Global Information Assurance Certification Paper

The State of Honeypots: Understanding the Use of Honey Technologies Today

GIAC (GCIA) Gold Certification

Author: Andrea Dominguez, andrea.dgz@gmail.com
Advisor: Stephen Northcutt
Accepted: November 2017

Abstract

The aim of this study is to fill in the gaps in data on the real-world use of honey technologies. The goal has also been to better understand information security professionals' views and attitudes towards them. While there is a wealth of academic research in cutting-edge honey technologies, there is a dearth of data related to the practical use of these technologies outside of research laboratories. The data for this research was collected via a survey which was distributed to information security professionals. This research paper includes details on the design of the survey, its distribution, analysis of the results, insights, lessons learned and two appendices: the survey in its entirety and a summary of the data collected.

1. Introduction

Honeypots, systems designed to deceive adversaries in a network, have evolved into many different flavors which have been widely explored and written about. Their levels of interactivity and complexity, return on investment versus their required effort and upkeep, as well as different attempts to automate while still fooling advanced threat actors have been minutely detailed. From honeypots have come the more specialized ideas of honeytokens, honeynets, honeywalls, and a myriad of other related concepts (Smith, 2016). Honeytokens in particular have a wide range of uses, applications, and nomenclatures. Honeycredentials, honeycreds, honeyhashes, canary credentials, honeyproxies, and honeytraps are some of the many terms used to describe these artifacts, which are used to lure and identify threat actors.

There is currently no authoritative study or comprehensive data on the use of honey technologies in real-world applications. By setting out to answer the research question – "How are honey technologies being used by organizations in real-world applications today?" – this research sought out to gather valuable and insightful data for the benefit of the academic and information security communities. This data can provide a glimpse into the use of intrusion detection technologies such as honeypots and honeytokens in network security programs across a variety of organizations. It can also help security professionals to better gauge the interest, use, and understanding of honeypots and other intrusion detection technologies. The hope is that future researchers will be better equipped to ask questions and to find direction for research based on the data collected in this study.

1.1 Honeypots

Honeypots are computer systems designed to lure attackers by simulating real systems within a network (Spitzner, 2003). While these systems appear to be real, they must have no production value. Any interaction with them should, by definition, be illicit. Anyone connecting to or interacting with such a system is suspicious. There are many kinds of honeypots – from low-interaction systems to the high-interaction, complex systems which are designed to attract and lure more advanced attackers (Pallarés, 2016). Honeypots can come in many forms, including endpoints, web servers, databases, and

email servers, among others. They are often used to gather intelligence on attackers, uncover insider threats, and to try to distract and trap attackers before they can reach valuable resources in a network.

1.2 Honeynets

Honeynets are complex honeypot networks designed to simulate a real network (Pouget, 2003). Honeynets connect and interact in the same way a real network would – none of the connections between systems are emulated. This requires a lot more time and resources to set up, but can provide valuable information about how the attacker can move laterally between different devices on a network. One honeynet includes several honeypot instances (Smith, 2016). Due to their complexity, it may also take the attacker a longer time to realize that the systems are all traps within a honeynet.

1.3 Honeytokens

The term "honeytoken" first appeared around 2003 (Spitzner, 2003), and gained traction in more recent years – well over a decade after they were first described in white papers and government intelligence papers. While the concept of honeytokens as we know them today is relatively new, the idea of inserting fake data to catch thieves and bootleggers has been around for a long time. Mapmakers were known to create so-called "paper towns" in their maps to identify whenever a rival mapmaking company stole their maps (Spitzner, 2003). Other forms of fake entries, or copyright traps, have existed in the dictionary, encyclopedia, map and directory industries for a long time – well before the age of computers and the widespread digital distribution of information. Honeytokens are the latest incarnation of this well-established concept.

There are many definitions of honeytokens in existence, but they all include the following three conditions: it must be an object, not a system; it must have no production value; it has no legitimate use and is therefore unauthorized (Pouget, 2003). It can be defined as any resource stored such as a text file, an email message or a database record which would not be accessed for regular production purposes (Grudziecki, 2012). Honeytokens must be unique and highly unlikely to appear in legitimate traffic to avoid false positive alerts (Smith, 2016). They must also be difficult to identify as bait by an adversary (Nicholson, 2015).

2. Research Question and Hypotheses

2.1 Research Question

The question that spurred this research was: "How are honey technologies being used by organizations in real-world applications today?" The idea of surveying as many organizations as possible to gather data came about from a dearth of data on the different applications of these technologies. The goal was to gain a better understanding of the real-world applications of, attitudes towards and knowledge about honey and other deception technologies today.

2.2 Hypotheses

Before starting out, a few hypotheses were formulated about honey technologies and their real-world use. They are as follows:

1. Honey technologies are not widely deployed in production systems.
2. The most commonly used honey technologies are honeypots, followed by honeynets, and finally honeytokens.
3. Honey technologies are viewed as desirable but difficult to implement.
4. Interest in honey technologies is growing (out of scope).

The first hypothesis – that honey technologies are not widely deployed in production systems – came about from the author's personal experience, as well as that of security professionals with which the author has had conversations. The second hypothesis is that within the different categories, the most widely deployed would be honeypots, followed by honeynets, and in third place would be honeytokens. The reasoning behind this hypothesis is that honeypots have been around for longer and are more well-known among security professionals. The other reason is that they have been tested and perfected more due to this advantage.

The third hypothesis is that honey technologies are seen as desirable but tricky to implement. Given the first hypothesis that their use in production systems is limited given their benefits, this third hypothesis would explain this problem. That is, people's perceptions (be it security professionals themselves, or the C-level executives who approve budgets) – not the technology's inherent nature –limits their use.

Finally, the fourth hypothesis is that interest in honey and other deception technologies is growing. This idea is based not on any hard evidence, but on the general buzz in past years throughout the industry. There are a growing number of vendors offering these types of products. Unfortunately, the answer to this question is beyond the scope of this survey, since it requires recurring surveys across time. However, it is the hope of this study that it will be the first of many surveys that will gauge the interest in and use of honey technologies across organizations in the years to come.

3. Research Methods

To answer the research question, a survey was created to gather data about the familiarity with, usage of, and attitude toward honeytoken technologies from a range of organizations. The survey was used to gather as much data as possible to improve the understanding of the use of honeytokens. The goal was to gauge the understanding of and interest in these technologies. The scope was to cast as wide a net as possible across different sectors (academic, non-profit, businesses, and government) for maximum survey response. The research methods as initially devised, and the research methods used in practice, differed – the sections below aim to give a clearer picture of what was planned versus what ended up being the best course of action in practice for the application of honey technology in real-world contexts.

3.1 Sample Calculation

Originally, the number of survey requests was to be determined following this methodology. The number of target responses would be a representative sample size of companies and organizations that manage their network. Additionally, the companies surveyed had to be at least partially in control of managing their network security. Were this not the case, the respondent would have been asked to forward the survey to the entity managing their network security program. Having determined the target number of responses, and assuming an average response rate of 20%, the plan was to send a number of invitations five times the number of expected responses to achieve the representative sample size. The sample exercise below is based on the number of public companies in the U.S. as of 2016 (Mauboussin, 2017).

Sample calculation:

Number of companies in stock exchange (2016): 3,671

Sample size (confidence 95%, margin of error 5%): 348

Response rate assumption: 20%

Number of surveys to be sent out: 1,740

- Sample size calculator (https://www.surveymonkey.com/mp/sample-size-calculator/)

In the end, this proved to be an impractical approach for many reasons. Chief among these is the scale of such an endeavor. As detailed in the sections on the survey distribution methodology, receiving 100 survey responses proved to be a gargantuan task. It would require the concerted effort of a team with that single focus in mind, and several months of careful planning and execution to gather the thousands of responses needed. Furthermore, it would require a marketing budget, as well as some incentive for participants to complete the survey.

An effort of such scale would be best undertaken by a large organization with significant clout and resources. These efforts are likely already being made by companies interested in offering relevant products, without making the results public. This research effort relied entirely on the goodwill of many generous individuals within the information security field. Its goal is to share the results with a diverse audience to further the understanding of these technologies and to spur further academic research.

4. Survey Design and Creation

The survey design process started with a list of questions and their answers, grouped into sections. The advisor then gave feedback, which was incorporated into the design. Feedback from other SANS instructors was gathered as well, which was reflected in the second draft. This time, the main feedback was to cut back on the length. This resulted in a painstaking process of deciding what was absolutely necessary to include and what could be removed. The process of elimination helped limit the scope and made the survey more focused.

Privacy was a focal point from the start. It was part of the design that the answers would be anonymized, and each survey would have a unique ID. The survey was designed to answer the following questions for the organizations surveyed:

1. *What is the organization's engagement with honey technologies?*
2. *Do they actively use honey technologies? (What kind? Since when?)*
3. *Do they plan to use honeytokens in the future? (When? What kind?)*
4. *What is their opinion of their efficacy based on experience?*
5. *What kinds of security controls do they have in place in their network?*
6. *Information about their organization in general (for analytical purposes)*

The survey was created using the survey creation and response-gathering service Survey Monkey. A few other services were tested, but Survey Monkey offered the most intuitive and easy to use design tools for including rules and other advanced features. Additionally, it offered the most competitive price for access to such features. The rules feature was especially important in helping guide participants to questions relevant to them so they did not waste time viewing questions not applicable to their organization. Since the survey had a total of 40 questions, it was deemed important that only questions relevant to them be shown, with the goal of avoiding survey fatigue and maximizing the completion rate. See Appendix A for a detailed view of the survey questions, answer options, sections, and rules.

The plan was to email the link along with a cover letter to information security professional associations, ISACs and other such groups as detailed in the section below. The email included a cover letter explaining the purpose of the survey, the significance of the results, as well an explanation of who should respond within any given organization. The email also included a statement explaining that the results would be anonymized for privacy. The instructions also included contact information for the author and advisor for any questions or concerns.

4.1 Survey Distribution – in Theory

The survey was to be distributed by leveraging different information security and threat intelligence sharing groups such as the SANS Technology Institute, the National Council of ISACs (Information Sharing and Analysis Centers), the NCFTA and others. The survey was to be distributed via targeted emails and the distribution lists of the groups in

the scope of the study. The idea was to maximize the distribution by leveraging groups with many members. This method would be much more efficient than researching many unique organizations, finding their most promising points of contact, and then reaching out to them one by one.

The surveys were to be responded to by the CISO, the CISO's office, or any such equivalent department head within an organization – more specifically, by the department in charge of information security within an organization. The email requested that the person in charge respond. Barring that, it would be requested that a senior individual with privileged access to and comprehensive understanding of the company's network and its previous, existing, and future security controls respond. In cases where an organization had a network but did not manage its own network security, the survey was to be sent on to the entity managing the network security of that company and filled out by a knowledgeable senior individual as detailed above.

4.2 Survey Distribution – in Practice

There were two phases to the survey distribution. The first involved careful research and planning – the author reached out to a total of 48 different ISACs (Information Sharing and Analysis Centers), ISAOs (Information Sharing and Analysis Organizations), and professional organizations to formally request their help in distributing the survey, along with a cover letter, to their members. The second phase came about due to the failure of the first method. It was a more ad hoc and informal method which involved reaching out to security professionals directly both via email and Twitter.

4.2.1 Survey Distribution Phase 1: ISACs

The author researched as many U.S.-based ISACs, ISAOs, and professional organizations as possible. The logic behind this was that each one of these organizations must have a list of members numbering from the dozens to the hundreds – each a unique company or organization that could potentially respond to the survey. A spreadsheet that included the contact emails (or contact form URLs) was created and used to keep track of who had been contacted, when, what the reply was and when it was received. The author

wrote up a formal email request to have each of these organizations distribute the survey to their members.

The email contained an in-depth background of the survey, the research question, as well as a guarantee of privacy to individual respondents. It also listed contact information for both the researcher and advisor. These emails were sent well before the survey was ready, in order to gauge interest and gain permission before sending it out. The email included an approximate date when the survey would be ready for distribution. The timing worked well, as some organizations took over a week to respond, by which time the author could send them the participant-specific write-up and the link to the survey within a few days' time.

4.2.2 Survey Distribution Phase 2: Twitter

Unfortunately, the first method did not yield promising results. Out of the 48 entities contacted, only ten replied. Out of those ten, only four agreed to distribute the survey to their members – more specifically, three agreed to send it out to their members directly, and one allowed the author to post the survey link and write-up on their Linked In group. This represented a much smaller response rate than anticipated (8%) and put the validity of the survey results at risk. A re-evaluation of distribution methods was urgently needed.

As an avid Twitter user, the author follows a good number of security professionals on the platform. The author decided to reach out to as many of them as possible to request their help in distributing the link to the survey to their Twitter followers. The author started with the big names, those with 50,000 or more followers, then moved onto those with fewer than that. Whenever possible, the author wrote an individualized message – a shortened version of the formal cover letter sent to the ISACs in the first phase – via direct messaging. Whenever appropriate, the author mentioned their interest in the security professional's work or research. In some cases, this was a podcast or blog. These were sincere appeals; the author avoided reaching out to people they had not been following on the platform for at least one year.

Whenever a user had direct messaging disabled, the author tweeted at them and added the appropriate tags and hashtags. Surprisingly, this method worked just as well as direct appeals via the more personalized direct messages. Additionally, the author reached out

to individuals directly via their preferred communication – often via email. This method also yielded positive results, although at a more limited scale.

5. The Survey

The survey was live from September 11 to October 11, 2017. During this time, 124 individual responses were received. The survey was designed in such a way as to limit the respondents to those who work at an organization that owns its network and has a network security program in place. The questions were designed, and results analyzed given those two initial assumptions. The survey can be found in its entirety, including questions, answer options and rules in Appendix A.

The survey was comprised of six distinct sections – qualifying questions, security controls, honeypots, honeynets, honeytokens, and demographics. Qualifying questions were designed to keep the survey within scope by weeding out respondents whose organizations did not meet the minimum criteria for participation in the data collection. The security controls section asked participants questions about what security controls were currently in place in their organizations. This question provides important context around which better analysis can be conducted. For example, are organizations with IDS systems in place more or less likely to have honeynets in their network? This question and others can be answered, and interesting correlations can be made.

The third through fifth sections were the focus of the survey. They were the sections addressing familiarity with, use of, and attitude towards honeypots, honeynets, and honeytokens respectively. Before each of these, a qualifying page was included. If the respondent answered in such a way that disqualified them from answering, they were not directed to the page with the questions regarding that technology, but instead skipped that page and moved onto the next section's qualifying page. The qualifying question in each case was "Do you have x in your network?" If the answer was "yes" or "no, but we plan on deploying them in the next 12 months," they would move on to the main questions.

The sixth and final section of the survey was titled "Just a few final questions…" and it was a set of questions on demographic and other topics of interest. The questions ask about the type, industry, and size of the organization, as well as the role of the person responding. The final question is an open-ended request for feedback.

6. Results

The results of the survey included a lot of interesting findings. Among these findings was that not very many of the respondents had honey technologies in their organizations, but many of them were at least familiar with the different concepts related to honeypots. The section below details some of the more significant findings in the different sections of the survey. For per-question results, please see Appendix B for a summary of the results.

6.1 Respondents Metadata

One of the considerations when analyzing the results was to look at the countries and regions of the responses, as a whole, before diving into the results themselves. Since the survey did not directly ask respondents to identify their location or their organization's location, this was gathered via IP address metadata. For privacy reasons, they will not be shared, nor their answers tied to those locations. The vast majority of responses came from the United States (78). The second country with most responses was the United Kingdom (12), and Canada was the third (5). Together, these three countries accounted for over 75% of all responses.

There were responses from a total of 27 different countries. Grouping the rest of the countries into regions, these were the totals from each: Europe (21), Asia (3), Africa (2), Latin America (1), Australia/ New Zealand (1), Middle East (1). The totals are broken down in the two charts below:

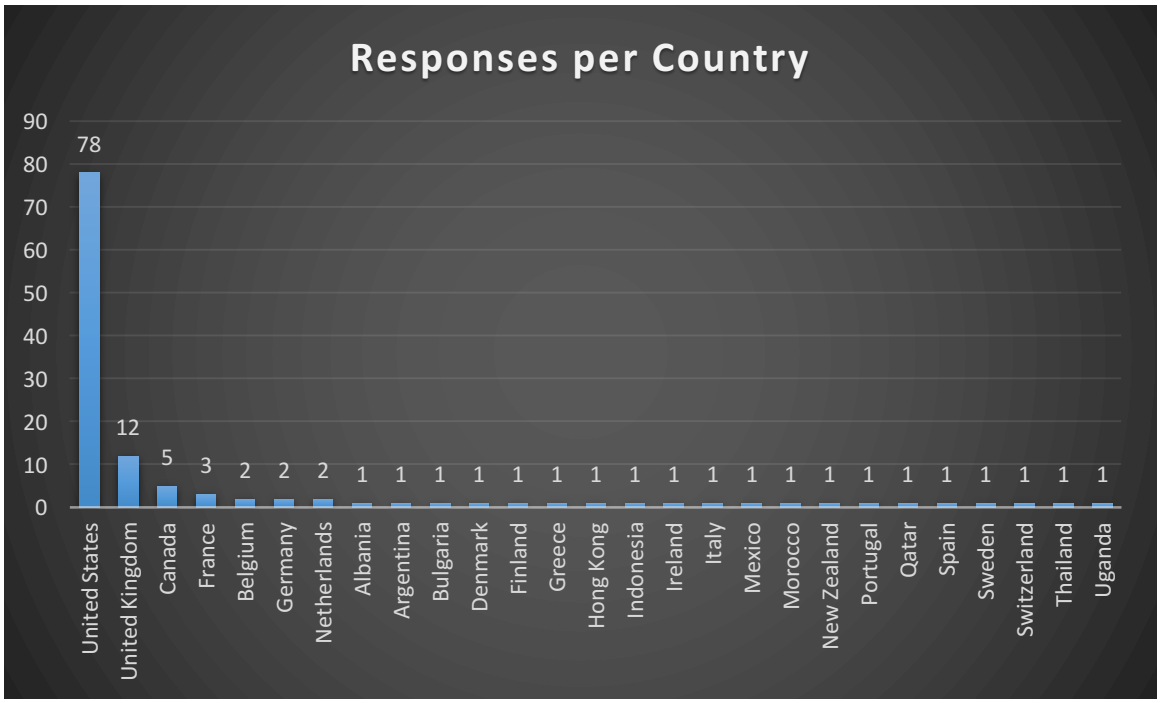

Table 6.1a "Number of Responses Received by Country"

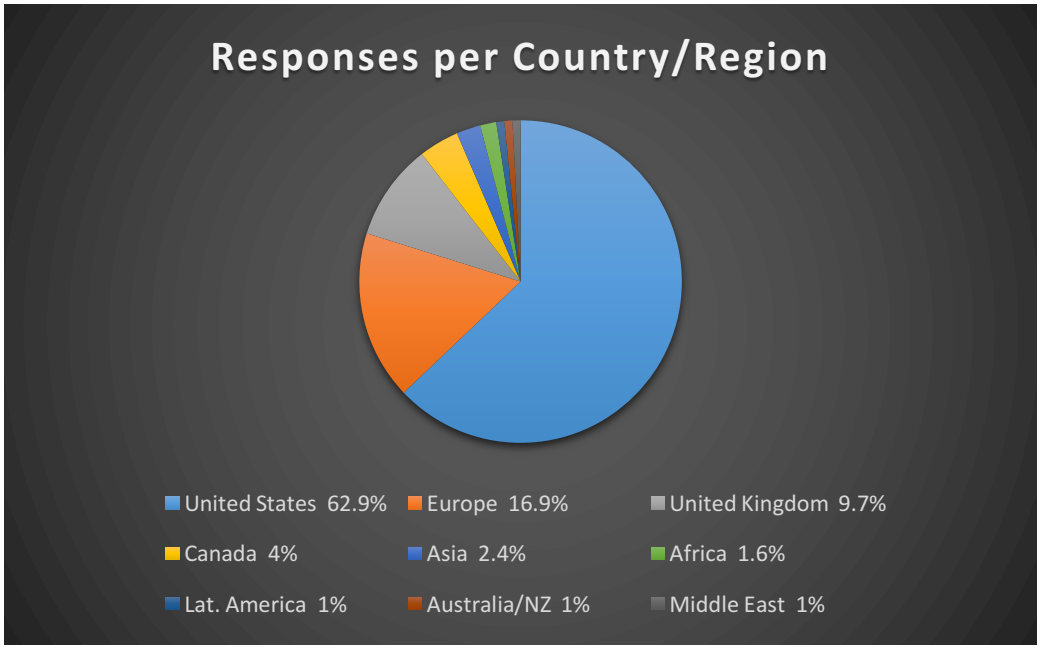

Table 6.1b "Percentage of Total Responses by Country or Region"

Another aspect worth looking at was the breakdown of the responding organizations. Interestingly, the top three organization types were: privately-held corporations, publicly-traded corporations, and government agencies – respectively. Academic institutions, utilities, and non-profits also made an appearance. It's worth noting that 34.6% of

respondents did not identify the type of organization, however, so this data point might not offer extraordinary insights on its own. The chart below shows the organization type breakdown by percentage:

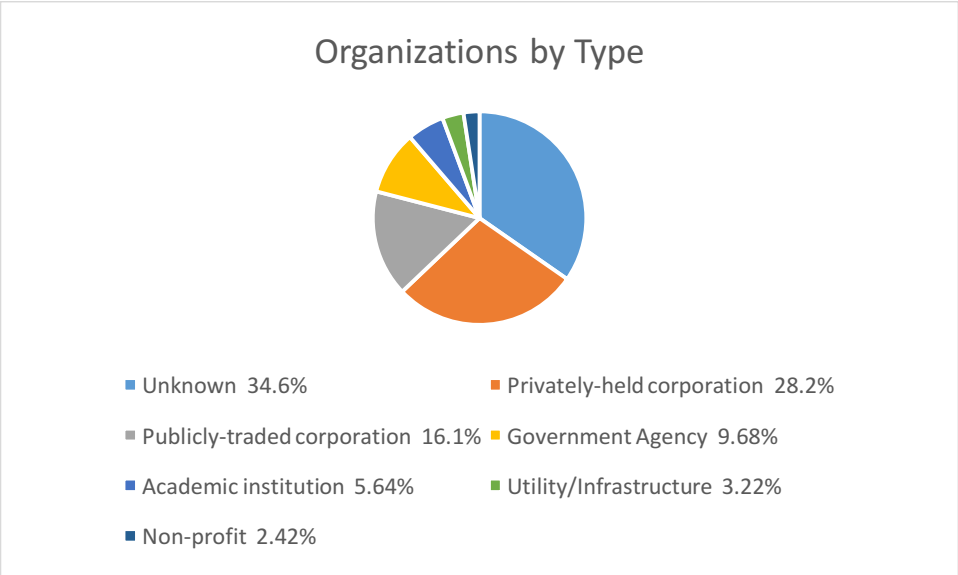

Table 6.1c "Organizations by Type"

One final dimension of the respondents' metadata is the responding individual's role and responsibilities within their organization. Two questions in the demographics section addressed this element/factor. Question 38 asked users to select the answer that most closely reflected their position within their organization. The majority (31%) responded "security administrator." This is a catch-all phrase for a number of roles within an information security department, including analysts, engineers, and incident responders. The second and third most common answers were director-level and IT staff, respectively. This configuration is detailed below in Table 6.1d.

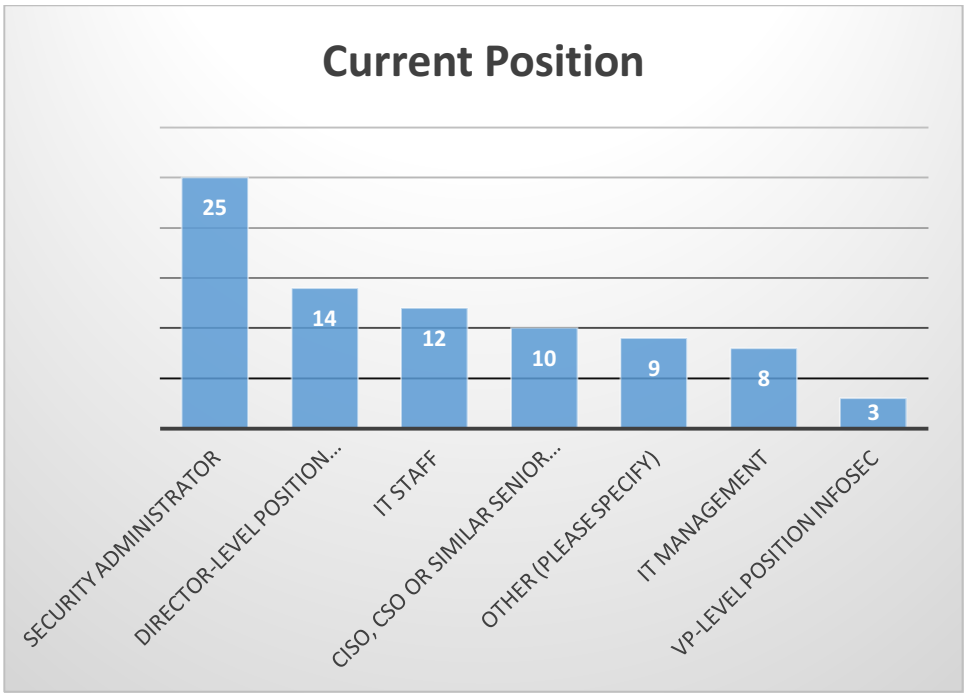

Table 6.1d "Respondents' Self-Reported Positions"

Question 39 asked about the specific responsibilities of the respondent within their organization as far as information security was concerned. The question allowed for participants to select more than one answer and listed 16 different options. The question instructed participants to select just the top two, but the average number of answers selected per person was 3.6. The three most common responsibilities selected were: Security analysis and planning, Incident response, and Network security, respectively. The chart below shows the responsibilities in order from most to least common:

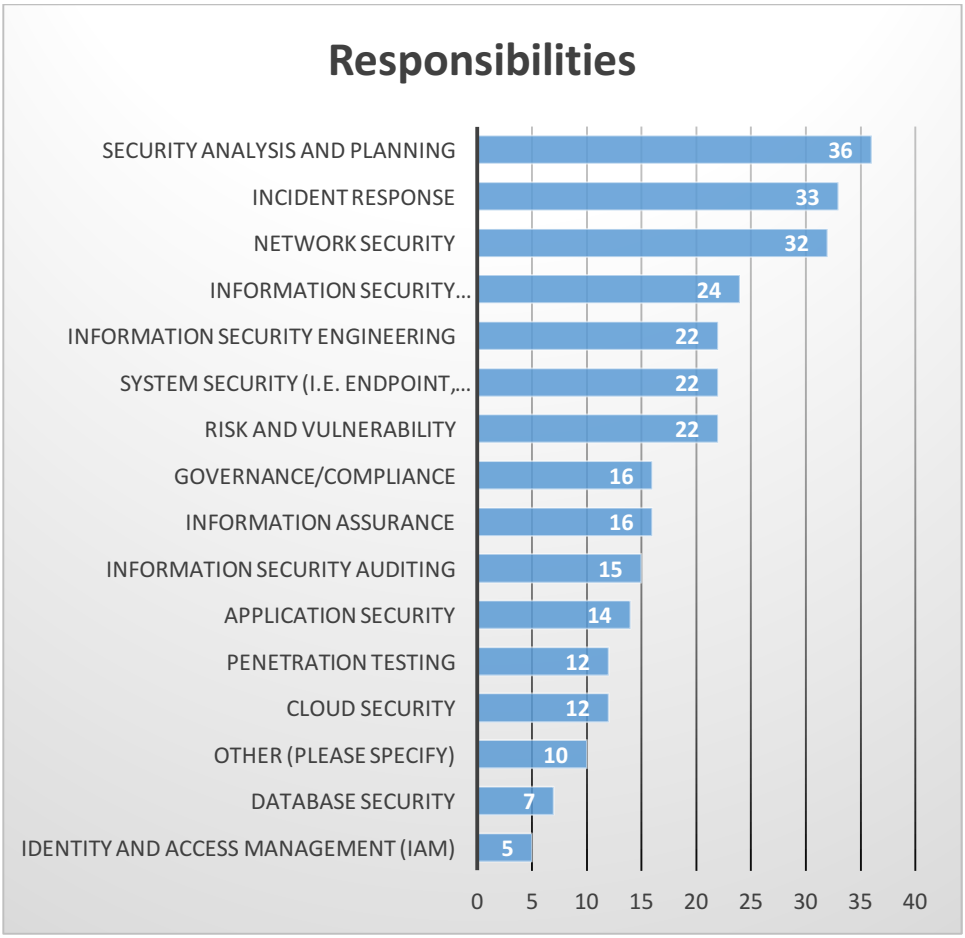

Table 6.1e "Respondents' Self-Reported Responsibilities"

6.2 Qualifying Questions

After an introductory page which gave participants an overview of the survey background, privacy, and contact information, the first page with questions was called "Qualifying Questions." At the top of this page is a message that reads "If your answer to Questions 1 or 3 is 'No' your company is not within the scope of this study. If your answer to Questions 2 is 'No' or to Question 4 is 'Externally,' this may be best answered by someone involved in the management of your network security program."

The very first question in the survey was "Does your organization own its own network?" This question, along with four other qualifying questions, were up front at the very beginning of the survey. If a company does not have a network of its own, is there any value in asking about their experiences with honey technologies? Surely there would be some value in asking individuals' opinions regardless, but for the scope of this survey,

the decision was made to limit the scope to organizations that have a network and at least partially manage their own security controls. If the answer to the first question was "No," the survey participant would be unable to go on to the next page, and were instead shown an exit screen thanking them for their participation. Only ten of those surveyed answered this question – they were all located in the United States.

The second question on this page was "Does your organization manage its own network?" This was not necessarily a disqualifier, but intended more for data-gathering purposes. The third question was a qualifier – it read, "Do you have a network security program in place?" The vast majority of respondents answered: "Yes" (93.6%). The final two questions in this section addressed the question of whether the security program was managed internally, externally, or both – and if both, how much of it was internally managed. Overall, 73% responded that their security program was internally managed, and 25% that it was both. Of those that replied that it was both internally and externally managed, 20% said over 75% of it was internally managed. Only less than 2% responded that their security program is externally managed. The survey suggested that those individuals consider forwarding the link to the entity in charge of managing their security to respond to the survey.

6.3 Security Controls

The second section of the survey covered the security controls currently in use at the participant's organization as well as their knowledge of different security controls. The first question asked respondents to select all of the security controls currently in place in their organization out of a list of 11 options, including "Other," which allowed them to write in their response. Out of the 95 respondents who answered this question, 23 selected "Honeypots/Honeynets/Honeytokens" as one of their answers (24.2%). Almost 99% indicated having a firewall, and 92% answered that they had an anti-virus solution in place. Other options selected by the majority of respondents included: Email filter (88%), IDS/ IPS (75%), Proxy (72%), Endpoint Security (67%).

Questions 8 and 9 asked about deception technologies and the respondent's knowledge of and experience with such technologies. Overall, 42% of respondents rated themselves as "knowledgeable," 36% said they knew about them but had never used them, and 13% categorized themselves as subject matter experts. As far as practical

experience, 33% said they had never participated in their implementation or use, 30% said they had limited experience, and 21% said they had been involved with at least one phase in the implementation/use of these tools. Only 16% responded that they had extensive experience with such technologies.

When asked whether they use any deception storylines to lure and guide attackers, 67% responded that they did not. The other 33% responded that they currently did, or were planning on building such storylines. Finally, when asked to rate the effectiveness of honey technologies on a scale of one to 10 (one being not effective at all and 10 being extremely effective), the average of all responses was 6.42. This shows a slight skepticism towards honey technologies among the security professionals surveyed. The perception should not considered to be entirely negative, but the score is still quite low when averaged across professionals with all levels of experience and engagement with the technologies.

6.4 Honeypots

The third section of the survey covered the use of honeypots at the participant's organization. Before questions on honeypots, there was a simple qualifying question: "Do you have honeypots in your network?" If the answer was "yes" or "no, but we plan on deploying them in the next 12 months," the survey would then lead them to the section on honeypots. Otherwise, the survey would skip to the qualifying question for the honeynets section. A total of 23 respondents said they did have honeypots in their network, while another 16 said they were planning on implementing them within the next 12 months.

The setup for the sections on honeypots, honeynets, and honeytokens was very similar. Of the 39 respondents navigated to the honeypots section, the most common answer to "How long have you had honeypots in your network?" – other than "N/A" was "2-3 years" (17%), followed by "Over three years" (8%). Those that had experience with honeypots rated their effectiveness an average of 7.35 out of ten. When asked about the frequency with which honeypot events are triggered, 27% said "daily," 23% said "weekly," 20% said "a few times per year," and 17% said, "rarely/seldom." When asked how many events had triggered in the past 12 months, surprisingly, the top two responses were "more than 15" (38%) and "none" (21%).

On the question about deployment on internal versus external systems and applications, the vast majority replied "internal" (54%), followed by "both" (40%). Regarding the monitoring mechanisms, 59% responded that both network-based and host-based tools were used, while 21% responded that they used network-based tools only, and 15% responded that they used host-based tools exclusively. These results indicate that the vast majority of implementations include internally deployed honeypots and that in most cases there is more than one monitoring mechanism.

6.5 Honeynets

The fourth section of the survey covered the use of honeynets at the participant's organization. Before questions on honeynets, there was a qualifying question: "Do you have honeynets in your network?" If the answer was "yes" or "no, but we plan on deploying them in the next 12 months," the survey would then lead them to the section on honeynets. Otherwise, the survey would skip to the qualifying question for the honeytokens section. Overall, 76% of respondents (68 in total) answered that they did not have honeynets in their network. Of the remaining 22, 63.6% were planning on implementing them in the next 12 months, while 36.3% said they already had honeynets in their network. Based on these findings, it can be concluded that there is a lot of growth coming to the use of this technology in the coming year.

When asked to rate their effectiveness, the average response given was 7.56 out of ten. When asked which team managed their honeynets, 73% answered that the SOC (Security Operations Center) was in charge, while 47% replied that the network engineering team did. On the frequency of events, the most common answer was "never" (31%), followed by "daily" (23%). In line with that, the question about the number of events in the past 12 months had the top three responses: "none" (42%), "more than 15" (33%), and "4-10" (17%). The inconsistency of responses may indicate that implementations may vary widely across different organizations.

6.6 Honeytokens

The fifth section of the survey covered the use of honeytokens at the participant's organization. Before questions on honeytokens, there was a qualifying question: "Do you have honeytokens in your network?" If the answer was "yes" or "no, but we plan on

deploying them in the next 12 months," the survey would then lead them to the section on honeynets. Otherwise, the survey would skip to the final section which asked about demographics – both of the organization they represent and about their role and responsibilities within that organization.

When asked about whether they had honeytokens in their network, 72% replied "no," while 15% replied "yes" and the remaining 13% replied, "no, but we plan on deploying them within the next 12 months." Of those who had them in their network already, the majority had had them for three months or less. When asked to rate their effectiveness, the average was 7.23 out of 10. The top two teams in charge of managing honeytokens were reported to be the SOC (57%) and the hunting team (43%). The vast majority of respondents (71%) reported that their honeytokens are hosted internally within their networks.

Half of all respondents (50%) claimed that events were triggered "rarely/seldom," while another 19% claimed that they were "never" triggered. This trend might be explained in part by the fact that many of these implementations have been around for only three months or less. Similarly, most respondents claimed that in the past 12 months no events had been triggered at all (56%). The second most common answer was "1-3" which was selected by 19% of respondents. These answers point towards a lack of experience with this technology. Perhaps in the future, these answers will change as honeytokens become more widely adopted. The fact that 13% of respondents were planning on implementing them in their network in the next 12 months may be indicative of an upward trend in their deployment in production systems.

7. Conclusion

7.1 Did the Hypotheses Hold Up?

The hypothesis that honey technologies are not widely deployed in production systems turned out to be correct. Of those surveyed, only 25% had honeypots, 9% had honeynets, and 15% had honeytokens. However, the second hypothesis was only partially correct. The hypothesis was that the most widely deployed would be honeypots followed by honeynets, and finally, the least deployed would be honeytokens. This hypothesis was an educated guess based on the length time these technologies have been around. The

logic went that the longer it's been around, the more commonly it will be seen in production. However, this turned out to be incorrect - honeytokens were more common than honeynets. In hindsight, this makes sense for at least two reasons. One being that honeytokens are much simpler to deploy. The second is that there are many commercially available products to quickly and easily deploy honeytokens.

The third hypothesis was that the lack of familiarity with honey technologies is limiting their use in production systems – rather than an inherent limitation of the technology. While it is difficult to establish causation, especially in a study of this size, there were signs of correlation. Individuals in organizations with no deployments of the various technologies and no plans for immediate deployment rated themselves as less knowledgeable about these technologies than the overall average. The tables below show comparisons of the self-reported knowledge levels of honey technologies of those who do not have them deployed versus the overall average response:

Honeypots	No	All	Difference
Subject matter expert	3.7%	12.6%	-8.9%
Knowledgeable/ Have used	35.0%	42.1%	-7.1%
I know what they are but have never interacted with them	46.0%	35.8%	10.2%
I know about one or more, not in-depth knowledge	11.0%	7.4%	3.6%
This is the first time I hear of them	3.7%	2.1%	1.6%

Table 7.1a "Familiarity with Honey Technologies – Honeypots not deployed vs. Average"

Honeynets	No	All	Difference
Subject matter expert	7.5%	12.6%	-5.1%
Knowledgeable/ Have used	44.0%	42.1%	1.9%
I know what they are but have never interacted with them	37.0%	35.8%	1.2%
I know about one or more, not in-depth knowledge	9.0%	7.4%	1.6%
This is the first time I hear of them	3.0%	2.1%	0.9%

Table 7.1b "Familiarity with Honey Technologies – Honeynets not deployed vs. Average"

Honeytokens	No	All	Difference
Subject matter expert	8.2%	12.6%	-4.4%
Knowledgeable/ Have used	39.0%	42.1%	-3.1%

I know what they are but have never interacted with them	41.0%	35.8%	5.2%
I know about one or more, not in-depth knowledge	8.2%	7.4%	0.8%
This is the first time I hear of them	3.3%	2.1%	1.2%

Table 7.1c "Familiarity with Honey Technologies – Honeytokens not deployed vs. Average"

The fourth and final hypothesis – that interest in honey technologies is on the rise – was not truly in the scope of this study. This is because the nature of this question requires a comparison across time. However, the results did give some insight thanks to the answer option "No, but we plan on deploying in the next 12 months". Through the responses to these questions, we learned that 17.2% of respondents plan on deploying honeypots in the next 12 months, while 16% said so of honeynets and 13% said so of honeytokens. Further studies are needed to explore this question more fully.

7.2 Insights

The most significant insight gained from this survey was concerning attitudes towards honey technologies depending on a respondent's personal experience with them. Perceptions of effectiveness varied whether the respondent had deception technologies in place in their environment or not. The average score for the question "What is your opinion of the effectiveness of honeypots, honeynets, honeytokens and other similar deception technologies – regardless of whether you've deployed them in your environment?" was 6.42 overall. The question asked respondents to answer based on a scale of one to ten (1 not effective at all - 10 extremely effective).

However, when looking at the specific honeypot, honeynet and honeytoken sections, the average scores were 7.35, 7.56, and 7.23 respectively. To have access to answer these questions, the respondents must have replied "yes" or "no, but we plan on deploying them in the next 12 months" to the qualifying question before each corresponding section. The conclusion could be drawn that those who are already familiarized with these honey technologies – either from direct experience or from knowledge about them – are more likely to rate them as effective. As further evidence, all of those who responded that they would be deploying honey technologies within the next 12 months rated themselves as "knowledgeable" or higher on Question 7: "How

knowledgeable are you about deception technologies such as honeypots, honeynets, and honeytokens?"

Honeypots – deploying in next 12 months	
Subject matter expert	12.5%
Knowledgeable/ Have used	37.5%
I know what they are but have never interacted with them	50.0%
I know about one or more, not in-depth knowledge	0%
This is the first time I hear of them	0%

Table 6.2a "Familiarity with Honey Technologies – Deploying Honeypots in 12 months"

Honeynets – deploying in next 12 months	
Subject matter expert	14.0%
Knowledgeable/ Have used	36.0%
I know what they are but have never interacted with them	50.0%
I know about one or more, not in-depth knowledge	0%
This is the first time I hear of them	0%

Table 6.2b "Familiarity with Honey Technologies – Deploying Honeynets in 12 months"

Honeytokens – deploying in next 12 months	
Subject matter expert	0.0%
Knowledgeable/ Have used	45.0%
I know what they are but have never interacted with them	55.0%
I know about one or more, not in-depth knowledge	0%
This is the first time I hear of them	0%

Table 6.2c "Familiarity with Honey Technologies – Deploying Honeytokens in 12 months."

7.3 Lessons Learned

Thanks to the generous feedback of participants, as well as the benefit of hindsight, a list was compiled of improvements for future versions of the survey to better capture relevant data. One of these is gathering data about the use of commercial products. Early on in the design phase of this survey, the author made a conscious choice to limit its

scope and purposefully left out any questions about specific commercial products and their use.

Part of the reasoning was to limit the length – at 40 questions in total, the survey was already quite long. The decision was made alongside the advisors keep it as short as possible. The author thought it more important to gain an initial understanding of the extent to which these technologies are used (or not) before delving into specifics. For the next iteration, these questions should be asked; specifically, two questions: "Is your implementation of x that of a commercial product, an in-house product, or a combination of both?" followed by "If commercial, please name the vendor and product used."

Several changes to the final section on demographics might help future research capture relevant data on respondents and their organizations more comprehensively. These improvements were identified thanks to respondents who selected "Other" and wrote in an answer. One is on Question 36, asking about the type of organization. The next iteration should include a category for "Law enforcement," and "Publishing" should be changed to "Media/Publishing". Relatedly, there should be a question asking the respondent to identify which country and region their organization is headquartered. This additional question would help offer better insight into the countries of operation of these organizations. The use of the IP addresses of correspondents to discover their location can easily be inaccurate if a person is using a VPN service, traveling, or if they are remote workers not located in the same country or region as the headquarters.

In the respondent-specific questions there were a few improvements that can be made as well. Question 38, which asks about the respondent's position should have two separate answers for C-level executive and other management roles. It should also list a more accurate option instead of simply "security administrator." One option might be to reword that option to "security analyst/engineer" and include "systems or security architect" and "threat intelligence/hunting" as two additional options.

7.4 Future Research

The aim of this research is to spur future research. A promising start would be a continued effort to gather data on the usage of honey technologies, as well as security professionals' knowledge about and attitude towards them. Gathering data across time could help academic researchers determine trends across different points in time. Case

studies of individual companies or industries would also greatly expand the dataset on these technologies' usage. Most importantly, these efforts yield the most benefit when they are timely – meaning, ongoing and consistent research is likely to be the most beneficial.

In the spirit of the free exchange of data and academic collaboration, the author is open to sharing the per-respondent results, stripped of identifying metadata such as timestamps and IP addresses. Sharing this dataset might be helpful to any researcher looking to identify any further correlations in the dataset not included in this paper. Please contact the author for more information.

References

1. Akiyama, M., Yagi, T., Hariu, T., & Kadobayashi, Y. (2017). HoneyCirculator: distributing credential honeytoken for introspection of web-based attack cycle. International Journal of Information Security. doi:10.1007/s10207-017-0361-5

2. Bowen, B. M., Kemerlis, V. P., Prabhu, P., Keromytis, A. D., & Stolfo, S. J. (2012). A system for generating and injecting indistinguishable network decoys. Journal of Computer Security, 20(2-3), 199-221. doi:10.3233/jcs-2011-0439

3. Chovancová, E., Ádám, N., Baláž, A., Pietriková, E., Feciľak, P., Šimoňák, S., & Chovanec, M. (2017). Securing Distributed Computer Systems Using an Advanced Sophisticated Hybrid Honeypot Technology. Computing and Informatics, 36(1), 113-139. doi:10.4149/cai_2017_1_113

4. Gandotra, V., Singhal, A., & Bedi, P. (2012). Threat-Oriented Security Framework: A Proactive Approach in Threat Management. Procedia Technology, 4, 487-494. doi:10.1016/j.protcy.2012.05.078

5. Jurado Pallarés, D. (2016). Análisis y estudio de Honeypots Complejos: Honeynets. UAM. Departamento de Ingeniería Informática. Retrieved from https://repositorio.uam.es/handle/10486/676952.

6. Lazarov, M., Onaolapo, J., & Stringhini, G. (2016). Honey Sheets: What Happens to Leaked Google Spreadsheets? [Dataset]. UCL Computer Science: London, UK. USENIX Security Symposium. doi:10.14324/000.ds.1502241

7. Mauboussin, M. J., Callahan, D., & Majd, D. (2017). The Incredible Shrinking Universe of Stocks: The Causes and Consequences of Fewer U.S. Equities (pp. 1-29, Publication). doi:16. http://www.cmgwealth.com/wp-content/uploads/2017/03/document_1072753661.pdf

Andrea Dominguez, andrea.dgz@gmail.com

8. Maybury, M., Chase, P., & Cheikes, B. (2005). Analysis and Detection of Malicious Insiders. 2005 International Conference on Intelligence Analysis. Retrieved from https://www.mitre.org/sites/default/files/pdf/05_0207.pdf

9. Nicholson, A., Janicke, H., Watson, T., & Smith, R. (2015). Rolling the Dice - Deceptive Authentication for Attack Attribution. ICCWS 2015 - The Proceedings of the 10th International Conference on Cyber Warfare and Security, 223-231. Retrieved from https://books.google.com/books?id=piikBwAAQBAJ&lpg=PA223&ots=EWzUpkyAcy&dq=rolling%20the%20dice%20-%20deceptive%20authentication&pg=PA223#v=onepage&q=rolling%20the%20dice%20-%20deceptive%20authentication&f=false.

10. Onaolapo, J., Mariconti, E., & Stringhini, G. (2016). What Happens After You Are Pwnd: Understanding The Use Of Leaked Account Credentials In The Wild. Proceedings of the 2016 ACM on Internet Measurement Conference - IMC 16. doi:10.1145/2987443.2987475

11. Pouget, F., Dacier, M., & Debar, H. (2003). "Honeypot, Honeynet, Honeytoken: Terminological issues. 1-26. Retrieved May 25, 2017, from http://www.eurecom.fr/en/publication/1275/download/ce-pougfa-030914b.pdf Institut Eurécom

12. Sanders, M. E. (2015). Unknown Threat Detection with Honeypot Ensemble Analysis Using Big Data Security Architecture. Illinois State University ISU ReD: Research and eData, Paper 360. Retrieved from http://ir.library.illinoisstate.edu/etd/360/

13. Shabtai, A., Bercovitch, M., Rokach, L., & Elovici, Y. (2014). Optimizing Data Misuse Detection. ACM Transactions on Knowledge Discovery from Data, 8(3), 1-23. doi:10.1145/2611520

14. Smith, S. D. (2016). Catching Flies: A Guide to the Various Flavors of Honeypots. SANS Reading Room. Retrieved June 21, 2017, from https://www.sans.org/reading-room/whitepapers/attacking/catching-flies-guide-flavors-honeypots-36897.

15. Spitzner, L. (2003). Honeypots: catching the insider threat. 19th Annual Computer Security Applications Conference, 2003. Proceedings. doi:10.1109/csac.2003.1254322

16. Spitzner, L. (2003). Honeytokens: The Other Honeypot. Symantec. Retrieved May 25, 2017, from https://www.symantec.com/connect/articles/honeytokens-other-honeypot

Appendix A: Full Survey

Survey Design – Question, Answers and Rules

Note: each section corresponds to one page of the survey, which must be submitted before the participant is taken to the next section. This applies for subsections as well – meaning that Sections 3a and 3b are on separate pages.

Honeypot Survey

Section 1: Qualifiers

Q1. Does your organization own its own network?
Yes
No
Rule: If "No," survey ends once Section 1 page is submitted.

Q2. Does your organization manage its own network?
Yes
No
Partially

Q3. Do you have a network security program in place?
Yes
No
Rule: If "No," survey ends once Section 1 page is submitted.

Q4. Is your network security program internally or externally managed?
Internally
Externally
Both

Q5. If it is managed both internally and externally, please describe how much of it is handled internally vs. externally:
Under 25% internally
25% internally
50% internally
75% internally
Over 75% internally
N/A

Section 2: Security Controls

Q6. Which security controls do you have in your network? (Select all that apply)
Firewall
Proxy
IDS
IPS
Endpoint Security (HIDS/HIPS)
Anti-Virus
Email Filter
Network IDS/IPS
Data Loss Prevention
Honeypots/Honeynets/Honeytokens
Other (please specify)

Q7. How knowledgeable are you about deception technologies such as honeypots, honeynets, and honeytokens?
Subject matter expert
Knowledgeable/ Have used
I know what they are but have never interacted with them
I know about one or more, not in-depth knowledge
This is the first time I hear of them

Q8. What is your experience with deception technologies such as honeypots, honeynets, and honeytokens?
Extensive – I've helped design and deploy several times
Medium – I've been involved one of the phases of design/implementation
Limited – I've used; not involved in design/implementation
None – I've never participated in their implementation or use

Q9. What is your opinion of the effectiveness of honeypots, honeynets, honeytokens and other similar deception technologies – regardless of whether you've deployed them in your environment? (1 not effective at all - 10 extremely effective)
1 – 10

Q10. Do you build "deception storylines" to guide attackers in a particular direction?
Yes
No
No, but we plan on building them

Section 3a: Honeypots Qualifier

Q11. Do you have honeypots in your network?
Yes
No
No, but we plan on deploying them within the next 12 months
Rule: If "No," survey skips to Section 4a once Section 3a page is submitted.

Section 3b: Honeypots

Q12. How long have you had honeypots in your network?
Under three months
3 – 6 months
6 – 12 months
1 – 2 years
2 – 3 years
Over three years
N/A

Q13. If you have honeypots in your network, how would you rate their effectiveness based on your experience? (1 - not effective at all, 10 - extremely effective)
Answer Choices
1 – 10

Q14. Do you deploy honeypots in public facing applications, internal applications & systems, or both?
Public-facing only
Internal applications only
Both

Q15. Are honeypots monitored on the host or via network-based tools?
Host-based monitoring
Network-based monitoring
Both
Other (please specify)

Q16. Are honeypots managed and deployed by a hunt team, a SOC team or someone else in the organization?
SOC
Network Engineering
Hunting Team
Other (please specify)

Q17. How often are honeypot events triggered on average?
Daily
Weekly
Monthly
A few times a year
Rarely/Seldom
Never

Q18. Approximately how many honeypot events have been triggered in the past 12 months?
None

1 – 3
4 – 10
10 – 15
More than 15

Section 4a: Honeynets Qualifier

Q19. Do you have honeynets in your network?
Yes
No
No, but we plan on deploying them within the next 12 months
Rule: If "No," survey skips to Section 5a once Section 4a page is submitted.

Section 4b: Honeynets

Q20. How long have you had honeynets in your network?
Under three months
3 – 6 months
6 – 12 months
1 – 2 years
2 – 3 years
Over three years
N/A

Q21. If you have honeynets in your network, how would you rate their effectiveness based on your experience? (1 - not effective at all, 10 - extremely effective)
1 – 10

Q22. Do you deploy honeynets in public facing applications, internal applications & systems, or both?
Public-facing only
Internal applications only
Both

Q23. Are honeynets managed and deployed by a hunt team, a SOC team or someone else in the organization?
SOC
Network Engineering
Hunting Team
Other (please specify)

Q24. How often are honeynet events triggered on average?
Daily
Weekly
Monthly
A few times a year

Rarely/Seldom
Never

Q25. Approximately how many honeynet events have been triggered in the past 12 months?
None
1 – 3
4 – 10
10 – 15
More than 15

Section 5a: Honeytokens Qualifier

Q26. Do you have honeytokens in your network?
Yes
No
No, but we plan on deploying them within the next 12 months
Rule: If "No," survey skips to Section 6 once Section 5a page is submitted.

Section 5b: Honeytokens

Q27. How long have you had honeytokens in your network?
Under three months
3 – 6 months
6 – 12 months
1 – 2 years
2 – 3 years
Over three years
N/A

Q28. If you have honeytokens in your network, how would you rate their effectiveness based on your experience? (1 - not effective at all, 10 - extremely effective)
1- 10

Q29. Do you deploy honeytokens in public facing applications, internal applications & systems, or both?
Public-facing only
Internal applications only
Both

Q30. Are honeytokens monitored on the host or via network-based tools?
Host-based monitoring
Network-based monitoring
Both
Other (please specify)

Q31. Are honeytokens managed and deployed by a hunt team, a SOC team or someone else in the organization?
SOC
Network Engineering
Hunting Team
Other (please specify)

Q32. Do you use externally hosted systems to deploy honeytokens or are these systems hosted in your infrastructure?
External
Internal
Both

Q33. How often are honeytoken events triggered on average?
Daily
Weekly
Monthly
A few times a year
Rarely/Seldom
Never

Q34. Approximately how many honeytoken events have been triggered in the past 12 months?
None
1 – 3
4 – 10
10 – 15
More than 15

Section 6: Demographics

Q35. Please select the type of organization you're responding about.
Publicly-traded Corporation
Privately-held Corporation
Government Agency
Non-profit
Academic Institution
Research/Think Tank
Utility/Infrastructure
Other (please specify)

Q36. Please describe the industry your organization belongs to:
Administrative Services
Agriculture
Construction
Courier/Messenger Services

Education/Research
Entertainment
Environmental/Waste Management
Finance/Banking
Food services
Forestry/Logging
Gaming
Government (state/local)
Government (national)
Healthcare
Information Technology
Insurance
Legal
Leisure/Hospitality/Tourism
Life Sciences (biotech, etc.)
Manufacturing
Natural resources/Mining
Non-Profit
Publishing
Real Estate
Retail/Wholesale
Sports
Telecommunications
Transportation/Logistics
Utilities
Other (please specify)

Q37. How many total employees does your organization have worldwide?
Fewer than 100
100 – 200
200 – 500
500 – 1,000
1,000 – 2,999
3,000 – 9,999
10,000 – 19,999
20,000 or more

Q38. Which of the following best describes your current position within your organization? (Please check one)
CISO, CSO or similar senior cybersecurity position
VP-level position in an information security department
Director-level position in an information security department
Security administrator
IT management
IT staff
Non-IT business manager

Other (please specify)

Q39. Which of the following most closely describes your primary responsibilities within your organization? (Please check up to two)
Information assurance
Cloud security
Risk and vulnerability
Identity and access management (IAM)
Security analysis and planning
Network security
Information security auditing
Governance/compliance
Incident response
Database security
Information security analyst/investigator
Penetration testing
System security (i.e., endpoint, server, virtual server, etc.)
Information security engineering
Application security
Other (please specify)

Q40. Last question! Do you have any comments or feedback? [Optional]
(Open-ended)

Appendix B: Full Survey Results (by Question)

Honeypot Survey

Q1. Does your organization own its own network?

Answer Choices	Responses	
Yes	91.94%	114
No	8.06%	10
	Answered	**124**
	Skipped	**0**

Q2. Does your organization manage its own network?

Answer Choices	Responses	
Yes	85.48%	106
No	4.84%	6
Partially	9.68%	12
	Answered	**124**
	Skipped	**0**

Q3. Do you have a network security program in place?

Answer Choices	Responses	
Yes	93.55%	116
No	6.45%	8
	Answered	**124**
	Skipped	**0**

Q4. Is your network security program internally or externally managed?

Answer Choices	Responses	
Internally	72.95%	89
Externally	1.64%	2
Both	25.41%	31
	Answered	**122**
	Skipped	**2**

Q5. If it is managed both internally and externally, please describe how much of it is handled internally vs. externally:

Answer Choices	Responses	
Under 25% internally	1.68%	2
25% internally	4.20%	5
50% internally	5.04%	6
75% internally	10.92%	13

Over 75% internally	20.17%	24
N/A	57.98%	69
	Answered	**119**
	Skipped	**5**

Q6. Which security controls do you have in your network? (select all that apply)

Answer Choices	Responses	
Firewall	98.95%	94
Proxy	71.58%	68
IDS	66.32%	63
IPS	62.11%	59
Endpoint Security (HIDS/HIPS)	67.37%	64
Anti-Virus	92.63%	88
Email Filter	88.42%	84
Network IDS/IPS	74.74%	71
Data Loss Prevention	54.74%	52
Honeypots/Honeynets/Honeytokens	24.21%	23
Other (please specify)	9.47%	9
	Answered	**95**
	Skipped	**29**

Q7. How knowledgeable are you about deception technologies such as honeypots, honeynets, and honeytokens?

Answer Choices	Responses	
Subject matter expert	12.63%	12
Knowledgeable/ Have used	42.11%	40
I know what they are but have never interacted with them	35.79%	34
I know about one or more, not in-depth knowledge	7.37%	7
This is the first time I hear of them	2.11%	2
	Answered	**95**
	Skipped	**29**

Q8. What is your experience with deception technologies such as honeypots, honeynets, and honeytokens?

Answer Choices	Responses	
Extensive – I've helped design and deploy several times	15.79%	15
Medium – I've been involved one of the phases of design/implementation	21.05%	20
Limited – I've used; not involved in design/implementation	30.53%	29
None – I've never participated in their implementation or use	32.63%	31
	Answered	**95**

Skipped	**29**

Q9. What is your opinion of the effectiveness of honeypots, honeynets, honeytokens and other similar deception technologies – regardless of whether you've deployed them in your environment? (1 not effective at all - 10 extremely effective)

Answer Choices	Average Number	Total Number
(no label)	6.418	584
	Answered	**91**
	Skipped	**33**

Q10. Do you build "deception storylines" to guide attackers in a particular direction?

Answer Choices	Responses	
Yes	13.83%	13
No	67.02%	63
No, but we plan on building them	19.15%	18
	Answered	**94**
	Skipped	**30**

Q11. Do you have honeypots in your network?

Answer Choices	Responses	
Yes	24.73%	23
No	58.06%	54
No, but we plan on deploying them within the next 12 months	17.20%	16
	Answered	**93**
	Skipped	**31**

Q12. How long have you had honeypots in your network?

Answer Choices	Responses	
Under three months	5.56%	2
3 – 6 months	5.56%	2
6 – 12 months	5.56%	2
1 – 2 years	5.56%	2
2 – 3 years	16.67%	6
Over three years	8.33%	3
N/A	52.78%	19
	Answered	**36**
	Skipped	**88**

Q13. If you have honeypots in your network, how would you rate their effectiveness based on your experience? (1 - not effective at all, 10 - extremely effective)

Answer Choices	Average Number	Total Number
(no label)	7.346	191
	Answered	**26**
	Skipped	**98**

Q14. Do you deploy honeypots in public facing applications, internal applications & systems, or both?

Answer Choices	Responses	
Public-facing only	5.71%	2
Internal applications only	40.00%	14
Both	54.29%	19
	Answered	**35**
	Skipped	**89**

Q15. Are honeypots monitored on the host or via network-based tools?

Answer Choices	Responses	
Host-based monitoring	14.71%	5
Network-based monitoring	20.59%	7
Both	58.82%	20
Other (please specify)	5.88%	2
	Answered	**34**
	Skipped	**90**

Q16. Are honeypots managed and deployed by a hunt team, a SOC team or someone else in the organization?

Answer Choices	Responses	
SOC	54.29%	19
Network Engineering	28.57%	10
Hunting Team	25.71%	9
Other (please specify)	8.57%	3
	Answered	**35**
	Skipped	**89**

Q17. How often are honeypot events triggered on average?

Answer Choices	Responses	
Daily	26.67%	8
Weekly	23.33%	7
Monthly	6.67%	2
A few times a year	20.00%	6
Rarely/Seldom	16.67%	5
Never	6.67%	2
	Answered	**30**
	Skipped	**94**

Q18. Approximately how many honeypot events have been triggered in the past 12 months?

Answer Choices	Responses	
None	20.69%	6
1 – 3	17.24%	5
4 – 10	17.24%	5
10 – 15	6.90%	2
More than 15	37.93%	11
	Answered	**29**
	Skipped	**95**

Q19. Do you have honeynets in your network?

Answer Choices	Responses	
Yes	8.89%	8
No	75.56%	68
No, but we plan on deploying them within the next 12 months	15.56%	14
	Answered	**90**
	Skipped	**34**

Q20. How long have you had honeynets in your network?

Answer Choices	Responses	
Under three months	12.50%	2
3 – 6 months	12.50%	2
6 – 12 months	0.00%	0
1 – 2 years	6.25%	1
2 – 3 years	0.00%	0
Over three years	12.50%	2
N/A	56.25%	9
	Answered	**16**
	Skipped	**108**

Q21. If you have honeynets in your network, how would you rate their effectiveness based on your experience? (1 - not effective at all, 10 - extremely effective)

Answer Choices	Average Number	Total Number
(no label)	7.556	68
	Answered	**9**
	Skipped	**115**

Q22. Do you deploy honeynets in public facing applications, internal applications & systems, or both?

Answer Choices	Responses	
Public-facing only	15.38%	2
Internal applications only	46.15%	6
Both	38.46%	5
	Answered	**13**
	Skipped	**111**

Q23. Are honeynets managed and deployed by a hunt team, a SOC team or someone else in the organization?

Answer Choices	Responses	
SOC	73.33%	11
Network Engineering	46.67%	7
Hunting Team	13.33%	2
Other (please specify)	6.67%	1
	Answered	**15**
	Skipped	**109**

Q24. How often are honeynet events triggered on average?

Answer Choices	Responses	
Daily	23.08%	3
Weekly	15.38%	2
Monthly	0.00%	0
A few times a year	15.38%	2
Rarely/Seldom	15.38%	2
Never	30.77%	4
	Answered	**13**
	Skipped	**111**

Q25. Approximately how many honeynet events have been triggered in the past 12 months?

Answer Choices	Responses	
None	41.67%	5
1 – 3	8.33%	1
4 – 10	16.67%	2
10 – 15	0.00%	0
More than 15	33.33%	4
	Answered	**12**
	Skipped	**112**

Q26. Do you have honeytokens in your network?

Answer Choices	Responses	
Yes	15.29%	13
No	71.76%	61
No, but we plan on deploying them within the next 12 months	12.94%	11
	Answered	**85**
	Skipped	**39**

Q27. How long have you had honeytokens in your network?

Answer Choices	Responses	
Under three months	14.29%	3
3 – 6 months	0.00%	0
6 – 12 months	9.52%	2
1 – 2 years	9.52%	2
2 – 3 years	4.76%	1
Over three years	4.76%	1
N/A	57.14%	12
	Answered	**21**
	Skipped	**103**

Q28. If you have honeytokens in your network, how would you rate their effectiveness based on your experience? (1 - not effective at all, 10 - extremely effective)

Answer Choices	Average Number	Total Number
(no label)	7.231	94
	Answered	**13**
	Skipped	**111**

Q29. Do you deploy honeytokens in public facing applications, internal applications & systems, or both?

Answer Choices	Responses

Public-facing only	4.76%	1
Internal applications only	57.14%	12
Both	38.10%	8
	Answered	**21**
	Skipped	**103**

Q30. Are honeytokens monitored on the host or via network-based tools?

Answer Choices	Responses	
Host-based monitoring	9.52%	2
Network-based monitoring	33.33%	7
Both	52.38%	11
Other (please specify)	4.76%	1
	Answered	**21**
	Skipped	**103**

Q31. Are honeytokens managed and deployed by a hunt team, a SOC team or someone else in the organization?

Answer Choices	Responses	
SOC	57.14%	12
Network Engineering	33.33%	7
Hunting Team	42.86%	9
Other (please specify)	4.76%	1
	Answered	**21**
	Skipped	**103**

Q32. Do you use externally hosted systems to deploy honeytokens or are these systems hosted in your infrastructure?

Answer Choices	Responses	
External	4.76%	1
Internal	71.43%	15
Both	23.81%	5
	Answered	**21**
	Skipped	**103**

Q33. How often are honeytoken events triggered on average?

Answer Choices	Responses	
Daily	6.25%	1
Weekly	0.00%	0
Monthly	6.25%	1
A few times a year	18.75%	3

Rarely/Seldom	50.00%	8
Never	18.75%	3
	Answered	**16**
	Skipped	**108**

Q34. Approximately how many honeytoken events have been triggered in the past 12 months?

Answer Choices	Responses	
None	56.25%	9
1 – 3	18.75%	3
4 – 10	12.50%	2
10 – 15	0.00%	0
More than 15	12.50%	2
	Answered	**16**
	Skipped	**108**

Q35. Please select the type of organization you're responding about.

Answer Choices	Responses	
Publicly-traded corporation	24.69%	20
Privately-held corporation	43.21%	35
Government Agency	14.81%	12
Non-profit	3.70%	3
Academic institution	8.64%	7
Research/Think Tank	0.00%	0
Utility/Infrastructure	4.94%	4
Other (please specify)	0.00%	0
	Answered	**81**
	Skipped	**43**

Q36. Please describe the industry your organization belongs to:

Answer Choices	Responses	
Administrative Services	2.53%	2
Agriculture	0.00%	0
Construction	1.27%	1
Courier/Messenger Services	0.00%	0
Education/Research	11.39%	9
Entertainment	2.53%	2
Environmental/Waste Management	0.00%	0
Finance/Banking	13.92%	11
Food services	0.00%	0
Forestry/Logging	0.00%	0
Gaming	0.00%	0

Government (state/local)	3.80%	3
Government (national)	5.06%	4
Healthcare	2.53%	2
Information Technology	26.58%	21
Insurance	3.80%	3
Legal	3.80%	3
Leisure/Hospitality/Tourism	0.00%	0
Life Sciences (biotech, etc.)	0.00%	0
Manufacturing	3.80%	3
Natural resources/Mining	1.27%	1
Non-Profit	0.00%	0
Publishing	0.00%	0
Real Estate	1.27%	1
Retail/Wholesale	3.80%	3
Sports	0.00%	0
Telecommunications	2.53%	2
Transportation/Logistics	0.00%	0
Utilities	3.80%	3
Other (please specify)	6.33%	5
	Answered	**79**
	Skipped	**45**

Q37. How many total employees does your organization have worldwide?

Answer Choices	Responses	
Fewer than 100	13.58%	11
100 – 200	4.94%	4
200 – 500	9.88%	8
500 – 1,000	8.64%	7
1,000 – 2,999	16.05%	13
3,000 – 9,999	13.58%	11
10,000 – 19,999	14.81%	12
20,000 or more	18.52%	15
	Answered	**81**
	Skipped	**43**

Q38. Which of the following best describes your current position within your organization? (Please check one)

Answer Choices	Responses	
CISO, CSO or similar senior cybersecurity position	12.35%	10
VP-level position in an information security department	3.70%	3
Director-level position in an information security department	17.28%	14

Security administrator	30.86%	25
IT management	9.88%	8
IT staff	14.81%	12
Non-IT business manager	0.00%	0
Other (please specify)	11.11%	9
	Answered	**81**
	Skipped	**43**

Q39. Which of the following most closely describes your primary responsibilities within your organization? (Please check up to two)

Answer Choices	Responses	
Information assurance	19.75%	16
Cloud security	14.81%	12
Risk and vulnerability	27.16%	22
Identity and access management (IAM)	6.17%	5
Security analysis and planning	44.44%	36
Network security	39.51%	32
Information security auditing	18.52%	15
Governance/compliance	19.75%	16
Incident response	40.74%	33
Database security	8.64%	7
Information security analyst/investigator	29.63%	24
Penetration testing	14.81%	12
System security (i.e., endpoint, server, virtual server, etc.)	27.16%	22
Information security engineering	27.16%	22
Application security	17.28%	14
Other (please specify)	12.35%	10
	Answered	**81**
	Skipped	**43**

Q40. Last question! Do you have any comments or feedback?

Answered	**12**
Skipped	**112**