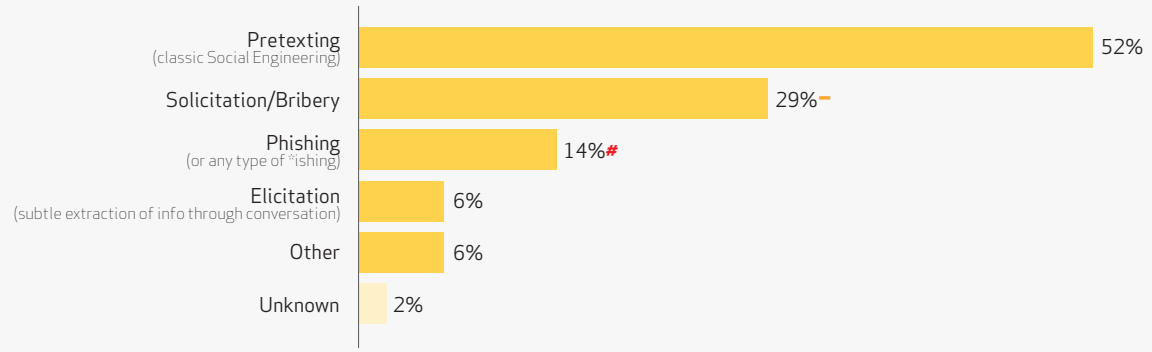

As depicted in Figure 23, pretexting (a.k.a. classic social engineering) moves to the top for the first time since we began this report. There are countless ways that imaginative and resourceful criminals can utilize pretexting, but the ploy most observed by us in 2011 targeted a string of small business owners over the phone by posing as support services for the payment card brands. Solicitation, on the other hand, was down markedly from 74% of all social tactics in 2010 to 29% in 2011. This is mostly due to fewer cases in 2011 involving organized, large-scale infiltration of banks in order to solicit employees to steal information. Quite a few arrests were made last year that busted up some of the more prolific crime rings behind such activities.

The "carbon layer" of information assets (the user) is notoriously susceptible to social tactics such as deception, manipulation, and intimidation, and savvy threat agents know how to use this to their advantage.

Phishing was relatively stable in the past year, edging ever so slightly up from 11% in 2010. It's quite a bit higher, however, when examining breaches affecting larger organizations. This is fairly interesting, and in line with many recent media reports detailing the use of malware-baited phishing lures cast toward some bigger and well-known enterprises. We believe this is a strategy designed to circumvent the typically more mature security measures in place at larger organizations. Why spend time searching for a way to exploit the specific technologies and weakness of a single company when every company contains people with the same basic vulnerabilities? You only need get one over-curious user to click, and we all know there's a clicker in every crowd (if you don't believe us, see exhaustive research here).

Figure 24. Social vectors by percent of breaches within Social

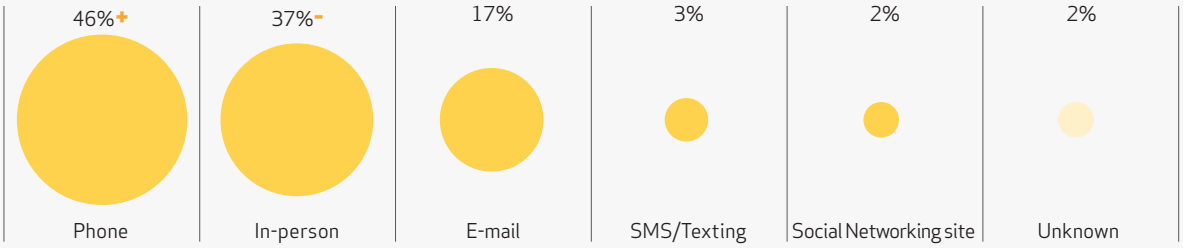

As seen in Figure 24, the top three vectors for conducting social attacks remain the same as in previous years, albeit in different order. The good ol' telephone served as the medium for nearly half of these cases. This coincides with the earlier observation of utilizing pretexting for extracting information from unsuspecting victims. Coming in second place were in-person tactics, typically used for soliciting waiters and cashiers within small businesses.

It should come as no surprise to learn that e-mail and phone share the top place as the most common vector of choice for orchestrating social tactics against larger organizations—this concurs with higher levels of pretexting and phishing against them. We observed attacks in 2011 that appeared to specifically target e-mail addresses at larger (or well-known) enterprises for use in subsequent attacks and targeted phishing campaigns. In some of these, the time separating the theft of e-mails and leveraging them in another attack was quite short. Law enforcement information shows some are offloaded to fill an order from a generic phisher, while some are used by the original thief for more directed or specific purposes.

Table 9. Social targets by percent of breaches within Social

Regular employee/end-user	43%
Cashier/Teller/Waiter	33%
Call center staff	8%
Finance/Accounting staff	6%-
System/network administrator	5%
Executive/Upper Management	3%
Helpdesk staff	3%
Customer (B2C)	2%
Human resources staff	2%-
Unknown	6%

Not much has changed this year with regard to the targets of social tactics as listed in Table 9. Unsurprisingly, regular employees and "money handlers" continue to be the main targets. While it's certainly a good strategy to use policies, procedures, training, and technical controls to protect human assets from falling for these schemes, there is no 100% effective patch for people. In one memorable case, we saw a phishing message successfully directed to a spam filter, only to later be retrieved and executed by an employee (but honestly—who can resist a simple click to see Nyan Cat streaking colorfully through space one last time?). Unfortunately, this particular ill-fated click kicked off a chain of

events leading to a substantial amount of data loss and quite a few upset customers. Scenarios such as these serve as reminders that humans can, in various ways, bring harm to the organization despite diligent efforts to control them. Nevertheless, doing nothing is an even worse control measure. We still favor well-thought out and well-delivered ways to increase awareness with regard to social tactics.

Misuse (5% of breaches, <1% of records)

Misuse is defined as the use of entrusted organizational resources or privileges for any purpose or in a manner contrary to that which was intended. These actions can be malicious or non-malicious in nature. This category is exclusive to parties that enjoy a degree of trust from the organization, such as insiders and partners. In 2009, Misuse was the most common of all threat actions, but dropped substantially in 2010 (48% to 17%). In 2011 the percentage of breaches involving misuse continued to decrease. For the entire caseload, misuse was found to be involved in only 5% of cases and was responsible for less than 1% of compromised records.

Although the percentage of cases declined substantially, the categories of misuse we did encounter are quite similar to what we have seen in past years. Embezzlement, skimming, and related fraud, for instance, topped the list with regard to frequency, appearing in 57% of incidents that involved misuse. Despite this relatively high percentage, the amount of records stolen in this manner was almost nil. Use of unapproved hardware and devices came in second at 50%. These two actions mostly consist of payment card handlers (cashiers, restaurant servers) using small handheld skimmers to capture data (though scenarios involving other kinds of rogue devices occur as well). Finally, abuse of system access and privileges came in third with 43% of all misuse. As you may recall, these have been the top three for the last two years, only in a slightly different order (e.g., last year: Embezzlement/Skimming—75%, Abuse of system access—49% and Use of unapproved hardware—39%). We anticipate that we will continue to witness these three varieties, as they correlate with financially motivated insiders with physical or logical access to valuable information.

Figure 25. Misuse varieties by percent of breaches within Misuse

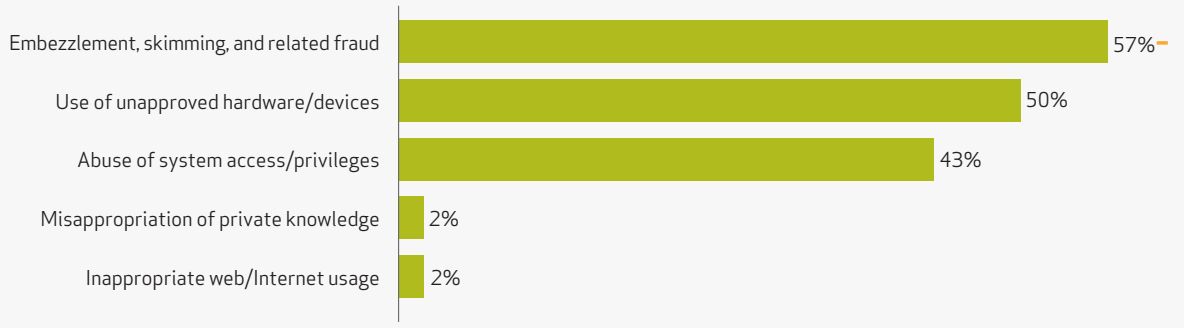

This year we began to record a new data point: the path or vector involved in misuse. While we realize that one year of tracking a few cases will not enable us to provide much in the way of conclusions, we hope that over time it will. Our aim is to ascertain where criminals are perpetrating their crimes so that corrective measures may be taken to prevent them from being successful. Based on our data, about half of all misuse was carried out through physical access within the corporate facility. This is intuitive, and is also closely related to the high amount of embezzlement and skimming crimes seen in our caseload. Interestingly, internal network access and remote access (e.g., VPN) show identical numbers as a vector of misuse (21%). It's interesting to note that these results do not seem to support the common concern that working from home is riskier than working from within the corporate perimeter. It will be interesting to see if the same results hold true in future samples.

In spite of all the warnings given by us and other industry insiders, we still see a number of organizations that fall victim to data theft through the malice of disgruntled ex-employees. In VERIS, these are classified as external agents (because they are no longer insiders) but when they are able to breach the victim because their privileges were never actually revoked, it still constitutes misuse. As we have said before, if you don't want these folks on your payroll or representing your company then DON'T LET THEM RETAIN ACCESS TO YOUR SYSTEMS.

Physical (10% of breaches, <1% of records)

The physical threat action category encompasses human-driven threats that employ physical actions and/or require physical proximity. Prior to the 2010 caseload, physical attacks were historically one of the least prevalent and least damaging threat actions. We attribute this to the nature of many physical actions, such as stolen user devices, which often result in data-at-risk rather than confirmed data compromise. Also, if physical access to a device is available as part of an insider's normal job duties, then it is classified under Misuse in VERIS rather than Physical (see the Misuse section and the discussion around embezzlement).

Beginning in the 2010 caseload, however, widespread ATM and gas pump skimming operations were responsible for a significant increase in physical attacks by percentage of breaches (29%) and compromised records (10%). Organized criminal groups usually conduct these operations, and one "spree" can hit numerous separate locations and/or businesses. These cases fall into a category that would not typically be pursued by Verizon investigators, but certainly fall under the jurisdiction of USSS and other law enforcement agencies.

There was no shortage of payment card skimming in 2011, and there were many notable arrests. "Operation Night Clone" is an example of how large and far-reaching skimming groups can become.

Accounting for only 10% of the combined caseload, the statistics this year show a decline in the percentage of incidents involving physical actions, including skimmers. The question becomes: "Was last year an outlier?" The answer: "Probably not." There are explanations behind the drop in the percentage of incidents and data loss associated with physical actions.

There was no shortage of payment card skimming in 2011, and there were many notable arrests. "Operation Night Clone" is an example of how large and far-reaching skimming groups can become. In July 2011, over 60 arrests in at least five countries were conducted as part of a multi-national law enforcement operation.[15] This year's caseload included many incidents, both U.S.-based and otherwise, that were carried out by organized skimming groups, and it's likely that the number of victims affected is higher than our totals reflect.

Also, difficulty in quantifying record loss associated with payment card skimming contributed to this decline since our last report. Skimming attacks do not target data in storage and when card numbers are totaled, it is often based on the amount used in fraudulent activity. This year featured more skimming cases where a confirmed incident was established, but the specific number of payment cards compromised was not known. However, the biggest reason for the lower proportion of data loss associated with physical attacks is simply the massive increase in the total amount of data stolen (by other threat actions) in the 2011 dataset.

Physical threat actions observed over the past year were limited to tampering, surveillance, or theft. Tampering is a physical action type defined as unauthorized altering or interfering with the normal state or operation of an asset. Tampering was found in every physical breach this year. Installation of skimming devices falls under this action and is responsible for the majority of tampering cases. Another form of tampering is operations where organized criminal groups swap legitimate PIN entry devices and Point of Sale (POS) terminals with "new" devices.

15 https://www.europol.europa.eu/content/press/major-international-network-payment-card-fraudsters-dismantled-1001

These devices are identical in appearance and designed to continue to perform their intended functions, but they are also redesigned to capture payment card data. They discreetly collect input from the swipe reader and/or the PIN entry keypad. Next to tampering, surveillance was the second-most common action. While the definition is much broader, all forms of surveillance observed in this dataset were associated with the installation of pinhole cameras to capture and transmit users' entry of their PINs on ATM machines. Used in conjunction with skimming devices, they are present in 35% of physical incidents, double the previous year (17%). In this context, theft (3%) refers to physically taking a device. Obviously, everything in this report falls under the label of "theft," but this specifically addresses forms of physical theft. All physical actions took place in public indoor or outdoor areas where, as one would expect, most gas pumps and ATMs are located.

Error (<1% of breaches, <1% of records)

Error is likely the most difficult threat category for which to provide hard guidelines when classifying its security incidents using VERIS. If our application of error is too liberal, then every incident contains some type of error. If we're too strict, then we'll miss identifying possible trends and opportunities in layer eight technology (that's the human layer, for you strict seven-layer OSI model adherents). We define an error rather broadly, as anything done (or left undone) incorrectly or inadvertently. That's the easy part. It gets a little trickier in practice when you're trying to achieve consistency among numerous data sharers who may all have different notions of when to include error in the chain of events.

> We're getting better at defining what is considered a formal error within VERIS. Take this statement for example: *"We have a procedure for locking down access to the server by source IP address of our third-party vendor and changing the default password, but in this instance we accidentally let this system slip through the cracks."* That would be labeled as a contributory error. Conversely, if the explanation was, *"So, you're telling me we're responsible for changing passwords and restricting by IP—wait, what?"* That would conjure a different and unofficial label in our minds, but we would not label it error.

Over time, we've taken an increasingly narrow interpretation on recording error when classifying an incident. Instead of trying to judge whether every bad decision, poor practice, or mental lapse deserves the label of "error," we only focus on critical errors that were either the primary cause of the incident or a significant contributing factor to it (in other words, error as a threat rather than a vulnerability). Furthermore, it must be something that represents a deviation from the normal operations of the organization. If their standard way of doing things offends all of our security sensibilities, is that an "error" or just systemic bad security?

Here's the important point from all of the above—only breaches in which the error was a primary cause of disclosure are included in the statistics displayed in this report.

This year, only four incidents out of 855 were associated with error as a primary cause, all of which involved disclosure of sensitive data on servers due to publication or misconfiguration errors. Another 12 incidents were flagged as including contributory errors. If you were to compare this to last year's report, you would see it is a far cry from the 200+ contributory errors we listed there. Unfortunately, this drop has nothing to do with improvements in security processes and procedures by last year's offenders, and everything to do with removing the label of error from "default credentials by default" and other insecure standard operating procedures.

Instead of trying to judge whether every bad decision, poor practice, or mental lapse deserves the label of "error," we only focus on critical errors that were either the primary cause of the incident or a significant contributing factor to it (in other words, error as a threat rather than a vulnerability).

Environmental (0% of breaches, 0% of records)

The environmental category not only includes natural events such as earthquakes and floods, but also hazards associated with the immediate environment or infrastructure in which assets are located. The latter encompasses power failures, electrical interference, pipe leaks, and atmospheric conditions. As in the last few years, environmental hazards did not factor into any breach scenarios across our combined dataset. There is, however, a public example of an environmental action that contributed to both the loss of confidentiality and availability this past year.

Tragically, in May 2011, an EF5 tornado stuck Joplin, MO.[16] Included in the tornado's path of destruction was St. John's Regional Medical Center, which was heavily damaged. The twister scattered debris—including medical records and X-rays from the hospital—across a wide geographical area, some of it found as far as 60 miles away.[17] The hospital's strategy to move to Electronic Health Records with a backup site outside of their geographical area appears to have paid large dividends with regard to their recovery.

Compromised Assets

This section offers insight into the types of information assets compromised by breaches during 2011. This is an important piece of the A^4 model because it focuses on the target of the agents and actions already discussed in this report. It turns the conversation from an external "Who are they and what are they doing?" to a more personal "What do we have that needs protecting and what harm might arise if we fail?"

Figure 26. Categories of compromised assets by percent of breaches and percent of records

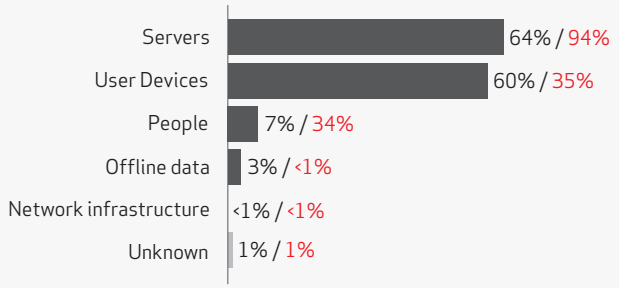

Figure 27. Categories of compromised assets by percent of breaches and percent of records – LARGER ORGS

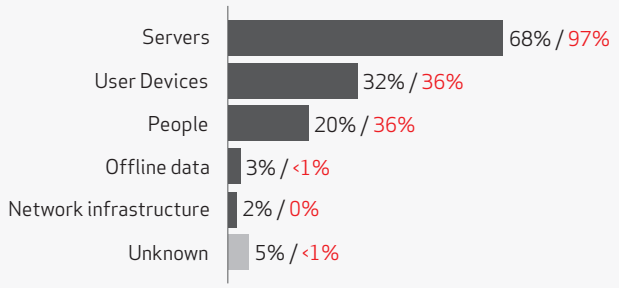

We know the question burning in your mind since the last DBIR: "Will 2011 be the year that user devices finally overtake servers as the most oft-compromised asset category?" Well, the votes are in, and the Pwny award goes to… servers! They managed to preserve their undefeated streak by a slim 4% margin of victory (2010 was an even slimmer 1% in case you've forgotten). One thing to keep in mind with this narrow victory is the three user devices that ranked the highest: POS terminals (35%), Desktops (18%), and ATMs (8%). Laptops were right down the bottom at 1%. Therefore desktops and laptops make up only 19% of compromised assets. Like all statistics, it depends on how you split them as to the story you are telling. We're sorry to disappoint all you fans of user devices out there (aren't we all?), but cheer up; it's been a good run, and there's always next year.

Getting a tad more serious, it's really not surprising that servers seem to have a lock on first place when it comes to the types of assets impacted by data breaches.

16 http://en.wikipedia.org/wiki/2011_Joplin_tornado
17 http://www.news-leader.com/article/20110523/NEWS11/110523001/Missouri-officials-say-tornado-killed-least-89

Figure 28. Percent of records compromised in breaches involving servers

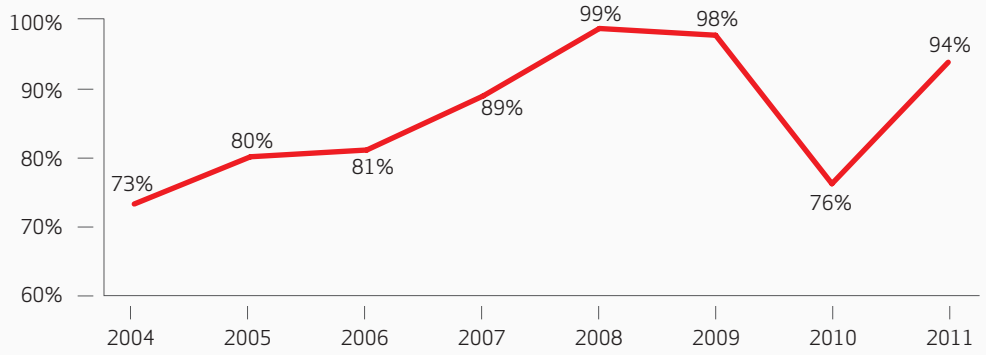

They store and process gobs (technical term) of data, and that fact isn't lost on data thieves. This point is even clearer when one considers the total amount of data stolen when servers are involved versus other asset categories. The running comparison is shown in Figure 28. If you're wondering about the dip in 2010, we attribute this to the lack of "mega-breaches" during that year (a topic we discussed extensively in the 2011 DBIR). Almost all incidents in which very large amounts of data are compromised involve servers. Since we observed no such incidents in 2010, data loss was a little more distributed. As a testament to this, Figure 35 in the Compromised Data section reveals a remarkably similar pattern.

We all know, of course, that user devices store and process information too. Furthermore, most organizations have a lot more of them than they do servers, and they're often widely distributed, highly mobile, less restricted, and—perhaps more importantly—controlled by end users (a shudder travels down the spine of all the admins out there). For all of these reasons and more, user devices frequently factor into data breaches in some manner or another and contribute to a hefty chunk of overall data loss. Sometimes they are the endpoint from which data is taken, but more often they simply provide an initial "foothold" into the organization, from which the intruder stages the rest of their attack. A common scenario—especially for larger organizations—involves the installation of a keylogger on a workstation or laptop in order to steal the user's username/password for an internal application server.

We can't leave the topic of asset categories without a mention of people, offline data, and network infrastructure. Some ask us why "people" are considered an information asset, and we usually reply by asking how they manage productivity with a workforce of brainless zombies that

Table 10. Compromised assets by percent of breaches and percent of records*

Type	Category	All Orgs		Larger Orgs	
POS server (store controller)	Servers	50%	1%	2%	<1%
POS terminal	User devices	35%	<1%	2%	<1%
Desktop/Workstation	User devices	18%	34%	12%	36%
Automated Teller Machine (ATM)	User devices	8%	<1%	13%	<1%
Web/application server	Servers	6%	80%	33%	82%
Database server	Servers	6%	96%	33%	98%
Regular employee/end-user	People	3%	1%	5%	<1%
Mail server	Servers	3%	2%	10%	2%
Payment card (credit, debit, etc.)	Offline data	3%	<1%	0%	<1%
Cashier/Teller/Waiter	People	2%	<1%	2%	<1%
Pay at the Pump terminal	User devices	2%	<1%	0%	<1%
File server	Servers	1%	<1%	5%	<1%
Laptop/Netbook	User devices	1%	<1%	5%	<1%
Remote Access server	Servers	1%	<1%	7%	<1%
Call Center Staff	People	1%	<1%	7%	<1%

*Assets involved in less than 1% of breaches are not shown

don't store and process information. Because people know things that should not be disclosed and do things (e.g., to other people or assets) that could contribute to disclosure (or some other form of harm), they must be protected. It's worth noting from Figure 27 that people contribute to 20% of incidents and 36% of data loss in larger organizations. Moving on, the percentage of breaches involving offline data declined for the third year in a row (from 25% to 12% to 3%). We believe this relates to another trend during that timeframe in which insider misuse has also seen a decline; such activities often targeted documents, media, and other things scattered about the workplace. Attributing some of this drop to the "green initiative" is probably a stretch, but it does give us opportunity to contribute to a healthier, happier planet by reminding everyone that less printing leads to fewer stolen documents (which, of course, begs the question of whether databases are more secure than paper, but never mind that). Finally, the following explanation should clarify why "networks" are tied to so few breaches. It's absolutely accurate that a huge proportion of breaches infiltrate networks. However, we're specifically referring to network infrastructure devices like routers, switches, cabling, etc. Traffic might pass through these devices during a breach, but the devices themselves are not compromised.

Table 10 gives a more detailed accounting of the kinds of assets compromised across the 850+ breaches we studied. Calling out larger organizations from the rest of the dataset is particularly helpful here. POS servers and terminals account for a large share of incidents across the whole dataset, but are basically a non-factor for larger organizations. Intuitively, web applications and databases play a much larger role at the enterprise level.

Some readers may find the lack of mobile devices like tablets and smartphones in Table 10 surprising. We confess that we also expected a stronger trend to emerge, but so far it has not. To be clear, we have conducted forensic engagements on mobile devices (for malware, suspected misuse, tampering, etc.), but confirmed data compromises remain rare. We can't help but think, however, that given the explosion of mobile users, applications, payments, etc., things may pick up in the future.

Because working definitions of "the cloud" are legion, it can be difficult to answer questions about how this paradigm factors into data breaches.

Ownership, Hosting, and Management

This sub-section touches on two very hot topics in IT and security circles today: cloud computing and the bring-your-own-device (BYOD) movement. The latter is what it sounds like: a model of allowing employees to use their personal devices for work purposes. This notion is both scary and attractive at the same time. It could boost employee happiness while saving money and administrative overhead, but it also exposes enterprise data and systems to devices that may not conform to enterprise security policies (and that's being gracious). Due to the buzz around BYOD in 2011, we've added a datapoint in VERIS regarding the ownership of assets involved in data breaches. Our findings are shown in Figure 31.

Because working definitions of "the cloud" are legion, it can be difficult to answer questions about how this paradigm factors into data breaches. Do we see breaches that compromise assets in an externally-hosted environment that is not managed by the victim? Yes; absolutely. Do we see successful attacks against the hypervisor in the wild? No; not really. We've said it before, and we'll say it again here: it's really more about giving up control of your assets and data (and not controlling the associated risk) than any technology specific to the cloud.

That a question is difficult to answer is not an excuse to avoid asking it or attempting to answer it. In thinking about how to study the relationship between data breaches and the cloud, our approach has been to identify traits inherent to the cloud, and then to observe how those traits coincide with breaches. Some will consider it over-simplistic, but ownership, hosting, and management are three important characteristics of cloud-based assets. An asset that is "in the cloud" is often one that you do not own, is hosted in an external facility, and is managed or administered by a third party (or some combination of these).

Figure 29. Hosting of assets by percent of breaches*

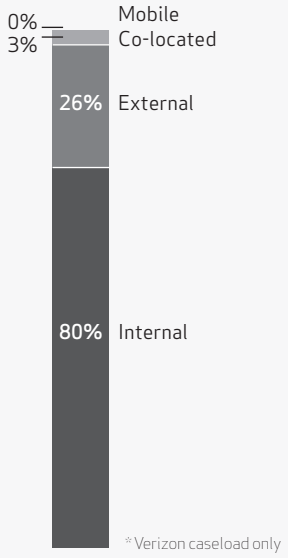

Figure 30. Management of assets by percent of breaches*

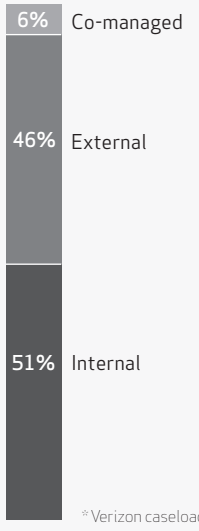

Figure 31. Ownership of assets by percent of breaches*

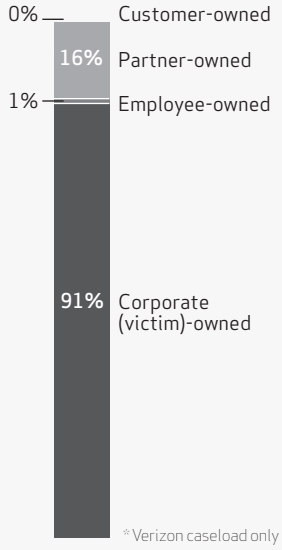

We all know that correlation is not causation, but we also know that patterns can be pertinent. With this in mind, we turn our attention to Figures 29, 30, and 31, which show statistics around ownership, hosting, and management of assets across our caseload. Whether these particular patterns are relevant is unclear from the data before us, but we can make some observations. For one, this marks the third year in which we've seen an increase in the proportion of externally-hosted and managed assets involved in data breaches. That may be important, especially if you cross-reference this with issues we bring up every year regarding third-party POS vendors that can't seem to figure out how to configure an access control list or change default credentials. It's also worth nothing that assets within larger organizations were more likely to be internally hosted and managed (around the 80% mark for both). Employee-owned devices were rare, but if the BYOD movement catches on, we should see that begin to rise.

Compromised Data

Over the years, we've received feedback questioning the value of counting the number of records compromised in this report. After all, not all records are created equal. Should a record containing a name and e-mail address be counted on the same scale as the source code to a flagship product? This is a valid concern and we encourage the reader to consider this when promoting or using this number in derivative works. Nevertheless, there is value in tracking the number. First, it is a number required by some third parties (e.g., payment card brands) and counts help keep perspective within the specific record types. Plus, we do suspect there is a relationship between number of records lost in a breach and the impact of the breach (at least for certain data varieties), and we won't be able to know if we don't collect the data and ask the questions. Finally, it acts not only as a kind of "wet finger in the wind" measurement, but the variety and amount of data loss also contributes to our understanding of motivation and activities of the attacker.

So, long story short, the "mega-breach" is back after a one-year hiatus, albeit with new types of threat agents and motivations associated with them. In lieu of focusing first on the really big scary number of compromised records (174 million, by the way), we will start by comparing the number of incidents and amount of losses associated with each variety of data. The distribution of record loss, the inherent imperfections around quantification of compromised records, and comparisons to prior years are also reviewed in this section.

Varieties of data compromised

Payment card information was again involved in more breaches (48%) than any other data type, but unlike previous years, it was not a runaway winner. Authentication credentials were a close second (42%) in the current dataset. This may allow payment cards to retain the title of "most stolen" but the title of "largest hauls" now belongs to the personal information variety, which includes name, e-mail, national IDs, etc. While only 4% of breaches included the loss of personal information, it comprised 95% of the records lost in this report. This is an enormous change from previous years; 4% of record loss in 2009 involved personal information and 2010 showed only 1%.

Table 11. Varieties of data compromised by percent of breaches and records

		All Orgs		Larger Orgs	
Variety	Label in Fig 31	Breaches	Records	Breaches	Records
Payment card numbers/data	CardData	48%	3%	33%	1%
Authentication credentials (usernames, pwds, etc.)	Credentials	42%	1%	35%	1%
Personal information (Name, SS#, Addr, etc.)	Personal	4%	95%	27%	98%
Sensitive organizational data (reports, plans, etc.)	OrgData	2%	<1%	22%	<1%
Bank account numbers/data	BankData	2%	1%	10%	1%
System information (config, svcs, sw, etc.)	SysInfo	2%	<1%	15%	<1%
Copyrighted/Trademarked material	Copyright	1%	1%	3%	<1%
Trade secrets	TradeSecret	1%	<1%	12%	<1%
Classified information	Classified	<1%	<1%	2%	<1%
Medical records	Medical	<1%	<1%	2%	<1%
Unknown (specific type is not known)	Unknown	44%	<1%	2%	<1%

It is intriguing that personal information was stolen en masse by both activist and organized criminal groups in 2011, but for very different reasons. The former seek this kind of data from victims because it makes a statement; it's embarrassing and it's…personal. The scale of the breach, sensitivity of the data, and the manner in which it is released are far more important to their goal than the actual value of the stolen data.

Figure 32. Varieties of data compromised by number of breaches and records

Figure 33. Role of organization size on variety of data compromise

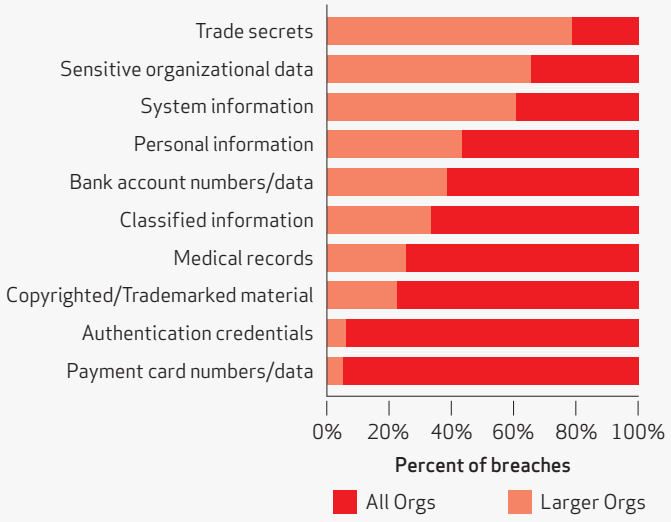

Organized criminals, on the other hand, are very concerned about the value of the data. The question, then, is: why would they want to steal a boatload of personal information, and what would they do with it once they have it? The answer lies not in the value of the information itself, but in its value for fraud and other schemes.

We observed criminals using personal information to launch phishing campaigns, commit identify theft, and engage in all manner of fraud. It might be a slightly different path, but in the end it still leads to money.

Another item of interest from the view in Table 11 is the number of "unknowns". In almost half of all breaches, it was clear that data left the premises, but it was not clear (or only partially clear) exactly what left. We will address this finding in more detail later in the section.

"[…] the criminal community has effectively been deterred from engaging in high-profile activity. Pulling off a huge heist might achieve fame and fortune, but it also attracts a lot of unwanted attention."

-DBIR 2011

Figure 33 compares data types compromised from larger versus smaller organizations. There are two glaring discussion points from this figure, each on opposite ends. We'll start at the bottom, where nearly all payment card breaches are shown to affect small businesses. This continues the trend established in last year's study, where the bulk of criminal activity targeting payment cards has shifted away from larger organizations to smaller ones, primarily because they can be obtained at a lower risk. At top of Figure 33, it is apparent that larger organizations were on the losing end in the majority of thefts involving trade secrets and other sensitive organizational data.

Distribution of Record Loss

Table 12 shows the amount of records lost per incident. While, like last year, we have a large share of smaller breaches, keeping the general shape in mind is important if probable losses are being estimated or modeled. Breaches with record losses in the hundreds or thousands are more likely than the headline-grabbing incident involving millions of records—though don't forget, we have a handful of those in here as well.

Over three quarters of breaches involved losses of less than 10,000 records and only 7 breaches out of the 855 lost more than one million

Table 12. Distribution of records compromised by percent of breaches

	All Orgs	Larger Orgs
<1k	39%	62%
1k-10k	44%	15%
10k-100k	11%	5%
100k-1m	3%	5%
>1M	2%	13%

records.[18] Yes, the mega-breach is back and they are occurring, as one might predict, in larger organizations. It should be no surprise that large organizations with more data have bigger breaches. In this year's data, organizations with over 1000 employees account for all breaches of more than 10 million records and over two-thirds of all breaches of more than one million records.

Figure 34. Unknown variety and amount of dataloss by percent of breaches

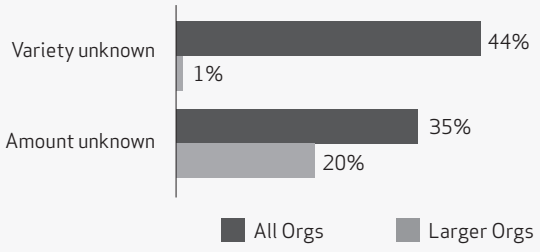

Quantifying Record Loss

The discussion of total records compromised is an attempt to step toward a quantitative assessment of the overall consequences of security incidents. There are a number of disclaimers required when having this discussion. First, "total records compromised" is only one metric in the measurement of the overall impact of a security incident. The differences in data types and their value to an organization were addressed earlier, but the point is worth repeating.

The second disclaimer is that, in some cases, it's pretty dang hard to do.

The struggle to count stolen records has several causes, some as a result of attack methods, some based on the data type involved, some due to lack of log collection by the victim. We (and our contributing partners) had many cases in which the existence of a breach was clear and evidence was found of data being exfiltrated (e.g., finding an encrypted payload file), but determining the record type and quantity proved difficult or impossible. Another common example is with organizations that suffer payment card breaches, but only the cards used in fraudulent activity can be confirmed and the total number of payment cards stolen remains unknown. Some types of information, such as trade secrets, are inherently difficult to quantify. Breaches of certain data varieties are more difficult to discover, as they don't come with the luxury of fraud detection to provide the victims with notification. Moreover, our law enforcement partners focus more on tracking and capturing the attacker than they do on counting an exact number of records lost with each and every victim. One other item to note, all but one of the unknown record types occurred in the smaller organizations (less than 1000 employees).

Now that we've looked at the gory details for last year, how does 2011 stack up against record loss over the previous seven years?

Figure 35 shows what looks to be the attackers regressing to being *mean* (see what we did there?). While 2011 featured a significant bounce back in record loss, the second-highest total in our study, we are still left having to explain the dip in 2010. We provided several different speculations and we feel that they still hold true; the leading theory remains that money-motivated organized crime is opting for more automated strikes as opposed to a single high-profile heist. We saw an increase in smaller breaches,

Table 13. Descriptive Statistics on records compromised, 2004-2011

	2011	All-Time ('04-'11)
Total Records	174,522,698	1,081,495,508
Mean	548,814	801,108
Median	1,270	1,270
Standard deviation	5,632,140	8,221,338
Percentiles		
10th	12	10
25th	100	45
50th	1,270	1,270
75th	3,500	10,559
90th	35,219	153,978
99th	2,830,000	10,000,001

18 RISK Team alumnus Alex Hutton is fond of referring to these less frequent but large incidents as "Black Swans" even though we've oft corrected him that this is an incorrect application of Talib's theory (this isn't actually true, but we couldn't resist a jab at him—ask him about it sometime).

a trend that continued into this year, perhaps indicating a preference for a steady trickle of records from many sources over sporadic fire hose blasts from a few. All but one of the large breaches (over 1M records) this year were attributed to activist groups rather than financially-motivated agents. This supports the theory that organized crime is quietly targeting the softer, smaller targets while activists are waging more of a "shock and awe" campaign.

Figure 35. Number of records compromised per year across combined DBIR dataset

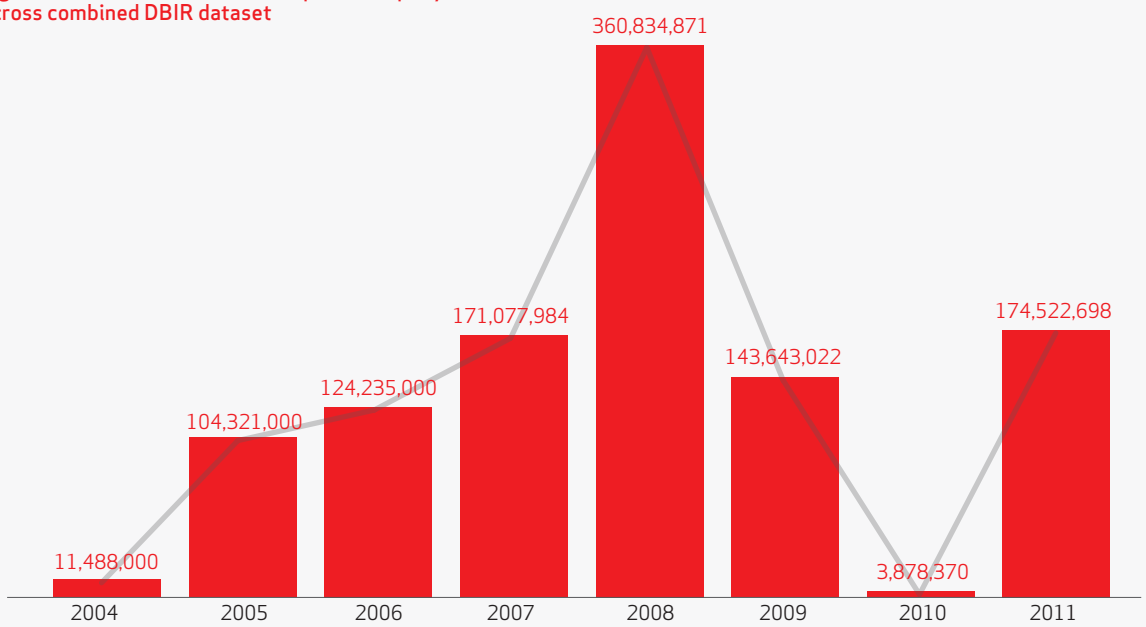

For the more statistically-minded reader (we know you're out there), we're again including Table 13. The mean records compromised this year fell to just under 550,000; the median rose to 1,270 records and the standard deviation is 5.6 million records. With the addition of this year's report we've surpassed **1 billion records** compromised since we've been recording breach data. A milestone, yes—but not one we're celebrating by breaking out the bubbly. We will, however, raise a toast to never seeing 1.1 billion.

Attack Difficulty

The best breach is the one that never occurs in the first place. Preventing a breach during the initial compromise would obviously have the least amount of repercussion. However, security controls are neither absolute nor binary; they range from quick-fixes all the way up to complex-and-costly. Understanding how much to allocate requires an understanding of the pressure the attacker is applying and where, during the attack, more pressure is used.

To explore this notion in a bit more detail this year, we opted to alter our usual assessment of the difficulty of the overall attack. We've provided separate ratings for the initial compromise and the subsequent actions that follow it. The *initial compromise* refers to any actions leading up to the attacker gaining unauthorized access to an asset, whereas *subsequent actions* covers anything done after that, including the compromise and exfiltration of data. We've often said the latter phase is usually more sophisticated, but we wanted to put our numbers where our mouths are. Before we look at those numbers, however, let's get the caveats and definitions out of the way.

Figure 36. Difficulty of initial compromise by percent of breaches and percent of records*

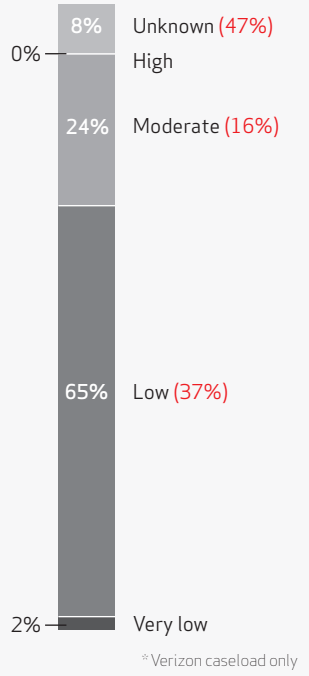

Figure 37. Difficulty of subsequent actions by percent of breaches and percent of records*

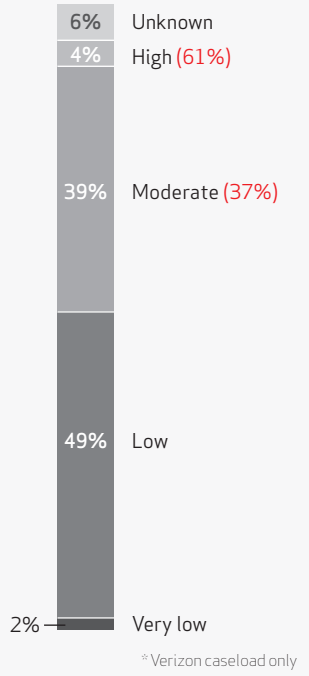

Rating the relative difficulty of the attacks we investigate admittedly involves some degree of subjectivity. However, in spite of this we find it is still useful as an indicator of the level of effort and expense required to breach corporate assets. Our investigators[19] assess the various details involved and then classify them 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.

With these ratings in mind, the reader can examine Figures 36 and 37. It is apparent that, as suggested earlier, the initial compromise tends to be easier to pull off than the stuff that happens afterwards. Gaining access is at least moderately difficult in just under a quarter of incidents, but close to double that for subsequent activities. By the way, don't let the "unknowns" be a distraction; sometimes logs are sparse and the exact techniques utilized simply aren't clear enough to assess their difficulty.

If one refers back to the list of common hacking varieties in Figure 21, this finding makes a lot of sense. Many of the methods listed there—especially those tied to first gaining access—are not particularly sophisticated. What occurs after that initial entry in a typical data breach scenario is where things usually get a bit more interesting. The assailant will often install various types of malware (which may have some degree of customization), escalate privileges, setup remote entry and control mechanisms, and explore the victim's network to find where the goodies are stored. And find them they do. As depicted in Figure 37, an astonishing 99% of all data was compromised via moderate to difficult post-infiltration methods. This statistic underlines our original assertion of the critical importance of preventing access in the first place.

So, what about larger organizations? Surely they're a lot more difficult to infiltrate, right? Sadly, our data seems to suggest otherwise; it does not appear that cybercriminals have to work much harder to compromise larger organizations than they do for smaller ones.

So, what about larger organizations? Surely they're a lot more difficult to infiltrate, right? Sadly, our data seems to suggest otherwise; it does not appear that cybercriminals have to work much harder to compromise larger organizations than they do for smaller ones.

19 Attack difficulty is not a part of the VERIS framework, and therefore, is not a data point collected by organizations partnering with us for the report. As a result, statistics in this section pertain only to Verizon's 2011 caseload.

Subsequent actions, however, did slide a bit more toward the difficult end of the spectrum, with about three of four in the moderate to high difficulty rating. This is probably the result of a more diverse range of assets, better internal defenses, and (typically) more mature detection and response capabilities.

As stated at the beginning of this section, rating the difficulty of attacks does involve a degree of subjectivity. However, it still serves as a useful indicator of threat agent capabilities and what must be done to protect against their attacks. The most effective and efficient approach is almost always to stop assailants before they get in the door. Most opportunistic criminals (and in the next section we'll see that most breaches are opportunistic) will not expend their resources on a hardened target while a softer one of similar perceived value is available. If information is highly desired—or targeted for other reasons—covering the basics, while still essential, may not be enough. Thus, understanding the motives, skills, and methods of our adversaries is important to any well-considered and well-prepared defense.

Attack Targeting

It's commonplace in the information security industry to classify attacks into two broad categories: opportunistic and targeted. The differences between the two categories are vast, but let's begin with a description of the terms:

- **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.

As best we can tell, becoming a target of choice is just one of those things comes with the territory of having a higher public profile and/or being known to possess valuable assets.

The ratio of opportunistic to targeted attacks in 2011, as seen in Figure 38, is very similar to what we reported the year before. However, this time we discovered two interesting facts related to opportunistic attacks: 1) 85% of targets of opportunity are organizations with fewer than 1000 employees, and 2) nearly three-quarters of them hit the Retail/Trade and Accommodation/Food Services industries. These observations would seem to support the conclusion we've drawn elsewhere in this report, namely, that large-scale automated attacks are opportunistically attacking small to medium businesses, and POS systems frequently provide the opportunity.

Targeted attacks, on the other hand, paint a very different picture. Finance/Insurance and the Information sectors are targets of choice more often than other industries, and larger businesses are also more likely to be singled out for attack. In fact, almost seven out of 10 targeted attacks leading to data breaches were against larger organizations. As Figure 39 confirms, they exhibit nearly twice as many targeted attacks as do those more opportunistic in nature.

Figure 38. Attack targeting by percent of breaches and percent of records*

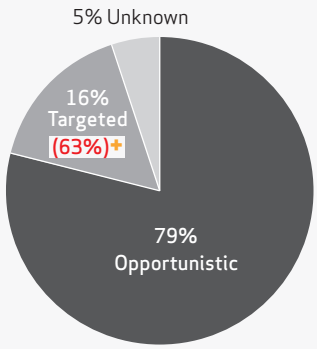

* Verizon caseload only

Figure 39. Attack targeting by percent of breaches and percent of records – LARGER ORGS*

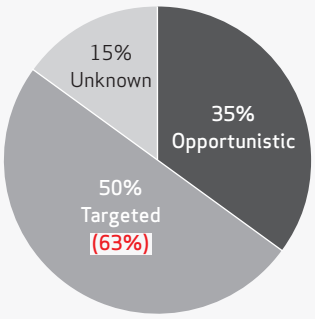

* Verizon caseload only

Unfortunately, the data doesn't have much else to say on this topic, and our crystal ball was in the shop this week. We can't be exactly certain why this disparity between smaller and larger organizations exists or whether it will continue in the future. As best we can tell, becoming a target of choice is just one of those things comes with the territory of having a higher public profile and/or being known to possess valuable assets. This is certainly the case for attracting the aggression of activist groups, as discussed earlier in this report. On the other end of the scale, smaller organizations often do not have the knowledge or resources necessary to address flagrant weaknesses in their Internet-accessible assets that cause them to be identified for opportunistic attacks. While larger organizations are typically better equipped to deal with security issues, they also have more assets, more distribution, and more complexity to manage. Thus, they are not immune to opportunistic attacks; organizations both large and small present far too many tempting opportunities to would-be assailants and inevitably pay the consequences.

Overall, the greater degree of targeting against bigger enterprises and the Financial and Information industries means these organizations are better able to defend against opportunistic attacks or they are more often singled out for attack, or perhaps a mixture of both. In any event, we believe that our rule of thumb still applies: 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).

A THREE-DAY WORK WEEK

During our data collection for this report, we received a list from one of our law enforcement partners containing the dates and locations of a large number of incidents tied to a small organized criminal group operating out of Eastern Europe. It provided us the opportunity to study their behaviors and activities over about a six-month period. We found it fascinating and include it here in the hopes that it helps drive home notions like "industrialized," "rapid," "large-scale," and "opportunistic," which we reference frequently in this report.

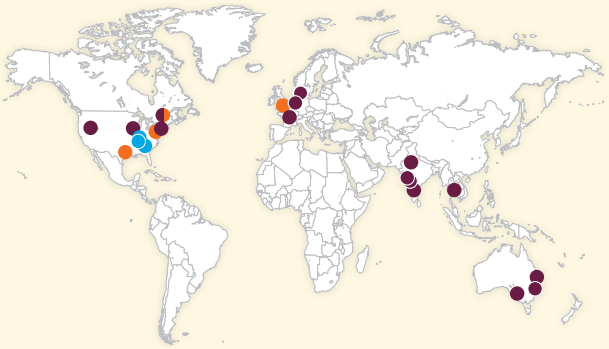

Saturday Sunday Monday

Analysis of the data showed the attackers not only had no routine workweek, but they only worked an average of three days a week. During one particular three-day work week, they punched the clock on Saturday, Sunday, and Monday. They compromised 22 organizations across nine countries; Monday was the most productive, with 15 confirmed breaches registered that day (in purple). We would joke about "nice work if you can get it" but the jail time these guys are facing doesn't make for very nice work at all.

Timespan of Events

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. To veteran incident responders it will come as no great surprise when we assert that the timespan of events in breach scenarios can vary greatly depending upon a multitude of factors. As with any other scientific endeavor, the longer one studies these factors, the more the model designed to capture them will evolve, along with our understanding of it. In past DBIRs, we separated events into three major phases, which we believed aligned quite well with the typical incident response process. Even so, we decided to give the timespans a bit of a tweak in this iteration. By splitting up the former "point-of-entry to compromise" phase, we hope to draw out a few more details around the initial compromise and data exfiltration events. Figure 40 shows the timespan phases along with the associated percentages. The corresponding findings are discussed in the paragraphs that follow.

Figure 40. Timespan of events by percent of breaches

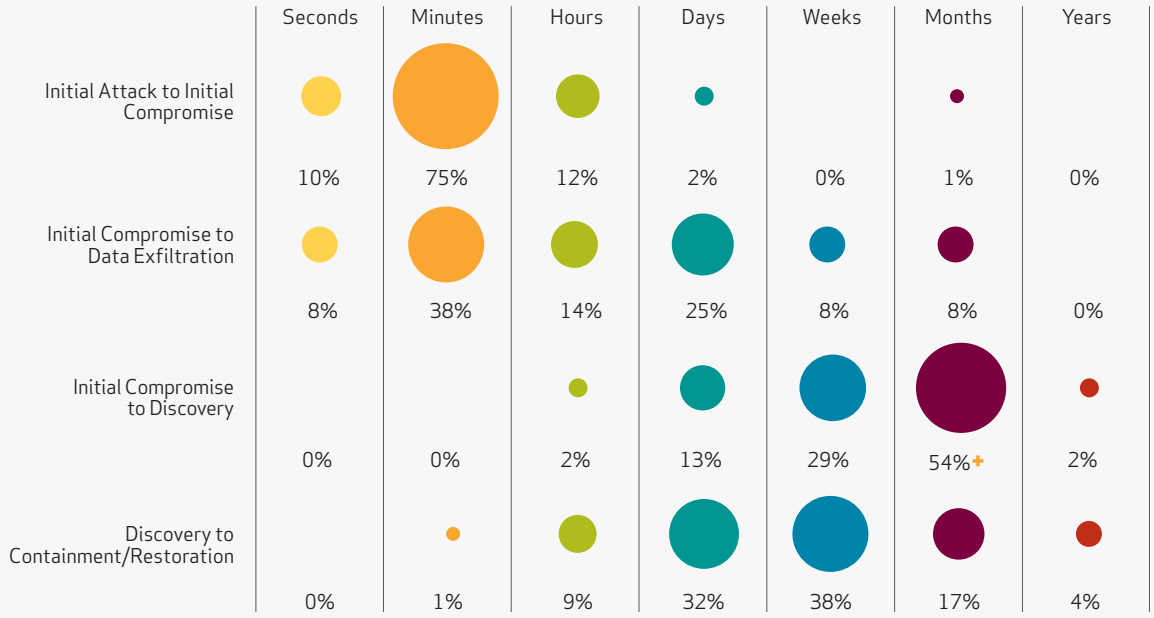

Initial Attack to Initial Compromise

This initial phase depicts the span of time from the first malicious action taken against the victim to the point at which a security attribute (e.g., confidentiality) of an information asset is negatively affected. This is not necessarily the culmination of the incident; it often represents the initial intrusion into the network. Please note that this phase is only applicable to deliberate and malicious actions by agents who *did not* already have access/privileges to the asset.

In the 2011 caseload, we observed that for the vast majority of incidents (85%), attackers are able to compromise the victim very quickly (minutes or faster). Keep in mind that the previously-published statistics for "point-of-entry to compromise" are not comparable here. Isolating the time it takes to gain unauthorized access highlights with startling clarity the speed at which a great many attacks occur. This result is largely (but not exclusively) the byproduct of the many automated, quick attacks against smaller organizations in the 2011 caseload. It just doesn't take that long to pop a POS system using a scripted list of known usernames and passwords.

Figure 41. Time between initial attack and initial compromise – LARGER ORGS

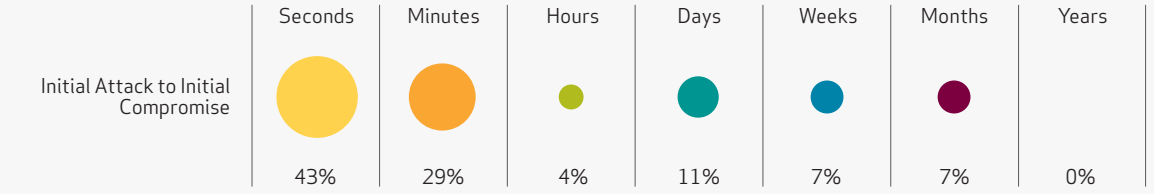

Lest larger organizations read this and believe themselves immune to such rapid compromises, they should take heed of the following statistic. A not-much-better 71% of them fell within minutes. Granted, the actions that felled them were often different (e.g., SQL injection), but does that really matter when the result is the same?

Lest larger organizations read this and believe themselves immune to such rapid compromises, they should take heed of the following statistic. A not-much-better 71% of them fell within minutes. Granted, the actions that felled them were often different (e.g., SQL injection), but does that really matter when the result is the same?

Initial Compromise to Data Exfiltration

The second phase covers the time period from when the first asset is negatively affected to the point when non-public data is removed from the victim's environment. One important distinction to note here is that the statistics presented for this phase are solely based on Verizon's 2011 caseload. Just as criminals are quick to get in the door, the data shows that not a great deal of time typically passes before the loot is in hand. Well over half the time, the victim's data was removed within hours of initial compromise. This drops to 25% for larger size organizations, but we're not quite ready to start patting them on the back yet. It's entirely likely that the difference is simply because a little more searching (and thus time) is required to find the loot rather than the existence of ironclad controls preventing data loss.

If we can say something positive about these findings, it is this: In over 40% of incidents we investigated, it took attackers a day or more to locate and exfiltrate data. This gives some hope that reasonable time exists for more than one shot at detecting/stopping the incident before data is completely removed from a victim's control.

Initial Compromise to Discovery

This phase deals with the span of time from when the first asset is negatively affected to the moment when the victim learns of the incident. It saddens us to report that, yet again, breach victims could have circumnavigated the globe in a pirogue (not from the bayou? Look it up.) before discovering they were owned. In over half of the incidents investigated, it took months—sometimes even years—for this realization to dawn. That's a long time for customer data, IP, and other sensitive information to be at the disposal of criminals without the owners being aware of it.

Because this topic is so important, and because we want to make it PowerPoint friendly, we've included Figure 42 dedicated to larger organizations. A little better, yes, but no cause for fanfare. You could still make it around the world faster than that by upgrading to a dingy with a small battery-powered motor.

Figure 42. Time between initial compromise and discovery – LARGER ORGS

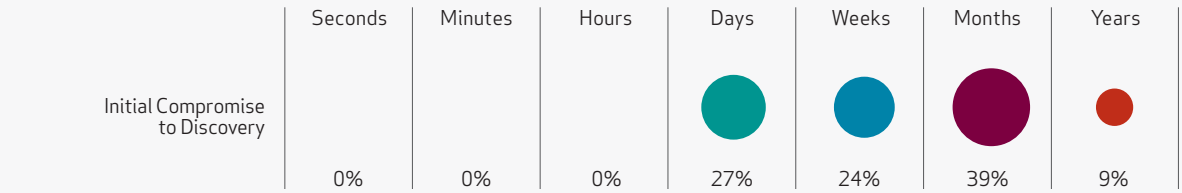

We see no real improvement over time with regard to this finding; nor are we observing any evidence that our data is somehow biased in this respect. With each passing year we examine more breaches from more sources, but like Bill Murray's alarm clock in "Groundhog Day," the song never changes. So, at this point we'll try to divert your attention away from harsh realities by inviting you to look at the pretty bubble picture and then hurry you right along. Nothing to see here folks.

Figure 43. Percent of breaches that remains undiscovered for months or more

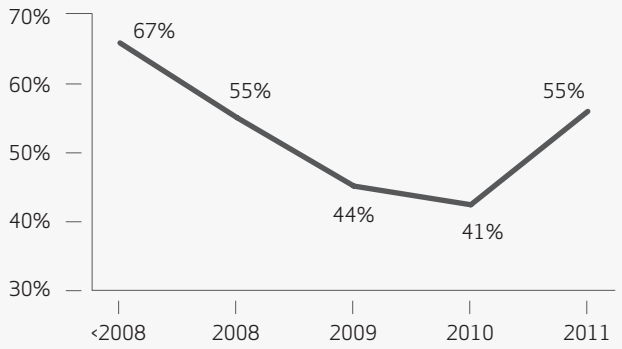

Discovery to Containment/Restoration

The final phase captures the span of time from when the victim learns of the incident to when it is contained (i.e., the "bleeding has stopped"), or the compromised system has been restored (i.e., the system is fully functional again).

We are cautiously optimistic to report that containment times are a teensy bit better than those of the previous year. 42% of breaches in 2011 were contained in within days compared 34% in 2010. Like we said—"teensy"—but definitely a shift in the right direction. Let's keeping building on that.

It's important to remember that incident response speed and execution are equally critical. Acting fast just for the sake of speed increases the risk of making mistakes, resulting in higher costs or needlessly extending the time necessary for full incident mitigation. Being effective in this respect does not mean having "x" number of incident responders on staff. Rather, it means the mechanisms are in place to properly initiate the response, classify the incident, and make timely decisions regarding when to engage external assistance or involve law enforcement. It means the practical elements—like the network diagram—are prepared and ready, or can be prepared when necessary. It also means working jointly with all parties involved towards the shared goal of containment. If this sounds like common sense, that's because it is. Nevertheless, organizations often fret over a myriad of possible bad outcomes rather than rolling up their sleeves and getting the incident resolved.

Breach Discovery Methods

Aside from each of the timespan metrics just discussed, the methods by which breaches are discovered account for some of the most interesting and sobering figures we have in our report. The time to discovery is intimately bound to the method of discovery, as some methods inherently take longer than others. Both metrics, however, are excellent indicators of the maturity of a security program since they directly reflect on the ability of an organization to detect and respond to threats. Unfortunately, as our research has shown for the last several years, third parties discover data breaches much more frequently than do the victim organizations themselves.

The Verizon RISK Team uses two main categories to group discovery methods: external and internal. External discovery is fairly self-explanatory; the breach was discovered by an external source and then reported to the victim.

Figure 44. Simplified breach discovery methods by percent of breaches

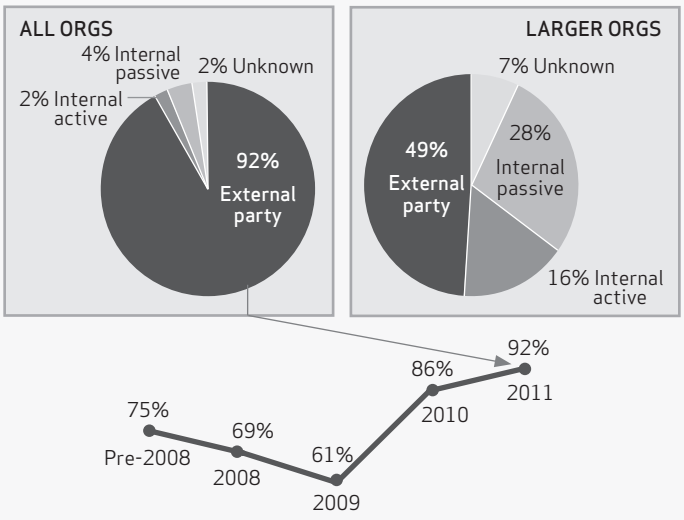

For internal discovery we further categorize methods as being active (those specifically designed/deployed for detection) or passive (those in which awareness of the incident arises from non-security processes).

Reviewing Figure 44, readers can readily observe that breaches discovered by external parties are at an all-time high. We attribute this to a couple of different factors, the foremost being the demographics of our dataset. Smaller businesses typically lack both the knowledge and the sophisticated technology required to monitor such events. Instead, they usually work along the lines of the well-known maxim, "If it ain't broke, don't fix it." And if there ain't nothing to say it's broke, it don't get fixed. The second reason is to some degree a function of the previous factor and to some extent is a by-product of the role law enforcement partners play in our dataset. We will discuss this in more detail in the external discovery section below.

When we compare the breach discovery results of the general population with those of large organizations, a very different picture emerges. Figure 44 illustrates that breaches discovered by external parties (49%) are substantially lower, while internal methods of discovery are much higher (44%) compared to the overall dataset. We have said in the past that larger organizations have the knowledge, capabilities, and resources to discover security breaches more effectively (if they would only use them), and these results bear that out. We will dive into this a bit more in a moment, but within the external discovery breach methods, there are a few good nuggets.

Figure 45. Breach discovery methods by percent of breaches

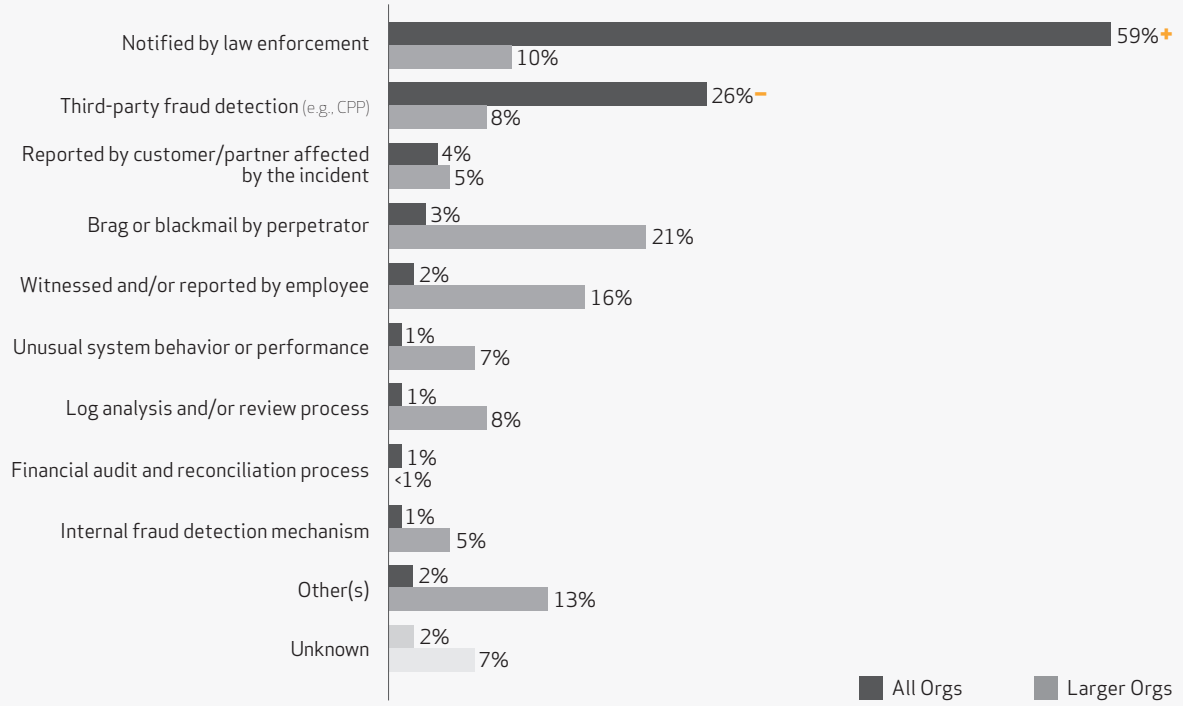

Another sobering thought with regard to CPP is that its effectiveness is predicated upon the ability to correlate certain fraudulent activities and patterns. Other types of valuable information, such as intellectual property or other sensitive organizational data, can be harvested from numerous sources but go completely unnoticed.

External Discovery

Notification by law enforcement was once again one of the main methods of breach discovery. In fact, it takes the top spot this year over the usual champion, third-party fraud detection. We believe one of the reasons for this increase is the larger role that law enforcement plays in the monitoring of criminal groups and their wayward activities (see Appendix B.). In such circumstances, law enforcement has knowledge that an attack or breach has occurred and can therefore notify the victim before fraud patterns emerge via common point of purchase (CPP) algorithms.

While we are on the subject of CPP, it should be noted that third-party fraud detection continues to be one the strongest methods of breach discovery. The advantage of CPP lies with its ability to see a broader pattern emerge across many organizations that are often dispersed over a large geographic area, thereby providing a more comprehensive vantage point than any single organization could possibly provide. Despite this, CPP is still an after-the-fact alert, and its success points to the hard-to-hear truth that fraudulent activity triggers detection better than many of our tools designed to do so before data loss or fraud ever occurs. Another sobering thought with regard to CPP is that its effectiveness is predicated upon the ability to correlate certain fraudulent activities and patterns. Other types of valuable information, such as intellectual property or other sensitive organizational data, can be harvested from numerous sources but go completely unnoticed. Thus, we believe the numbers surrounding non-payment card breaches are far worse than reported since there is no CPP-like mechanism to detect their loss.

External breach notification methods are much different for large organizations. While notification by law enforcement was the second most seen, at 10%, it was still far lower than that of the overall dataset. In most cases for large organizations notification occurred when the thief made the disclosure known. Perhaps we should create new breach discovery classifications of "YouTube," "Pastebin," and "Twitter" for the 2013 DBIR? (Of course, we're joking (sort of), but it is quite important to understand the role social networking plays in breach discovery, but also in how attacks are initiated using these tools. Perhaps we'll follow up with a blog post another time.) An interesting "what-if" scenario would be whether or not these organizations would have discovered these breaches through some sort of internal breach discovery method. In many cases, there is little evidence suggesting they would.

Internal Active

Internal active discovery relates to IDS/IPS/HIPS, log monitoring, anti-virus, and other like technologies that organizations typically use to prevent, detect, and respond to data breaches. Unfortunately, as referenced in several places, many smaller organizations do not have the awareness, aptitude, funding, or technical support to perform these tasks on par with the rapidity and/or sophistication of the threats they face.

Nevertheless, although large businesses are better at utilizing these investments compared to the overall dataset, they still have a difficult time arriving at any significant return. In 8% of breaches affecting large organizations, it was basic log-review and analysis that topped the internal active discovery list. Readers will probably already be aware that this is one of the methods that we tout yearly, and believe to be more effective than nearly all other methods. How do we know this? Well, when we conduct an investigation, that's how we find the breach—reading the logs (see Figure 46). We are also long-time proponents of searching for haystacks instead of needles within logs. Take a look at the side bar on the next page for more discussion.

Figure 46. Availability of log evidence for forensics by percent of breaches*

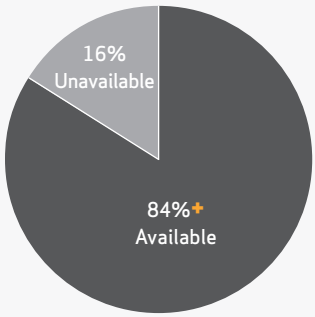

* Verizon caseload only

OF LOGS, NEEDLES, AND HAYSTACKS—PART DEUX

Some loyal readers may remember a sidebar from the 2010 DBIR titled "Of Logs, Needles, and Haystacks." We thought this year would be an opportune moment to present the second installment.

We often use what we refer to as the "Sesame Street" method of discovering a data breach. If you have kids, or have been a kid at some stage in your life, or, like some we know—never stopped being a kid, you may recall the Sesame Street song titled "One of These Things is Not Like the Other." In our role as investigators, we often employ this method when searching for evidence of a security incident. To work smarter rather than harder, we stopped looking for the needles, and began looking at the haystacks. If one of the haystacks is "not like the others," then that's a good place to start looking for needles. Following are a few real world examples from our cases.

When we present the data depicted below, we often ask the audience, "Can anyone see something anomalous with this chart?" After the chuckling subsides, we illustrate that the chart could be made the size of a postage stamp (hopefully we aren't dating ourselves too much here) and one could still immediately notice "one of these things is not like the other." The particular example below shows log file size, but this type of graph could represent any number of artifacts.

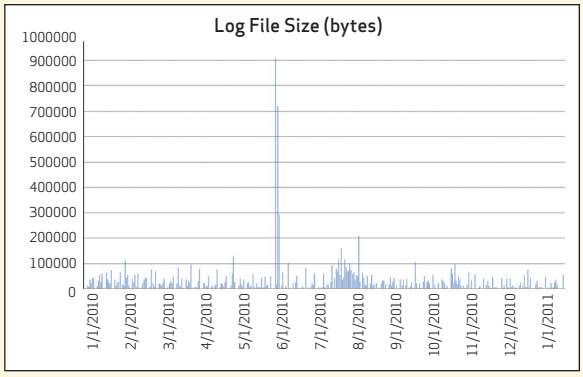

We have seen other very useful compromise indicators such as log file line count, log file line length (particularly good with SQLi in the URI), spikes in traffic types (e.g. ssh, FTP, DNS, or anything atypical to the organizational norms), number of events, country of origin of IP connection (or by protocol), bandwidth usage, and e-mail messages sent/received. The list is as long as the limits of the imagination. The really interesting thing about this type of monitoring is that it doesn't take a ton of cash to implement an effective solution. It can be done with a few commands on a Linux or Windows system.

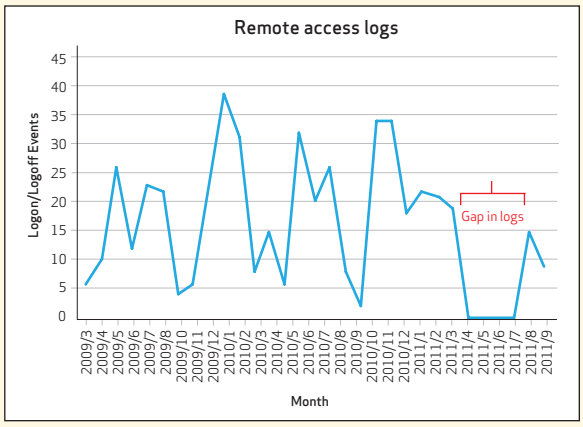

One thing to keep in mind when looking at number of events is that it's not just the dramatic increases that should ring the alarm bells, but the decrease of events may also be a cause for concern and an even louder ringing of the bells. The remote access log entry count below shows a spiky graph of logon/logoff events, but there is a gap that helped us determine when the data breach occurred.

All of this is, of course, obvious with glorious 20/20 hindsight, but that's a privilege of our job!

Internal Passive

Internal passive discovery is best described as when someone or something that is not responsible for security reports the signs of an incident. Typically, this involves a regular employee who, in the course of their daily responsibilities, notices something strange (e.g., slower system performance). More important than what they notice is that when they witness strange circumstances they make the conscious decision to ACT on it. Unfortunately, our investigators commonly learn through interviews that strange activity was noticed on a system, but nothing was said or done. Usually, the employee states that either they didn't know what to do, or they didn't feel it was important. Organizations that properly train personnel to recognize the tell-tale signs of a security incident and how to respond will reap the side benefits of having an army of roaming incident-detectors working for them. Last year, we compared this to the wonderful gift of free beer; we'll appeal to a wider audience this time and go with free wine.

Based on these results, it would appear that the wine flows more freely in larger organizations. 16% of their incidents were witnessed and reported by employees. This makes sense, as larger organizations have more security awareness campaigns due to the requirements of various compliance regimes, and training is often conducted annually. On the other end of the spectrum, the typical store clerk will not usually notice strange and unusual circumstances unless it hampers their Bejeweled game on Facebook. Although it may not keep employees from clicking on links and opening attachments from unknown senders, it can help an organization learn about a potential security incident faster if their employees recognize that something is amiss and react.

Anti-Forensics

The obvious problem with identifying anti-forensics techniques is: if they are effective, how would we know? Nonetheless, our investigators often do identify anti-forensics techniques in the field. And when they do, they make a note of it.

In 2011, anti-forensics were—and have been for several years running—observed in about one-third of Verizon's caseload. We still believe this represents a very conservative low-end estimate of the prevalence of anti-forensics. Naturally, these techniques vary in nature and purpose, but they're all centered on erasing, hiding, or corrupting evidence in some way or other. Specific common examples include wiping or tampering with logs, encrypting compromised data so that it cannot be examined, and rearranging the crime scene to mask the obvious fact that Mrs. Peacock killed Colonel Mustard in the study with the candlestick (just seeing if you're awake).

The obvious problem with identifying anti-forensics techniques is: if they are effective, how would we know? Naturally, these techniques vary in nature and purpose, but they're all centered on erasing, hiding, or corrupting evidence in some way or other.

Somewhat surprisingly, the use of anti-forensics does not seem to be markedly different for larger organizations. Anecdotally, however, it seems that criminals don't discriminate when it comes to the application of anti-forensics techniques. The obfuscation of evidence is just as much at home in "APTs" as it is in scareware. Timestomping, packing, encryption; these techniques are the rule, not the exception.

AN INTRODUCTION TO TIMESTOMPING

Building a timeline of system events is one of the primary tasks of any forensics investigator. Stomping all over that timeline is one of the most common anti-forensics techniques out there. We are all familiar with the standard modified, accessed, and created timestamps. However, in the NTFS file system there are a number of other timestamps associated with every file. By understanding how these timestamps work we are better able to identify when timestomping is occurring.

In NTFS a file is a collection of attributes indexed by a Master File Table (MFT). The timestamps we all know and love can be found in the $STANDARD_INFORMATION attribute. In this same location is a fourth timestamp: the MFT entry modified timestamp. All four of these timestamps are accessible via NtQueryInformationFile in the Windows API. Code that is stomping on these timestamps will likely be linking NtSetInformationFile.

Another attribute that can be found associated with an NTFS file is the $FILE_NAME attribute. This attribute is associated with particular hard links of a file, and each comes with its own unique set of all four timestamps. These timestamps are not accessible (or directly modifiable) via the Windows API.

The crux of this discussion is that by accessing the standard Windows API, an attacker with sufficient privilege can modify all of the timestamps in the file system to effectively render timeline analysis pointless. In the physical world, it would be akin to a hurricane blowing across your crime scene. Timestomping is not game over for a skilled investigator, but it is certainly a significant setback.

Figure 47. PCI DSS compliance status based on last assessment*

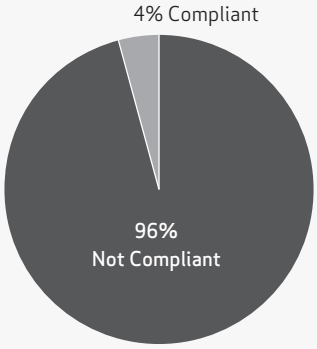

* Verizon caseload only. Largely based on victims' claims re their status (we don't force them to prove it).

PCI DSS

The Payment Card Industry Data Security Standard (PCI DSS) is a set of control requirements created to help protect cardholder data. Every year Verizon's caseload contains a number of organizations that are required to adhere to the PCI DSS. Because these are confirmed data breach victims, obvious questions arise with respect to the compliance status of these organizations. This section examines this important topic from several perspectives.

When reviewing this data, it is important to be mindful of the ways in which validation is performed. Take a moment and think about how most of us validate our own lives. Quite often this is based on the choices that we make and how they *compare with that of others*—whether it is peers, neighbors or celebrities. We often see that this also carries over into our professional lives. When we start talking about compliance validation, the top comments that we typically hear are "How do we compare to XYZ Corporation?" and "How do we fare in comparison to our industry?" On the surface this may seem like mere curiosity—or is it more telling than that?

However, there is a fundamental flaw with that comparative mindset—particularly for organizations that factor that into their IT security strategy. We allow ourselves to believe that being just slightly better than others also somehow equates to being more secure.

This is what it would look like on a "Good, Better, Best" scale:

- Good—"My security is better than many of my peers, but we're still not meeting our compliance requirements…"
- Better—"My security is better than most of my peers and also meets the *letter* of our compliance requirements…"
- Best—"My security is better than most of my peers, meets the *spirit* of our compliance requirements, and evolves with the changing threat landscape…"

What does this mean with respect to PCI DSS? When looking at Figure 47, we can clearly see that 96% of breach victims were not compliant as of their last assessment (or were never assessed/validated). The majority of these organizations don't make the "Good" category. Interestingly, this is the highest percentage of non-compliance that our report has ever noted. Although it can be interpreted a variety of ways, there are clearly few breach victims that fall in the "Better" or "Best" categories.

The majority of these organizations don't make the "Good" category. Interestingly, this is the highest percentage of non-compliance that our report has ever noted.

Does this mean that PCI DSS has failed? We think that is unlikely. If it were a failure of the PCI DSS, we would expect to see a much higher percentage of organizations being breached that at least fell into the "Better" category and a much lower non-compliance ratio. Rather, we are seeing a continuing trend whereby more of the organizations that fall in the 96% tend to be on the small side—a shift towards Level 4 merchants. In many cases these organizations have either failed to perform their (self) assessments or failed to meet one or more of the requirements. The most notable distinction here is that of the merchant failing to perform its assessment or achieving compliance versus a failure of the PCI DSS itself. It is also important to note that while

Verizon has investigated a significant number of breaches to larger organizations, many of those involved data compromises that *did not include* PCI data (i.e., large entity breaches do not always equate to large scale PCI data compromise).

Taking a closer look at the individual compliance requirements in Table 14, we see that 8 out of the 12 were found to be "Not in Place" at least two-thirds of the time. We believe this further reinforces the discussion above. While it is difficult to correlate every possible cause-and-effect relationship, in nearly every case we have concluded that such a significant deviation from PCI DSS (as well as many industry standards) was a major factor in these victims' breaches.

Table 14. Percent of relevant organizations in compliance with PCI DSS requirements based on post-breach reviews conducted by Verizon IR team*

				2011	
	2008	2009	2010	All Orgs	Lrg Orgs
Build and Maintain a Secure Network					
Requirement 1: Install and maintain a firewall configuration to protect data	30%	**35%**	18%	29%	71%
Requirement 2: Do not use vendor-supplied defaults for system passwords and other security parameters	**49%**	30%	33%	42%	86%
Protect Cardholder Data					
Requirement 3: Protect Stored Data	11%	**30%**	21%	18%	14%
Requirement 4: Encrypt transmission of cardholder data and sensitive information across public networks	68%	**90%**	89%	89%	86%
Maintain a Vulnerability Management Program					
Requirement 5: Use and regularly update anti-virus software	**62%**	53%	47%	23%	86%
Requirement 6: Develop and maintain secure systems and applications	5%	**21%**	19%	20%	57%
Implement Strong Access Control Measures					
Requirement 7: Restrict access to data by business need-to-know	24%	30%	33%	**36%**	29%
Requirement 8: Assign a unique ID to each person with computer access	19%	**35%**	26%	20%	57%
Requirement 9: Restrict physical access to cardholder data	43%	58%	65%	**73%**	100%
Regularly Monitor and Test Networks					
Requirement 10: Track and monitor all access to network resources and cardholder data	5%	**30%**	11%	11%	43%
Requirement 11: Regularly test security systems and processes	14%	**25%**	19%	6%	33%
Maintain an Information Security Policy					
Requirement 12: Maintain a policy that addresses information security	14%	**40%**	16%	16%	83%

Another interesting parsing of the dataset revealed that there was a noticeable difference in the way the PCI DSS compliance picture looks if you were to carve the 96% into two size categories: a) companies smaller than 1,000 employees, and b) companies of 1,000 or more employees. The majority of Level 4's discussed earlier fall into the sub-1,000 employee category. However, looking at the larger organizations you find that the number of "Not in Place" items using the same measurement as above drops from eight to only three. What might we be able to infer from this?

Possibly that larger organizations are doing a better job at complying with the standard (or implementing security fundamentals in general) and thus leaving themselves less exposed and subjected to being breached. This would also tend to align with some of the data presented by the Payment Card Industry regarding the compliance rates of the large-merchant population.

Another interesting area to examine is that of the four percent of victims found to be compliant but still suffered a breach. We still hear the common mantra "How could I have been breached?—I'm compliant!" We cannot stress enough that while compliance definitely helps drive security, compliance *does not equal* security. And we believe the perpetrators of these breaches share a similar philosophy. In fact, in many cases where we have the opportunity to be involved with the prosecution's interviews we learn that the perpetrators rarely know who they are hacking. In most cases it seems that they learn the identity of their victim *after* they have gained unauthorized access.

Due to the point-in-time nature of PCI assessments, it is possible that an organization deemed compliant at its last audit may not still be compliant at the time of the breach. Furthermore, there is always some degree of variability across individual assessors.

Overall, organizations both large and small seem to struggle the most with requirements 3, 7, 10, and 11. Interestingly, when looking at the numbers on a year-over-year basis we see mixed progress:

- Improved (x5)—Requirements 1, 2, 6, 7, and 9
- Declined (x4)—Requirements 3, 5, 8, and 11
- Remained the same (x3)—Requirements 4, 10, and 12

Initially we may look at the above results as a positive trend—after all, more of the requirements improved than declined. However, if we examine the summation of the overall improvements (+32%) versus the overall decline (-46%), it actually produces a difference of -14%. The most significant improvement was Requirement 1 (+11%)—"Install and maintain a firewall configuration to protect data." The most significant decline was Requirement 5 (-24%)—"Use and regularly update anti-virus software."

We still hear the common mantra "How could I have been breached?—I'm compliant!" We cannot stress enough that while compliance definitely helps drive security, compliance *does not equal* security.

What does this all mean with regard to the effectiveness of the PCI DSS? Overall, the standard attempts to set a bar of essential practices for securing cardholder data. Nearly every case that we have seen thus far has attributes of its breach that could have been prevented if the control requirements had been properly implemented. Of course, there is no way to be certain that new and different tactics could not have been used by the perpetrators to circumvent a compliant entity's controls. However, as long as we continue to see organizations with the sizable gaps that we see here, we will likely continue to see the perpetrators utilize such vulnerabilities as the path of least resistance to gain unauthorized entry. Organizations should continue striving to achieve the best possible security posture. For most, PCI DSS continues to be a reasonable (and required) yardsticks for measuring their progress.

The Impact of Data Breaches

Nary a year goes by without someone asking why the DBIR lacks information about the impact or consequences of data breaches. As the well known axiom suggests, "You can never please all the people all the time." Actually, we not only understand their question, but we're standing right next to them shouting for more data on the magnitude of loss. We spend a great deal of time and text touting the virtues of the DBIR for risk management decisions, yet we provide little to no information on one of the two major components of risk.[20]

This omission is not for lack of effort; there are legitimate reasons that the report does not contain such information. Foremost among them is that most organizations are focused primarily on fighting the fires caused by a breach and want to get back to "business as usual" as quick as possible. The process of recording the variety of ways money

20 Factor Analysis of Information Risk (FAIR) defines risk as "the probable frequency and probable magnitude of future losses." See "An Introduction to Factor Analysis of Information Risk (FAIR)."," Risk Management Insight LLC, November 2006. Stats in the DBIR almost entirely relate to the frequency of loss events.

and resources are spent during the event is a distant thought at best. Even after all the fires are extinguished, for reasons we haven't exactly nailed down yet, organizations are not motivated to collect data on their losses. Secondly, an investigation focuses on collecting evidence to prove or disprove compromise, assess exposure, and contain the breach. Analyzing and quantifying financial losses to the victim organization is simply not what we're paid to do. Although we do occasionally come across pieces of such information, we do not have the opportunities to gather near enough to complete the puzzle. Nor are we on the ground long enough after the breach to truly study the long-term consequences.

While we might be able to publish the bits and pieces that we collect on losses, we made a conscious decision at the very beginning not to do so. One of the aspects of the DBIR that we (and we hope many others) like is that, from cover to cover, it is filled with objective, credible, factual information. Since we do not collect data of that caliber on losses during an investigation, we have not felt it fits with the rest of the report.

That said, we're just as eager to gather and report the data as some readers are to see it. We realize that breach details describing the threat agents, actions, and loss events along with a credible account of financial losses are the Holy Grail of our field. Because of this, we've reconsidered our aforementioned decision to not publish "bits and pieces," and have begun to make whatever impact-related observations we can both during and after the investigation. Sometimes this yields almost nothing; sometimes it can be very informative. We do not have enough data compiled for 2011 to make this a statistics-driven section, but we can relay some of these observations to you.

We realize that breach details describing the threat agents, actions, and loss events along with a credible account of financial losses are the Holy Grail of our field. Because of this, we've reconsidered our aforementioned decision to not publish "bits and pieces," and have begun to make whatever impact-related observations we can both during and after the investigation.

While a few 2011 breach victims estimated their losses to be in the hundreds of millions of dollars, most did not get near to that amount. In fact, the large majority of them emerged relatively unscathed from their troubles. While they were inconvenienced and probably had a sleepless night or two, they returned to normal business operations before too long.

Hands down, the easiest data surrounding impact that we noted in 2011 was response and recovery losses, which include forensics. Naturally, we're biased here, since we don't typically work for free. For larger organizations, such fees don't normally affect the bottom line, but they can be hard to cover for smaller businesses. Fraud-related losses were also very common, which isn't surprising given the large number of breaches to payments cards and identifying information. These ranged from a few bucks on the low end to in excess of $100 million.

More rare, but also more substantial, were regulatory and legal fees that were noted (often through public domain information) for several victims. Sometimes it's difficult to discern whether these costs actually get paid out, but class action lawsuits and similar filings can require a sizeable amount at least be set aside.

Brand damage, declines in market value, and loss of competitive advantage are always the top of mind "WIBeHI" (Wouldn't it be horrible if…) fears for executives with respect to data breaches. For most breaches—even ones that seem rather bad—these fears are unfounded. Breaches don't appear to typically have a major long-term impact on stock value.

There are exceptions, however, and we observed some of them in 2011. We know of at least four victims that are no longer in business, either wholly or in large part because of their breach. One organization was unable to recover from collateral damages to their client-facing systems that occurred during the incident. Another's main line of business depreciated to such a degree they closed up shop. A smaller "click and mortar" retailer decided "clicks" weren't worth the additional risk exposure after a breach, and ceased e-commerce operations. Other examples include outraged customers, soured B2B relationships, governmental scrutiny, and other consequences nobody wants to experience. Fortunately, these occurrences are rare, but we feel it incumbent upon us to mention that periodically they do occur.

It is also important to remember that the impact is often felt by organizations that are not the breach victim. In payment card breaches, for instance, the card issuers (the banks that provide payment cards to their customers) can be severely affected by fraud on customer accounts. These downstream effects are often neglected when considering the overall consequences of data breaches.

As we close out this short section (with the hope that it will grow in the future), we would like to state our belief that assessing the impact of an incident is a worthwhile exercise. We understand that most organizations just want to "get back to good" following a breach rather than dwell on the consequences, but such crises can also be an opportunity for valuable learning experiences. Not only can this help to better understand the incident in question, but also the data derived can benefit future risk analyses and decision-making.

Organizations wishing to assess the impact of incidents are invited to—you guessed it—utilize VERIS. The schema specifies a few baseline data points that can help you get started. We offer it up to your review and consideration.

In our litigious age, Verizon investigators employed in our Electronic Discovery (e-discovery) unit often see legal action ensue almost immediately following a data breach. Best practices dictate that a corporation should issue a legal hold notice as soon as they reasonably anticipate litigation. Depending on the nature and the type of the breach, an organization would be wise to consider a breach the impetus to initiate their internal discovery response process. Shortly after commencing this process, an organization should identify relevant data custodians and data repositories, then issue a legal hold followed by a number of other processes.

E-Discovery has taken on a much larger role in the overall discovery process. Legal and IT teams need to understand their organization's ability to respond to litigation or an investigation. Having proven and well-defined discovery response processes will help lower the cost, time, and risk associated with discovery requests and reduce the overall risk associated with incident response. With an organized business process, corporate counsel can systematically and effectively manage the increasingly complex nature of incident and discovery response.

Far too often we see organizations that don't have many (or any) documented E-Discovery processes, with little training, limited technology, and, most importantly, no executive-level support to effectively and efficiently respond to discovery requests. Organizations serious about mitigating elements such as cycle time, cost, and risk should consider having a well-defined discovery response program (which includes a discovery response plan and a discovery response team) that can be invoked consistently. However, prior to creating such a program, it's imperative to start with an overall, objective assessment of the people, processes, and technologies associated with discovery response—in order to leverage what may already be in place and determine where the greatest gaps exist.

Got a question or a comment about the DBIR?
Drop us a line at dbir@verizon.com, find us on Facebook, or post to Twitter with the hashtag **#dbir**.

2012 DBIR: CONCLUSIONS AND RECOMMENDATIONS

This year, we're including something new in this section. However, being the environmentally conscious group that we are, we're going to recycle this blurb one more time:

> ***"Creating a list of solid recommendations gets progressively more difficult every year we publish this report. Think about it; our findings shift and evolve over time but rarely are they completely new or unexpected. Why would it be any different for recommendations based on those findings? Sure, we could wing it and prattle off a lengthy list of to-dos to meet a quota but we figure you can get that elsewhere. We're more interested in having merit than having many."***

Then, we're going to reduce and reuse some of the material we included back in the 2009 Supplemental DBIR, and recast it in a slightly different way that we hope is helpful. As mentioned, we've also produced something new, but made sure it had a small carbon (and page space) footprint. If you combine that with the energy saved by avoiding investigator travel, shipping evidence, and untold computational cycles, these recommendations really earn their "green" badge.

Figure 48. Cost of recommended preventive measures by percent of breaches*

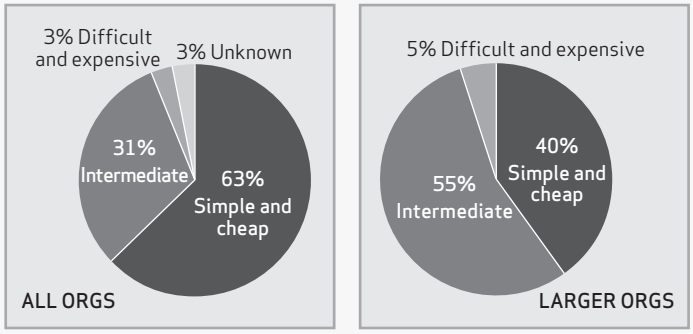

* Verizon caseload only

Let's start with the "something new." We've come to the realization that many of the organizations covered in this report are probably not getting the message about their security. We're talking about the smaller organizations that have one (or a handful) of POS systems. The cutout on the next page was created especially for them and we need your help. We invite you, our reader, to cut it out, and give it to restaurants, retailers, hotels, or other establishments that you frequent. In so doing, you're helping to spread a message that they need to hear. Not to mention, it's a message that the rest of us need them to hear too. These tips may seem simple, but all the evidence at our disposal suggests a huge chunk of the problem for smaller businesses would be knocked out if they were widely adopted.

Now we'll move on to the "reduce and reuse" part, which centers on larger organizations (which is not to say it's irrelevant to SMBs). As we have seen throughout this report, problems affecting larger organizations look different in many ways than those plaguing the masses of smaller businesses in our sample. It makes sense that the solutions to these problems would likewise be different. Hence, the dual focus of this section.

The cutout on the next page was created especially for smaller organizations and we need your help. We invite you, our reader, to cut it out, and give it to restaurants, retailers, hotels, or other establishments that you frequent.

POINT-OF-SALE SECURITY TIPS

Greetings. You were given this card because someone likes your establishment. They wanted to help protect your business as well as their payment and personal information.

It may be easy to think "that'll never happen to me" when it comes to hackers stealing your information. But you might be surprised to know that most attacks are directed against small companies and most can be prevented with a few small and relatively easy steps. Below you'll find a few tips based on Verizon's research into thousands of security breaches affecting companies like yours that use point-of-sale (POS) systems to process customer payments. If none of it makes sense to you, please pass it on to management.

✓ **Change administrative passwords on all POS systems**
  – Hackers are scanning the Internet for easily guessable passwords.

✓ **Implement a firewall or access control list on remote access/administration services**
  – If hackers can't reach your system, they can't easily steal from it.

After that, you may also wish to consider these:
- Avoid using POS systems to browse the web (or anything else on the Internet for that matter)
- Make sure your POS is a PCI DSS compliant application (ask your vendor)

If a third-party vendor looks after your POS systems, we recommend asking them to confirm that these things have been done. If possible, obtain documentation. Following these simple practices will save a lot of wasted money, time, and other troubles for your business and your customers.

For more information, visit www.verizon.com/enterprise/databreach (but not from your POS).

For those who don't remember (tsk, tsk), the 2009 Supplemental DBIR was an encyclopedia of sorts for the top threat actions observed back then. Each entry contained a description, associated threat agents, related assets, commonalities, indicators, mitigators, and a case study. To provide relevant and actionable recommendations to larger organizations this year, we're repurposing the "indicators" and "mitigators" part from that report.

- **Indicators:** Warning signs and controls that can detect or indicate that a threat action is underway or has occurred.
- **Mitigators:** Controls that can deter or prevent threat actions or aid recovery/response (contain damage) in the wake of their occurrence.

Our recommendations will be driven off of Table 7, which is in the Threat Action Overview section, and shows the top ten threat actions against larger organizations. Rather than repeat the whole list here, we'll summarize the points we think represent the largest opportunities to reduce our collective exposure to loss:

- Keyloggers and the use of stolen credentials
- Backdoors and command control
- Tampering
- Pretexting
- Phishing
- Brute force
- SQL injection

Hacking: Use of stolen credentials	
Description	Refers to instances in which an attacker gains access to a protected system or device using valid but stolen credentials.
Indicators	Presence of malware on system; user behavioral analysis indicating anomalies (i.e., abnormal source location or logon time); use of "last logon" banner (can indicate unauthorized access); monitor all administrative/privileged activity.
Mitigators	Two-factor authentication; change passwords upon suspicion of theft; time-of-use rules; IP blacklisting (consider blocking large address blocks/regions if they have no legitimate business purpose); restrict administrative connections (i.e., only from specific internal sources). For preventing stolen credentials, see *Keyloggers and Spyware, Pretexting,* and *Phishing* entries.

Malware: Backdoors, Command and Control Hacking: Exploitation of backdoor or command and control channel	
Description	Tools that provide remote access to and/or control of infected systems. Backdoor and command/control programs bypass normal authentication mechanisms and other security controls enabled on a system and are designed to run covertly.
Indicators	Unusual system behavior or performance (several victims noted watching the cursor navigating files without anyone touching the mouse); unusual network activity; IDS/IPS (for non-customized versions); registry monitoring; system process monitoring; routine log monitoring; presence of other malware on system; AV disabled. During investigations involving suspected malware we commonly examine active system processes and create a list of all system contents sorted by creation/modification date. These efforts often reveal malicious files in the Windows\system32 and user temporary directories.
Mitigators	Egress filtering (these tools often operate via odd ports, protocols, and services); use of proxies for outbound traffic; IP blacklisting (consider blocking large address blocks/regions if they have no legitimate business purpose); host IDS (HIDS) or integrity monitoring; restrict user administrative rights; personal firewalls; data loss prevention (DLP) tools; anti-virus and anti-spyware (although increased customization rendering AV less effective—we discovered one backdoor recognized by only one of forty AV vendors we tried); web browsing policies.

Physical: Tampering	
Description	Unauthorized altering or interfering with the normal state or operation of an asset. Refers to physical forms of tampering rather than, for instance, altering software or system settings.
Indicators	An unplanned or unscheduled servicing of the device. Presence of scratches, adhesive residue, holes for cameras, or an overlay on keypads. Don't expect tampering to be obvious (overlay skimmers may be custom made to blend in with a specific device while internal tampering may not be visible from the outside). Tamper-proof seal may be broken. In some cases an unknown Bluetooth signal may be present and persist. Keep in mind that ATM/gas skimmers may only be in place for hours, not days or weeks.

Physical: Tampering	
Mitigators	Train employees and customers to look for and detect signs of tampering. Organizations operating such devices should conduct examinations throughout the day (e.g., as part of shift change). As inspection occurs, keep in mind that if the device takes a card and a PIN, that both are generally targeted (see indicators). Set up and train all staff on a procedure for service technicians, be sure it includes a method to schedule, and authenticate the technician and/or maintenance vendors. Push vendor for anti-tamper technology/features or only purchase POS and PIN devices with anti-tamper technology (e.g., tamper switches that zero out the memory, epoxy covered electronics).

Keylogger/Form-grabber/Spyware	
Description	Malware that is specifically designed to collect, monitor, and log the actions of a system user. Typically used to collect usernames and passwords as part of a larger attack scenario. Also used to capture payment card information on compromised POS devices. Most run covertly to avoid alerting the user that their actions are being monitored.
Indicators	Unusual system behavior or performance; unusual network activity; IDS/IPS (for non-customized versions); registry monitoring; system process monitoring; routine log monitoring; presence of other malware on system; signs of physical tampering (e.g., attachment of foreign device). For indicators that harvested credentials are in use, see *Unauthorized access via stolen credentials.* During investigations involving suspected malware we commonly examine active system processes and create a list of all system contents sorted by creation/modification date. These efforts often reveal malicious files in the Windows\system32 and user temporary directories.
Mitigators	Restrict user administrative rights; code signing; use of live boot CDs; onetime passwords; anti-virus and anti-spyware; personal firewalls; web content filtering and blacklisting; egress filtering (these tools often send data out via odd ports, protocols, and services); host IDS (HIDS) or integrity monitoring; web browsing policies; security awareness training; network segmentation.

Pretexting (Social Engineering)	
Description	A social engineering technique in which the attacker invents a scenario to persuade, manipulate, or trick the target into performing an action or divulging information. These attacks exploit "bugs in human hardware" and, unfortunately, there is no patch for this.
Indicators	Very difficult to detect as it is designed to exploit human weaknesses and bypasses technological alerting mechanisms. Unusual communication, requests outside of normal workflow, and instructions to provide information or take actions contrary to policies should be viewed as suspect. Call logs; visitor logs; e-mail logs.

Pretexting (Social Engineering)	
Mitigators	General security awareness training; clearly defined policies and procedures; do not "train" staff to ignore policies through official actions that violate them; train staff to recognize and report suspected pretexting attempts; verify suspect requests through trusted methods and channels; restrict corporate directories (and similar sources of information) from public access.

Brute-force attack	
Description	An automated process of iterating through possible username/password combinations until one is successful.
Indicators	Routine log monitoring; numerous failed login attempts (especially those indicating widespread sequential guessing); help desk calls for account lockouts.
Mitigators	Technical means of enforcing password policies (length, complexity, clipping levels); account lockouts (after x tries); password throttling (increasing lag after successive failed logins); password cracking tests; access control lists; restrict administrative connections (i.e., only from specific internal sources); two-factor authentication; CAPTCHA.

SQL injection	
Description	SQL Injection is an attack technique used to exploit how web pages communicate with back-end databases. An attacker can issue commands (in the form of specially crafted SQL statements) to a database using input fields on a website.
Indicators	Routine log monitoring (especially web server and database); IDS/IPS.
Mitigators	Secure development practices; input validation (escaping and whitelisting techniques); use of parameterized and/or stored procedures; adhere to principles of least privilege for database accounts; removal of unnecessary services; system hardening; disable output of database error messages to the client; application vulnerability scanning; penetration testing; web application firewall.

Unauthorized access via default credentials	
Description	Refers to instances in which an attacker gains access to a system or device protected by standard preset (and therefore widely known) usernames and passwords.
Indicators	User behavioral analysis (e.g., abnormal logon time or source location); monitor all administrative/privileged activity (including third parties); use of "last logon" banner (can indicate unauthorized access).
Mitigators	Change default credentials (prior to deployment); delete or disable default account; scan for known default passwords (following deployment); password rotation (because it helps enforce change from default); inventory of remote administrative services (especially those used by third parties). For third parties: contracts (stipulating password requirements); consider sharing administrative duties; scan for known default passwords (for assets supported by third parties).

Phishing (and endless *ishing variations)	
Description	A social engineering technique in which an attacker uses fraudulent electronic communication (usually e-mail) to lure the recipient into divulging information. Most appear to come from a legitimate entity and contain authentic-looking content. The attack often incorporates a fraudulent website component as well as the lure.
Indicators	Difficult to detect given the quasi-technical nature and ability to exploit human weaknesses. Unsolicited and unusual communication; instructions to provide information or take actions contrary to policies; requests outside of normal workflow; poor grammar; a false sense of urgency; e-mail logs.
Mitigators	General security awareness training; clearly defined policies and procedures; do not "train" staff to ignore policies through official actions that violate them; policies regarding use of e-mail for administrative functions (e.g., password change requests, etc.); train staff to recognize and report suspected phishing messages; verify suspect requests through trusted methods and channels; configure e-mail clients to render HTML e-mails as text; anti-spam; e-mail attachment virus checking and filtering.

Got a question or a comment about the DBIR?
Drop us a line at dbir@verizon.com, find us on Facebook, or post to Twitter with the hashtag **#dbir**.

APPENDIX A: EXAMINING RELATIONSHIPS AMONG THREAT ACTIONS

We thought we'd take some time this year and dig into the threat actions with a little more rigor. In the main report, we called out the top threat actions in Table 7 and specifically for larger companies in Table 8, but we thought there might be more to the story. This appendix takes a look into the relationships between threat actions observed in 2011.

Upon close examination of their actions, it appears as though threat agents (like most everyone else) strive to get the most reward or benefit from the least amount of work or investment. While they are intelligent beings and capable of adaptation, parsing the data suggests that they are only as intelligent and adaptive as they need to be.

Remember from Figure 18 that threat actions fall into seven categories (hacking, malware, social, etc.) and VERIS also enumerates specific varieties under each category. Every incident involves at least one threat action, but many have more. Figure 49, which gives a distribution for the number of threat actions per incident, illustrates this point.

Figure 49. Number of threat actions per breach

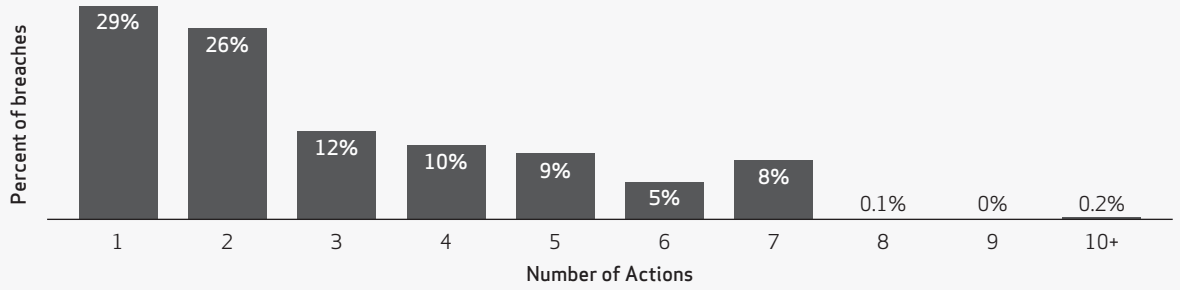

As you can see, 29% of breaches used only a single threat action, 26% of them used two threat actions, and so on. Note that over half successfully caused a loss through just one or two actions. The average number of actions per incident was 2.9 across the entire dataset compared to 2.75 in larger organizations. This in itself is interesting, but we're going to press on in an effort to detect patterns that help measure how (non) adaptive these adversaries are.

Breaches with a Single Threat Action

If we look at breaches containing only a single threat action, we see a bit of a different story than the overall information in the report (see Table 15). These one-dimensional threats typically bypass a single control layer that is between the attacker and the data they're after. For example, 60% of SQL injection attacks in the 2011 dataset were single-event incidents, meaning they exfiltrated data (or otherwise caused an incident) in the initial compromise and didn't continue beyond that. Single-event incidents are often over and done in a matter of seconds or even milliseconds.

The first thing to note is that Table 15 isn't a top 11 list. This is the entire list for incidents in which a single threat action was identified. That alone is informative; VERIS specifies over 150 threat actions, and yet all of the 240 single-event incidents can be described with only eleven of them. What's more, using just the top five covers 84% of the single-action incidents.

Table 15. Threat actions used in single-action breaches

Rank	Threat action	Category	All Orgs	Larger Orgs
1	Exploitation of default or guessable credentials	Hacking	104	0
2	Tampering	Physical	52	5
3	Pretexting (classic Social Engineering)	Social	24	3
4	Brute force and dictionary attacks	Hacking	21	0
5	SQL Injection	Hacking	12	1
6	Abuse of system access/privileges	Hacking	12	3
7	Use of stolen login credentials	Hacking	5	2
8	Exploitation of insufficient authentication (e.g., no login required)	Hacking	5	0
9	Publishing error	Error	3	3
10	Misconfiguration	Error	1	1
11	Embezzlement, skimming, and related fraud	Misuse	1	0

Breaches with Two or More Threat Actions

If we focus on breaches with two or more threat actions, some interesting patterns begin to emerge. In order to see the trees through the forest, we created all possible sets of pairs (for example a breach with two threat actions created a single pair, a breach with three threat actions created three pairs, four threat actions created six pairs, and so on). Using this method generates almost 4000 pairs of threat actions, of which 281 were unique. From this, we created Figure 51 which shows these action-to-action comparisons for breaches in which two or more threat events occurred.

While we don't expect you to be able to read the details (see that one yellow box right there?), what this shows is that 96% of the combinations are empty (grey) and only a handful of combinations are orange or red (indicating pairs that occurred frequently). That graph tells us that some patterns exist here and it's those patterns we really want to study in more detail. We plan to revisit this in a couple of follow-up blog posts in the near future, but for now, let's make some more eye candy with this data.

Figure 50. Cumulative Distribution of Threat Action Pairs

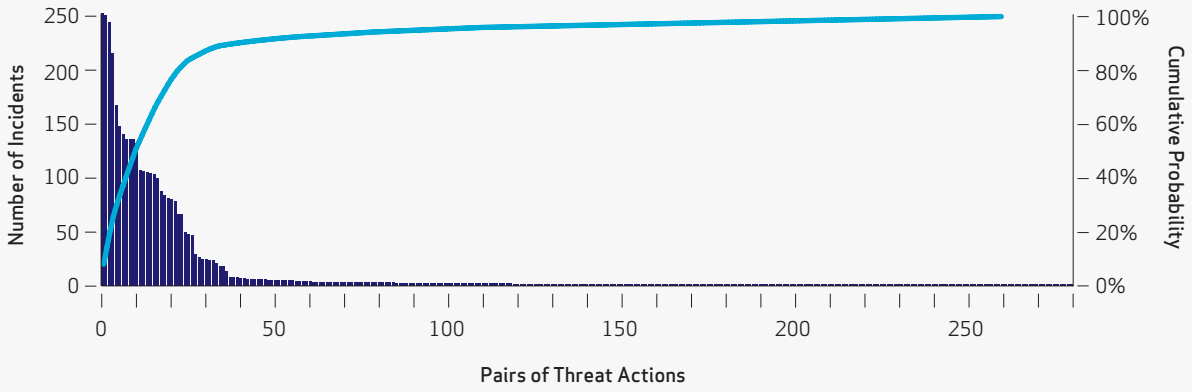

Figure 51. Pairs of Threat Actions

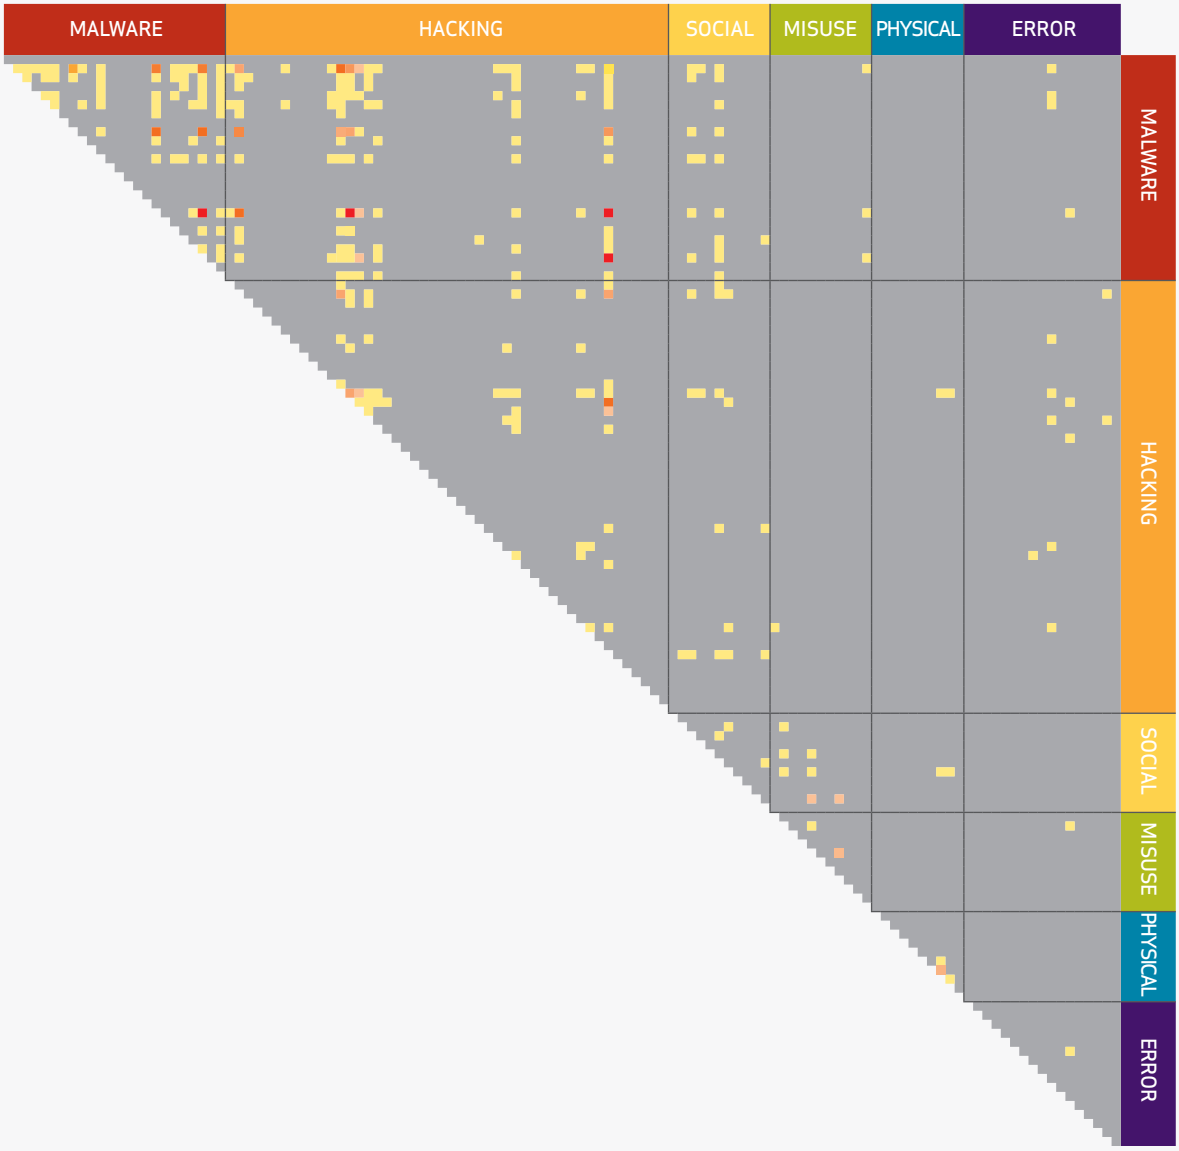

The dark blue area in Figure 50 represents the number of times a pair of threat actions were observed. For example, the paired action of "malware – keylogger" and "hacking – use stolen credentials" was leveraged the most, appearing together in 252 of the incidents in this year's report. That combination is the first blue line to the left and raises up to 252, and the other action pairs tail off in frequency from there. The light blue line in Figure 50 shows the cumulative percentage of pairs (the fancy term is a "cumulative distribution function"). We see a rather obvious turn in that light blue line around the twenty-fifth pair, so we'd probably see the best bang for our security buck if we focused on the top 25 pairings.

(As you read this next section, it's okay to say "wow" out loud.)

Table 16. Threat Actions from top 25 pairs

malware	Keylogger/Form-grabber/Spyware (capture data from user activity)
malware	Send data to external site/entity
malware	Backdoor (allows remote access/control)
malware	Disable or interfere with security controls
hacking	Exploitation of default or guessable credentials
hacking	Exploitation of backdoor or command and control channel
hacking	Brute force and dictionary attacks
hacking	Use of stolen login credentials

If we look further at these top 25 pairs (again, representing 81% of the almost 4000 pairs), we see that they are just different combinations of eight threat actions representing bread-and-butter for the majority of threat agents in 2011.

What this means is that the typical cybercriminal keeps reverting to the same permutations of a handful of tactics: basic hacking combined with malware. Sure, there are some intricate pieces of malware and some hacking techniques that stand above the rest, but by and large, we see attackers being only as intelligent and adaptive as they need to be. A slight rewording might drive the underlying point home more effectively: attackers are only as intelligent and adaptive as WE FORCE THEM TO BE. Clearly—as a community—we're not exactly forcing them to bring their A-game.

The threat-action pairs for larger organizations change in a couple of interesting ways (keep in mind the dataset is considerably smaller for larger organizations). First, malware appears to be more prominent and varied; 40% of the malware combinations (malware-malware) only appeared in larger organizations. Second, phishing attacks (combined with malware) pop into the top 25 in three different pairs and 84% of the pairs involving social only appeared in larger organizations. Finally, the list of pairs for larger organizations has much more variety, with only 36% duplicates (versus 93% overall). While this may be due to large organizations forcing the threat agents to bring at least their C-game, it's more likely a reflection on the automated, repetitious, and uniform attacks against smaller organizations.

But wait—we're not done; a few more tasty morsels can be pulled from this data. By comparing how often a specific threat action occurred overall with how often it was paired with other actions, we can make a series of statements. Anyone hoping to dazzle co-workers over drinks, impress a romantic prospect, or just have something to say during an uncomfortably lengthy elevator ride with higher-ups is welcome to give these a shot.

When malware was used to exfiltrate data, 98% of the time it was paired with keylogging functionality. This indicates that most keyloggers were set up to send out the information they collected. 251 incidents leveraged this combination and out of those, the attacker also used stolen credentials in all but eight of those incidents (yes, we ran this with groups of three threat actions as well).

98% of backdoors installed were paired with exploitation of the backdoor. This is not really surprising given that this report only covers breaches with known data loss. But still, that's a very high percentage, and shows backdoors aren't just scattered around haphazardly; they're likely to get used (something to think about when you see them in your network).

73.6% of people will believe a statement that includes a statistic, even if it is completely made up. Okay, we made that statement up. But now we can say that we've made up one statistic in the DBIR.

91% of the 276 breaches leveraging stolen credentials also had keyloggers installed. We mentioned this after the first statement, and this too may not be surprising since we've limited this data to only successful breaches. But these are "shock and awe" statements; let's take what we can from it.

When default or easily guessable credentials were used (in 379 incidents), 57% of the time a keylogger was installed and a backdoor was present in 47%. Remember though, that 104 incidents had the single threat action of leveraging default credentials.

Out of the 174 backdoors installed, 61% were also seen with a keylogger. Again, this isn't earth-shattering, but it is interesting to note that, more often than not, malware is multi-purpose.

To sum up, while a lot of this is just exploratory at this point, there are some important lessons for the present and directions for future study. For one, we've often stressed the importance of raising the cost of attack to the criminal. These results provide more than anecdotal evidence that their cost is not very high in the majority of breaches. We need to change that. Another thing to consider is the utility of pattern recognition to incident detection. While some monitoring technologies use various heuristics and pattern-matching, a great deal of our detection capabilities depends on one-dimensional signatures. The more we know about attack sequencing, behaviors, relationships, dependencies, etc, the better off we'll be in terms of recognizing badness when it's happening in front of our eyes.

Keep an eye on our blog (verizon.com/enterprise/securityblog) for more work like this. We hope to get some opportunity to dig into this type of data more and more as we go through the next year.

Got a question or a comment about the DBIR?
Drop us a line at dbir@verizon.com, find us on Facebook, or post to Twitter with the hashtag **#dbir**.

APPENDIX B: A USSS CASE STUDY OF LARGE-SCALE "INDUSTRIALIZED" CYBERCRIME

In April 2008, the U.S. Secret Service was contacted by a bank fraud investigator and advised that a number of the bank's debit cards had been compromised at a franchise restaurant in the New England area. All of the activity from the compromised cards took place in Europe and utilized manually punched transactions and white plastic cards.

Secret Service agents responded to the franchise location and conducted interviews with the owners, managers, and several employees in order to gain an understanding of the business process for accepting payment cards as well as who had access to the processing system. Preliminary investigations revealed that the point of sale (POS) server was connected to the Internet without a firewall in place or an antivirus program installed. Additionally, franchise employees were allowed to use the system to access the Internet to check personal web-based e-mail. Analysis of the POS server identified a keylogging program, which had been used to capture payment card data. This data was subsequently exfiltrated to a server owned by a web hosting company in Florida. Unfortunately, the web hosting company did not maintain appropriate log information to provide any assistance in this case.

The Secret Service was contacted again at a later date by a credit card provider and informed of additional fraudulent payment card activity related to this investigation. This time the fraudulent activity took place in Turkey, Bulgaria, Austria, and other European locations. The payment card data appeared to have been compromised at a separate location of the same franchise. Analysis of the POS server at this location again identified a keylogger, as well as Peer-to-Peer and password cracking tools. Further analysis of this POS server identified the login credentials as belonging to the same large web-hosting company mentioned above.

Meanwhile, Secret Service agents working a similar POS intrusion into a private business had again identified the same server and the same login credentials during their investigation. Based on this evidence, a Federal search warrant for this server was obtained and the analysis identified additional franchise locations that had been compromised throughout the United States.

Eventually, the Secret Service was contacted by the company that processed the payment cards for the entire franchise. The processor had investigated a POS intrusion at a location in Florida that identified payment card data that was exfiltrated to a server located at the same web hosting company identified by the Secret Service. Analysis revealed an additional 25 franchise locations that had been compromised by the same attacker using similar login credentials to those previously identified.

Secret Service agents worked together with the franchise headquarters to determine the origin and extent of the compromise. Analysis determined that the corporate network had not been compromised. It appeared that the attackers were compromising each location separately through remote administration services that were accessible from the Internet. Once the attackers gained access, keylogging software programs were installed to capture the payment card data.

Approximately 50 franchise locations had been compromised, and several FTP servers with similar login credentials had been identified at the same web hosting company. Analysis of these FTP servers also began to identify POS intrusions into other companies.

Additional analysis performed by CERT at Carnegie Mellon University identified artifacts originating from a country in Eastern Europe. CERT also identified the retrieval of keylogger files from an IP address that resolved to a private company's server. The company provided consent for the Secret Service to monitor the activity on this compromised server. This led to the identification of additional FTP servers at the same web hosting company that had received similar traffic as previously identified FTP servers, along with similar login credentials as the others.

Analysis of the traffic on the new FTP servers and the compromised private company server identified numerous connections resolving to the previously identified Eastern European country. Additional analysis of the FTP servers identified an additional 50 merchants with no affiliation to the franchise that initiated this investigation.

The traffic collected from the compromised private company server identified that it was being used for more than an intermediate collection point. Evidence collected from this server helped to identify the tools, tactics, compromised data, and the individuals attacking POS systems throughout the United States. A Remote Desktop Protocol (RDP) toolkit had been identified. The RDP tool appeared to be used to run automated scans of an IP range searching for computers with an open RDP port. When an open RDP port was identified the attackers would test a list of usernames and passwords or use password cracking tools to gain access to the POS system. In addition to hacking tools, several ICQ numbers were identified as the attackers would log into the server via Remote Desktop and start an ICQ chat client. They would log into web-based e-mail accounts, chat via ICQ, and retrieve files from the FTP servers.

Secret Service agents were able to identify several e-mail accounts and ICQ chat conversations from the suspected attackers. Agents served search warrants on these accounts and were able to start building a framework around this group of attackers. The connections made by the attackers to the compromised private company server and to web-based e-mail accounts were traced back to the previously identified country in Eastern Europe. Monitoring of the FTP servers that received payment card data and the compromised payment cards identified through e-mail search warrants continued to identify new victim businesses.

With the assistance of Secret Service offices located in Eastern Europe and their foreign law enforcement counterparts, the locations and real world identities of the attackers were pieced together from the evidence collected. Several of these individuals had been previously investigated in their country for "computer crimes" and were known to law enforcement.

During the course of the investigation, over 112,000 payment cards had been compromised from 163 of the franchise locations that initiated the case. Additionally, at least 800 other retail computer systems were compromised in numerous hotels, movie theaters, medical facilities, residential servers, pizzerias, bakeries, cafes, and various small businesses. Over $20 million in losses have currently been identified as a result of this investigation.

In May 2011, a Federal Grand Jury indicted four Eastern European individuals for conspiracy to commit computer-related fraud, conspiracy to commit wire fraud, and conspiracy to commit fraud in connection with access devices.

In a collaborative effort with law enforcement in Eastern Europe, three of the suspects were arrested. Two suspects were arrested by Secret Service agents while they were attempting to enter the United States. The third suspect was arrested by foreign law enforcement in Eastern Europe pursuant to a provisional arrest warrant issued by the Department of Justice and is awaiting extradition to the United States.

Got a question or a comment about the DBIR?
Drop us a line at dbir@verizon.com, find us on Facebook, or post to Twitter with the hashtag **#dbir**.

ABOUT THE 2012 DBIR CONTRIBUTORS

Verizon RISK Team

The RISK team is dedicated to Researching the ever-changing risk environment, Investigating and responding to all manner of security incidents, developing Solutions based on credible data and analysis, and cultivating Knowledge within Verizon, its clients, and the security community. The team has a wealth of experience and expertise, conducting over 1000 IR engagements over the last ten years. Included among them are many of the largest breaches ever reported. During these investigations, the team regularly interacts with governmental agencies and law enforcement personnel from around the world to transition case evidence and set the stage for prosecution. The expansive dataset generated through these activities offers an interesting glimpse into the trends surrounding computer crime and data compromise.

Australian Federal Police

The Australian Federal Police (AFP) Cybercrime Operations (CCO) Investigations Teams investigate offenses against Parts 10.7 and 10.8 of the Criminal Code Act 1995 (Cth). These offenses include computer intrusions (for example, malicious hacking), the creation and distribution of malicious software (for example, viruses, worms, trojans), and dishonestly obtaining or dealing in personal financial information (including identity theft). The AFPs focus in these areas is twofold:

1. The protection of Australian Government systems
2. The protection of Systems of National Importance, with a key element being the banking and financial sector

CCO works collaboratively with industry and government agencies in this task. It works closely with the Australian Intelligence Community through the Cyber Security Operations Centre (CSOC), and given the frequently transnational nature of such attacks, has developed close collaboration with international law enforcement agencies in order to identify offenders and mitigate the impact of the attacks.

The CSOC is a whole of government approach providing comprehensive understanding of the cyber threat and the security status of government networks and networks of national importance. It also identifies and analyzes sophisticated cyber attacks and provides government with responses and options. It assists responses to cyber events across government and critical private sector systems and infrastructure.

There are many challenges that face investigators of cybercrime. The legal challenges are difficult but, as more prosecutions are put before the courts, more understanding and use of the legislation will evolve. CCO has initiated an education program directed at the legal fraternity, which has been well received.

Dutch National High Tech Crime Unit

The Dutch National High Tech Crime Unit (NHTCU) is a team within the Dutch National Police Agency, dedicated to investigating advanced forms of cybercrime. The team's vision is to make the Netherlands an unsafe place for cybercrime. In addition to Dutch victims and criminals, this includes the use of Dutch infrastructure in criminal activities.

The team specializes in using out-of-the-box investigation methods and techniques to find and target the most important players in the criminal chain. The team has excellent contacts in North America and Western and Eastern Europe, and often plays the role of bridge builder between high tech crime units in different countries.

Another success factor is the advanced cooperation with other public and private partners, where information is freely shared and joint strategies are implemented. An example of such cooperation can be read in the description of the Bredolab case. The NHTCU recently started up the Dutch Electronic Crimes Task Force, a new cooperation with financial and other parties to institutionalize public-private partnership as a means to actively combat certain types of cybercrime.

Irish Reporting & Information Security Service

Ireland's economy has grown rapidly over the past decade thanks in no small part to the increasing use of information technology. Information technology and the Internet have enabled Irish consumers and businesses alike to better access and deliver services, create and access new markets, exchange information rapidly, and process information in more efficient means. However, this increasing reliance on information technology brings with it numerous risks and threats that if not properly addressed could result in significant negative impact on those businesses, individuals, and ultimately Ireland's economy. Just as business and consumers have found legitimate use for the Internet there are others who use the Internet for more nefarious purposes such as criminal, illegal, terrorist, and corporate and national espionage activities.

The Irish Reporting & Information Security Service (IRISSCERT) is Ireland's first Computer Emergency Response Team. IRISSCERT is an independent not-for-profit company limited by guarantee founded in 2008. Our goal is to provide a range of high quality information security-based services to aid Irish-based organizations and citizens to better secure their information technology facilities and services. We achieve this goal by providing alerting services on current and potential information security threats, providing guidance on information security prevention, response, and mitigation strategies, and providing a computer security incident handling and coordination service for national and international organizations.

These services are provided by a dedicated core of experts in the field of information security who provide their services to the Irish Internet community free of charge so that Irish businesses can better protect their information systems and to make the Irish Internet space a safer environment for all.

For more information please refer to https://www.iriss.ie/iriss/RFC_2350.htm or if you wish to report an incident e-mail info@iriss.ie.

Police Central e-Crime Unit

The Police Central e-Crime Unit (PCeU) was founded in September 2008, built around the existing Metropolitan Police Computer Crime Unit at Scotland Yard, to provide a national investigative capability for the most serious cybercrime incidents with the goal *"to provide a safe and secure online and networked computing environment that enhances trust and confidence in the UK as a safe place to live and conduct business."*

Initial successes proved that a fast-time response and co-operation with partners through the establishment of the virtual task force was vitally central to tackling cybercrime and reducing the levels of harm.

Following the British Government's Strategic Defence Security Review, in October 2010, cybercrime has been recognized as a Tier One threat and the National e-Crime Programme and the PCeU have secured increased funding to improve the police response to e-Crime by developing the mainstream capability of the Police Service across England, Wales, and Northern Ireland.

Today the PCeU has a proven track record with many successful operations behind it. It has grown in size and reputation to a dynamic and focused element of the law enforcement approach to combating cybercrime. The PCeU has demonstrated that through timely intervention and partner collaboration it can reduce financial harm to the country, achieving a savings of £140 million between April – October 2011.

February 2012 saw the launch of three regional PCeU hubs as part of the National e-Crime Programme's mandate of developing specialist capability and resources countrywide. This summer the PCeU, working alongside partner agencies, has the responsibility to provide the police response to cyber threats targeting the London Olympic Games.

United States Secret Service

As the original guardian of the nation's financial payment system, the United States Secret Service has established a long history of protecting American consumers, industries, and financial institutions from fraud. Over the last 145 years, our investigative mission and statutory authority have expanded, and today the Secret Service is recognized worldwide for our expertise and innovative approaches to detecting, investigating, and preventing financial and cyber fraud.

Today's global economy has streamlined commerce for both corporations and consumers. Financial institutions and systems are readily accessible worldwide. Today's financial fraud and cybercriminals have adapted to this new means of global trade and seek to exploit this dependence on information technology. Cybercriminals consequently have become experts at stealing stored data, data in transit, and encrypted data. They operate based on trust, long standing criminal relationships, high levels of operational security, and reliability. The culture also has evolved over the last decade and is now described as non-state sponsored, transnational, and is almost impossible to infiltrate due to its dynamic nature and operational security.

To combat these emerging threats, the Secret Service has adopted a multi-faceted approach to aggressively combat cyber and computer-related crimes by establishing a network of 32 Electronic Crimes Task Forces (ECTF), including two international ECTFs located in Rome, Italy and London, England, 38 Financial Crimes Task Forces (FCTF), and a Cyber Investigations Branch. This approach enables the Secret Service to detect, prevent, and aggressively investigate electronic crimes, including cyber attacks on the nation's critical infrastructures and financial payment systems.

For more information or to report a data breach, please contact your local Secret Service office at www.secretservice.gov.

Got a question or a comment about the DBIR?
Drop us a line at dbir@verizon.com, find us on Facebook, or post to Twitter with the hashtag **#dbir**.

2012 DATA BREACH INVESTIGATIONS REPORT

A study conducted by the Verizon RISK Team with cooperation from the Australian Federal Police, Dutch National High Tech Crime Unit, Irish Reporting and Information Security Service, Police Central e-Crime Unit, and United States Secret Service.

verizon.com/enterprise

© 2012 Verizon. All Rights Reserved. MC15244 03/12. The Verizon and Verizon Business names and logos and all other names, logos, and slogans identifying Verizon's products and services are trademarks and service marks or registered trademarks and service marks of Verizon Trademark Services LLC or its affiliates in the United States and/or other countries. All other trademarks and service marks are the property of their respective owners.