Toward a Taxonomy of Secrets

[redacted] (b)(3)-P.L. 86-36

Introduction

A great deal of the effort expended in the realm of information security is devoted to the protection of secrets. In the context of this paper, we will use the most general meaning of term – a piece of information that one individual or organization wishes to prevent another individual or organization knowing. This is a definition that is explicitly intended to be far broader than the association of the term "Secret" with traditional hierarchical military/governmental classification schemes. In some cases, the definition of "knowing" may include "being aware of the existence of."

While most of the literature on information security accepts the existence of secrets as a given and proceeds to develop various solutions for the preservation and handling of secrets, it is our intent to step back from this and consider briefly the basic nature of secrets. We will explore the various types of secrets and the different motivations that lead to the creation and keeping of secrets. The information presented on this topic may be viewed as a rudimentary taxonomy of secrets, which may be fleshed out further if it is deemed useful by the community. We will also consider the implications that might be drawn about how best to deal with secrets. In all cases where we refer to the keeper of the secret, it should be assumed that this term may refer to either an individual or an organization, such as a corporation, religious body, or government, unless other distinctions are explicitly made.

This paper will work within a multidimensional categorization structure. Each dimension will be discussed separately, and it is assumed that the taxonomical description of any given secret will be an n-tuple, with each dimension being described by one of a limited number of choices within that dimension.

Motivations for Keeping Secrets

As an initial step into the world of secrets, we will attempt to categorize secrets based on the motivation for creating and keeping them. At first glance, one may instinctively assume that all secrets are alike, especially if one has previously considered secrets primarily in the context of classic information security techniques. Most of information security is at its core based on a hierarchical model of classification that is an outgrowth of military environments and policies. Because of the nature of that area of study, there has been a lot of thought applied to the "how" of keeping secrets with relatively little analysis of the "who," "what," or "why" of the secrets themselves. In considering these questions, it seems that while there are several basic and largely distinct motivations for creating secrets, they can be grouped into two major categories. In the one group there are those cases where the regulation of the secret-keeping behavior is largely from the beliefs and motivations of the secret-keeper, and in the other are those cases where the regulating action is more external, in the form of an entity distinct from the secret-keeper. An evolutionary psychologist might assert that both of these reduce to the same thing, namely, "competitive advantage." In practice, though, humans seem to treat the two subclasses differently.

While the various subcategories are not purely divisible between these two general forms, we will postpone consideration of the nature of and reasons for the overlap, leaving those topics for those more philosophically inclined to debate.

The goal here is a rough generalization that can be refined further, as the community sees fit.

Self-Regulated Secrets

The first category in this dimension of secret space is a self-regulated secret. This is a secret that is kept based on an internal perception of risk. This perception of risk comes in two forms, rational and irrational. In other words, the perception of risk that the keeper of the secret acts upon may be plausible given the available information, or it may be an irrational fear, entirely out of proportion to the actual or plausible risk. In either case, the motivating factors tend to cluster around two centers – embarrassment and control.

Embarrassment

Embarrassment is a slight variation on privacy. In the case of embarrassment, the keeper believes (accurately or not) that the revelation of the information will lead to ridicule, derision, exclusion, or other repercussions based on social factors. Embarrassment does not entail the ability to do harm directly. Instead it deals with information that, when revealed, will change social dynamics and cause others to reassess their opinion of the subject of the secret, who may not be the keeper of the secret. In other words, the keeper of a secret motivated by embarrassment is motivated by the fear of "what people will think" (or "do") if the secret is revealed.

Fear and shame are two of the greatest motivators in human behavior, and both provide impetus for secrets, even though they are likely to be inaccurately calibrated. This type of secret is the driving force behind blackmail, scandal, PR disasters, and some of the less savory aspects of democracy. Oddly enough, experience shows us that embarrassment often leads the keeper of the secret to cause himself additional harm in attempting to avoid the revelation of a secret. In many scandals through history, the initial indiscretion might have been forgivable, had it been admitted to and dealt with. It is often the attempted cover-up that is deemed to have been unforgivable. And how many missteps have been turned into brilliant coups by those perceptive enough to highlight the actions taken to correct the mistake, and thus rise above it? One could write a book about the dynamics of confession and forgiveness, but suffice it to say that the dynamics of embarrassment do not yield easily to formulaic solutions, thus considerably complicating the handling of such secrets.

Control

Control is a slightly different motivation for secrets. In this case, the information being kept secret is believed by the keeper to relate directly to the control of assets, processes, or knowledge that might give others the ability to more directly do harm or gain advantage. This type of secret might include such things as military plans, financial data, bargaining positions in negotiations, trade secrets, safe combinations, etc. Even in situations where there are social dynamics involved, a control secret is one which gives the holder the ability to directly act – to build the bomb, to buy low and sell high, to "head 'em off at the pass," to call the other guy's bluff, or whatever.

Relationship of Embarrassment, Control, and Privacy

If one can envision a spectrum of harm, embarrassment secrets deal with the subjective end of the spectrum. Embarrassment brings in social and emotional factors, and is imprecisely measured, and often based on a seriously inaccurate assessment of likely outcomes. Control secrets tend to be on the objective end, dealing with more tangible, quantifiable factors. The estimation of harm that drives the valuation of a control secret may be inaccurate, but it is usually based in tangible factors.

Privacy secrets may be found anywhere along the spectrum, in that they are not based on the valuation of the secret, or the estimate of harm. Privacy, for the purposes of this discussion, is the keeping of secrets out of a belief that others simply have no need or right to know. While a concern for the harm (in either the embarrassment or control sense) caused by disclosure may coexist with the desire for privacy, we will use the term privacy to cover the principled or philosophical objection that may be overlaid onto the more pragmatic calculation of potential damage.

Externally Regulated Secrets

Externally regulated secrets are those where the keeping of the secret is dictated or regulated by somebody other than the keeper of the secret. If the keeper is an individual, the regulator may be a group to which the keeper belongs, or a third party not directly involved in the dynamics between the keeper of the secret and those from whom the secret is being kept. In the case of organizations, the regulator may be a larger organization, or the governing body of the organization. In either case, there is some sort of "higher authority" relationship in place, relating the individual to the group, or the members of the group to established group norms. Where the keeper's own assessment of the situation drove the behaviors involving self-regulated secrets, externally regulated secrets bring rules and the judgment of outside arbiters into play.

Legal Requirements

The most obvious instance of external regulation of secrets is law or policy. In such cases, there are extant rules that the keeper of the secret is expected to comply with. Examples of areas where such regulations regarding the keeping or revealing of secrets might be found include SEC disclosures, torts, clearances, corporate security policies, etc. The keeper is bound by the rules, and there is clear understanding that violations of the rules will bring punishment by a specified adjudication process, in addition to whatever harm results directly from the revelation of the secret.

Note that in these cases, it is entirely possible for the keeper of the secrets to find that the regulations covering the handling of a particular secret may be derived from multiple regulating regimes that are not mutually, or even internally, consistent. This has been a recurring problem for those designing information security tools. The laws, regulations, and policies as expressed are not always feasibly implemented using the available technology. If the policies are not correctly and unambiguously formulated, it may be impossible to implement them, regardless of technology, simply due to the inconsistencies.

Social Cohesion

Social cohesion is an interesting motivator for secrets, and one that goes to the core of the human dynamics of secrets. Social cohesion secrets are those whose existence is not necessarily predicated on the inherent value of the secret, or the potential harm if it is revealed. Instead, these secrets are kept more for their usefulness in delineating "us" and "them." Humans, whether individually or in groups, tend to define themselves in part by demarcating the differences between themselves and the "others." We are good, hardworking, honest members of the company, the group, the tribe, while "they" are outsiders who follow other norms.

A trivial example of a socially cohesive secret might be noted when ordering a hotdog in the city of Chicago. One might walk up to the hotdog vendor and ask for one hotdog with mustard, relish, and ketchup. Perhaps the order would be filled as specified, but the vendor knows a secret – you are from out of town. The vendor knows this because it is Chicago tradition that a true hotdog consists of a Vienna brand beef wiener (steamed, not grilled), on a Rosen's poppy-seed bun, with cucumber slices, tomato slices, diced onion, sport

peppers, pickle relish (a shade of near-fluorescent green seldom found in nature), yellow mustard, and celery salt. Any of the "sacred seven" condiments may be changed slightly or omitted, but one never asks for ketchup. It is just not done, except by somebody from out of town. If a Chicagoan were to do such a thing, he would feel compelled to make some comment about the fact that he was doing it. By commenting, he would indicate that even though he was violating the social norm, he was aware of it, and thus still a member of the group, even if a somewhat odd member.

This is precisely the sort of "secret" which binds people in shared experience, and separation from those who do not know the liturgy, the cultural references of a particular place and time, the jargon, the recipes, or whatever the secret may be. As one might guess, it is the holding of the secret, and the trust bonds established by doing so, that has value to the keepers of the secret. The actual objective content of the secret may be of little value, either to those holding the secret, or to any outsider. If such a secret is lost, it may easily be replaced – the value is not in the secret itself.

It may be noted that while there are many trivial examples of such secrets, one should not underestimate the fervor with which some secrets of this type may be defended. This is due to the enormous value placed on the trust of the group holding the secret and the exclusivity of it. No matter how meaninglèss or even silly the secret may appear, the powerful desire to be one of the included few may lead the keeper of such a secret to extraordinary lengths to preserve and protect the secret, simply as a matter of personal honor.

Tradition/Momentum

A last external regulatory force for secrets is tradition or momentum. This is the case where a secret continues to be kept beyond its useful life. Perhaps there has been no revocation of the rules rendering it a secret, despite the fact that the information has already become known through other means. Possibly the driving factor is the notion that "we've always done it this way." Whatever the reason, the basic characteristic is that the secret does not need to be kept any longer, but the processes, habits or regulations that governed the keeping of the secret continue on. Bureaucracies are especially suited to this type of secret, due to the lack of a mechanism for periodic review and revision of the regulatory and cultural structures that maintain the secret, either formally or by convention.

Note that it may not even be intentional for such secrets to endure. An example might be a database system originally designed to keep certain fields secret. After the need for the secrecy has gone away, it may still be more trouble than it is worth to revise the system to eliminate the protections on those fields. The owners of the system understand that they no longer need to give special protection to those fields, but the benefit of removing the protection is outweighed by the cost of doing so. Thus the fields remain protected by the system, even though the need to do so has been overtaken by events. This general case has the perverse twist that, over time, the understanding that the protected fields are no longer secret might be forgotten, and the protection may be carried over into new systems if the requirements are not reexamined and updated appropriately as the legacy system or process is replaced.

Valuations of Secrets

In this section, we will set aside the matter of motivation, and look at the actual or perceived value of the secret as another way of viewing secret-space. As previously mentioned, this classification is an overlay to the set of motivations, rather than somehow a further subdivision of it. By this, we mean that any of the secret types denoted here may be paired with any of the motivating factors noted previously. While it may be the case that some of the motivation categories

may tend toward a particular valuation type, they are by no means exclusively coupled, nor is there any type noted here which is inherently excluded from any of the motivation categories.

"Real" Value

This valuation is perhaps the simplest, yet it is still a somewhat fuzzy concept. The idea is that this type of secret is kept secret due to some relatively accurate mapping between the attempt to keep it secret, and the value judgment made in the process of responding to the motivating factors.

As an example, if the secret is being kept due to legal reasons, it is because the information being kept secret does, in fact, fall within the bounds of the law or policy in question. The law applies to the data, and the data are being handled correctly in accordance with the law. A situation where the law was mistakenly believed to apply to particular information, or where the legally mandated protection was not properly carried out, would not be considered to have "real" valuation in our sense of the term.

If the motivation is embarrassment, then categorizing the secret as having "real" value indicates that the secret is in fact not known by others, and is of a nature that it would cause a change in the social dynamics were it to be revealed.

In general terms, the exposure of a secret with "real" value will have an effect at least somewhat correlated to what was expected by the secret-keeper. There is a value to the secret, though the results of revealing it may not be of exactly the form or magnitude anticipated by the secret-keeper.

Illusory Value

A secret with illusory value is one where there is a markedly inaccurate valuation assumption made by either the secret-keeper or those on the outside attempting to derive or reveal the secret.

A classic example might be the Geraldo Rivera television special where, on live television, Mr. Rivera presided over the opening of "Al Capone's secret vault." The vault in question was a bricked-off area in the basement of a hotel in Chicago that had once been the headquarters of the notorious gangster Al Capone. Mr. Rivera became aware of the existence of the vault, and through a series of mistaken assumptions came to the erroneous conclusions that (A) the vault was "secret"; (B) it contained items put there by Al Capone; (C) it had been subsequently bricked over by Capone, or at his direction; and (D) it had not been opened since. In fact, the "vault" is now generally presumed to have been a coal bin of a type common to commercial buildings in Chicago of similar vintage. If this assumption is accurate, it is likely that the bin was bricked off not by Al Capone's gang, but by the owners of the building when the coal-fired furnace was upgraded. In this case, there was no "secret" to be found – only the illusion of a secret, pursued by the sadly mistaken Mr. Rivera.

Another example might be the substance abuser who, for fear of embarrassment, job repercussions, or other social stigma, keeps his addiction hidden. When it is at last revealed, the addict may discover that the response is not the vilification he expected, and that his "secret" was, in fact, already widely known to others, possibly even before he himself was aware of it. In this case, the "secret" had value that was illusory on two grounds – it was not at all well mapped to the reality of the consequences of revelation, and it was not even a secret at all.

It should also be clear that it is possible for a secret to be illusory in multiple ways. In such cases, both the keeper and those trying to find out the secret may be acting out their respective roles based on inaccurate assessments of what the

secret is, the fact that it is (or isn't) a secret, and the value of it.

Irrelevant Secrets

Irrelevant secrets are, as the name implies, secrets that nobody really cares about for their value. Examples might be the secret handshakes, costumes, and rituals of fraternal organizations, or the "secret sauce" (usually either Thousand Island salad dressing or some combination of two or more common condiments) touted by fast food restaurants. These are the secrets that one keeps not for what they are, but for the air of mystery, the sense of fun, or the fellowship surrounding the act of keeping the secret. There is ample room for overlap between the irrelevant and illusory valuation categories. It is entirely possible for either party (the keeper or the one excluded from knowing the secret) to deem a given secret irrelevant, while the other party continues to behave as though it is a secret with real value, as noted above.

Derivation of Secrecy

The next overlay, after motivation and valuation, is the derivation of secrecy. This categorization deals with how a secret becomes a secret. There are two options that are immediately apparent, and while others may exist, it is unclear at this time what they might be. The two categories are discretionary and mandatory secrets. The terminology is chosen for the rough analogy to discretionary and mandatory access control, which is a familiar concept in computer security, but better or more precise terminology may be substituted at a later date to prevent semantic overloading of these terms. Suggestions are welcomed.

Discretionary Secrets

A discretionary secret is found in a situation where there is a set of data items that are related, but the secret-keeper wishes to hide the relationship from the outsider. Another possible construct is where the total data set may not be known, but some subset may be safely revealed. In either case, as long as a sufficient number of the pieces of related information are kept secret, the adversary is denied the complete picture. A common means of protecting such data sets is to pick one or more data elements and keep them secret, while allowing others to be perused freely. This makes it easier to access and process the data, as only a subset of the total dataset is restricted in its usage and handling.

As an example, a record in a government or corporate database might contain a variety of fields pertaining to an individual – name, home address, phone number, social security number, gross annual income, marital status, race, etc. It might be desirable to allow public use of some of the data for statistical analysis reasons, while still keeping other data confidential. In these cases, it may be the secret-keeper's discretion (thus the name) as to which of the fields to reveal, and which to keep secret. It may not matter functionally, for example, if the name or the salary is kept secret, as long as the result is that a particular name/salary combination is not revealed.

Either-or, pick and choose, but the danger remains. If multiple parties are attempting to manipulate the same data, there is a good chance that they will not make the same choices. The result in such cases is that a clever adversary may structure her queries to infer the relationships over the whole that the individual subsystems are attempting to keep secret, barring adequate coordination of their protection efforts.

Mandatory Secrets

Mandatory secrets allow no discretion. It may be that the mandatory secret is, in itself, a unitary item of such value that it must be kept secret, even if there is no contextual linkage to any other information. Alternatively, mandatory secrets

may be the result of a comprehensive and cohesive approach to prior discretionary secrets.

As an example of the latter case, let us imagine a system made up of multiple subsystems, each of which starts with multiple data fields that are treated as discretionary secrets. Let us further assume that the secret-keeper in charge of the first subsystem makes choices about which data fields to keep secrets. Then the next subsystem secret-keeper makes further choices, and so on. At some point, there may be a situation where the choices made by prior subsystem secret-keepers will have limited the possibilities such that subsequent secret-keepers will be faced with mandatory secrets, rather than discretionary ones. The possible permutations will have been constrained in such a way that some fields must not be revealed, lest an adversary be able to infer or deduce the linkages between fields that are desired to be kept secret.

Whether a secret is mandatory due to its intrinsic nature, or due to its relationships to other data items and their respective openness or secrecy, the effect is the same. The stage has been set such that the revelation of a particular piece of information will logically complete the conditions required to allow or cause the predicted harm to occur. The nature and accuracy of the prediction of harm due to revelation are the subjects of the prior categorizations.

The essence of mandatory secrets is that there is no choice, for whatever reason. You do not get to take your pick of what item to focus on, and you do not get to pick whether or not to keep the particular item secret. The keeping of the specific data item as a secret is mandatory, or the penalty conditions will be satisfied.

The Perceived Nature of Secrets

The last categorization overlay or dimension that we will discuss is the way in which the secret is perceived and understood. While the previous categorizations dealt with why the secrets became secrets, their relative value to various parties, and the ways in which the secrecy might be allocated, this categorization deals with the essential being of the secret itself.

These categorizations may be more open to interpretation than those previous, and are inherently fuzzy. One could start all manner of philosophical debate over this segment of the paper, but it is included because of the importance we humans attach to these aspects of secrets, despite our frustration at the lack of precision. The three categories we have come up with are factual secrets, perceptual secrets, and attribution secrets.

Factual Secrets

Factual secrets are more or less discretely discernable and objective. Examples might be formulas, algorithms, laws of nature, and objective truths. These are the sorts of secrets that are demonstrably the "right" (or wrong) answer to a particular question. Even in areas where the measurement is imprecise, a factual secret is one where the application of the secret is not in doubt.

In the story of the lady and the tiger, a man is given a choice between two doors. One holds a lady, the other a deadly tiger. If the man were told the "secret," that a lion was behind the right hand door, he probably would not quibble about the lack of zoological accuracy about the specific genus and species of carnivorous feline, as long as the location of said feline was conveyed accurately. And if the gentleman were smart, he would find no fault with the lady, whatever her qualities, given the alternative! In either case, the secret is factual and may be verified. There is no opinion involved – it's a lady or a tiger.

Perceptual Secrets

Perceptual secrets are those that are less objective. They are subject to interpretation,

based on context, supposition, or the individual thought process of the secret-keeper or the secret-obtainer. If factual secrets are by their nature tolerant of some minor imprecision, perceptual secrets reside almost entirely in the realm of speculation and interpretation. A perceptual secret is seldom the answer in isolation, but is rather a part of a larger framework that may lead one to discernment.

For example, let us presume that a general in a given country holds a secret – he thinks the dictator is an obnoxious idiot. Let us then presume that a foreign government finds out this secret. The exact implications of the secret are still somewhat nebulous. If the general thinks the dictator is an obnoxious idiot, does that mean that the general is amenable to being disloyal and perhaps becoming the leader of a coup? Or does the general believe that his personal self-interest may be best served by staying loyal to a man who he finds personally loathsome? Knowing the fact of the general's loathing does not tell us – the fact is open to interpretation.

As extreme examples, mathematical secrets are likely to be factual, while religious secrets are almost certain to be perceptual by the intent of these categorizations. It is possible to ask a question, to which a perceptual secret is a factual answer, such as "Question: What does General X think of President Y? Answer: General X thinks President Y is an obnoxious idiot" in the example above. That said, it is the interpretation of the importance and context of that factual component of a perceptual secret that is of interest, more than the factual formulation.

Attribution Secrets

Attribution secrets are the basis of many a plot twist in bad novels and farce comedy movies. These are the secrets where the actual secret is of little importance compared to the knowledge of who knows it. These secrets are the foundation of chains of reasoning along the lines "I know that he knows that she knows that Fred doesn't know, but she knows that I know; therefore, I will hide the jewels in Clyde's suitcase, and throw them all off!"

Much of game theory ties in to the conundrums presented by attribution secrets, because if we can know what the other knows, we can often predict his behavior, regardless of whether the data he is basing his decisions on are correct or not. Attribution secrets may also be perceptual or factual, making this a somewhat messy genre for taxonomy purposes. That said, the problem of attributable secrets is of sufficient interest to warrant a place somewhere, and this is where we've stuck it for now. Let the formalists amend our structure as necessary in the future.

So, Why Should We Care?

From the discussion so far, we can see that we've got a few problems in our traditional approach. In the past the computer security community has attempted to develop mechanisms and models of security based on some assumptions that, unfortunately, do not hold true in all circumstances.

The first common assumption is that the quality of being secret is a largely static attribute that has a unitary value – something is a secret of a particular level, at a given moment in time, or it is not. Changes in the secrecy of a particular piece of information are few, and very infrequent, if they are allowed at all by the system.

The second assumption is that secrecy is an abstract quality which can be defined in the absence of context or relationships.

The third is that all secrets are secrets because of their intrinsic value, rather than due to any other cause, or for any other purpose.

Based on these assumptions, we in the security community have constructed very structured

ways of dealing with secrets based on mathematical rigor and automated precision. By doing so, we have attempted to distill the inherently messy illogic of human beings into something which may be predictably repeated in a deterministic fashion in silicon. All the sociological factors are of necessity approximated away in the reduction.

There are many situations where this is an appropriate strategy. There are contexts (particularly in hierarchically structured organizations) where the functional model of secrets is such that the mathematical structure maps perfectly to the actual behavior of the users of the system in that context. As an example, if one asks whether an "A1" system (using the TCSEC rating scale) was a useful thing, many will say that it was too limiting, and no useful work could be done with it. Others will counter that "A1" machines worked fine, and cite places where such systems were used to great effect. How can both be right?

The difference is not that the mathematical rigor went away, or failed in one usage but not in another. It is that in one instance, the model was appropriate to the context in which the system was used. The human system operated under the same assumptions as the digital system in that context. There was no conflict between how the human understood the problem at hand and viewed the elements of the problem, and how the machine behaved to support the human. In the other context, the model perhaps did not map to the needs of the humans. It is not that the humans didn't have secrets to protect, it is that the context in which the humans dealt with those secrets followed different rules, or changed the rules out of synch with the computing system. The mismatch caused more perceived harm than the perceived value of the rigor imposed by the system.

To further complicate matters, we have not even begun to map all the nonintuitive aspects of the rough taxonomy presented here into a model that can be elegantly automated or mathematically described. One can make the case that even in situations where the secret being kept has no intrinsic value, there is still a need to keep it. This is true, but the nondeterministic nature of some of the reasons that humans keep secrets makes this task particularly difficult, as the existence and importance of secrets may fluctuate in relation to context, in ways that we have not yet begun to formalize.

What is missing is the understanding that there might be a different approach to be taken. There is no absolute rule that all problems dealing with secrets must be mapped to, and solved in, the digital domain. There are some types of secrets in the taxonomy above that may defy any attempt at such logical expression. By examining a particular problem with due consideration to both the taxonomy and the discussion above, we may find there is another way. Our efforts might, in some cases, be better applied to understanding the motivations and dynamics that are creating the secrets in the first place, with an eye toward coming up with a system that better maps to the human processes and behavioral tendencies.

We may find that by shifting the structures, rules, and value systems in human space, we have fewer secrets to deal with. We may find that we can alter the human-space systems to clarify the context and valuation of secrets, and put them into a form more amenable to elegant automation. We must get back to the notion that, in dealing with secrets, the human/automation construct works better if the whole system is adapted to the behaviors, values, and motivations of the humans. We must stop expecting the humans to adapt to a model of secrets and behavior formulated for the ease of the digital designer.

It is not our position, however, that automation and mathematical rigor are of no value. We merely assert that such rigor must be in the service of the humans, and not an end in and of itself. It must support, not oppose, the needs and desires of the users. If the human system is incon-

sistent, security professionals may identify the problems and offer assistance in removing the ambiguities. The rigor and formalism of the field may be used as a tool to help the humans come up with systems that are both acceptable to the users and internally consistent. Ultimately, though, the revised system must be acceptable to the humans, or they will actively subvert the system.

Conclusions

We have laid out a very rough, multidimensional taxonomy of secrets, focusing on their nature and origins. We have proposed that this taxonomy will give security practitioners a framework that can be used in understanding human behavior in relation to secrets, which may vary noticeably from the common assumptions. It is the author's hope that this paper might lead those in the security community to further explore all the aspects of human behavior surrounding the motivations, creation, valuation, and handling of secrets. By doing so, we may gain more insight into how to make secure systems that actually meet the needs of the illogical, nondeterministic, oddly programmed system components known as "users."

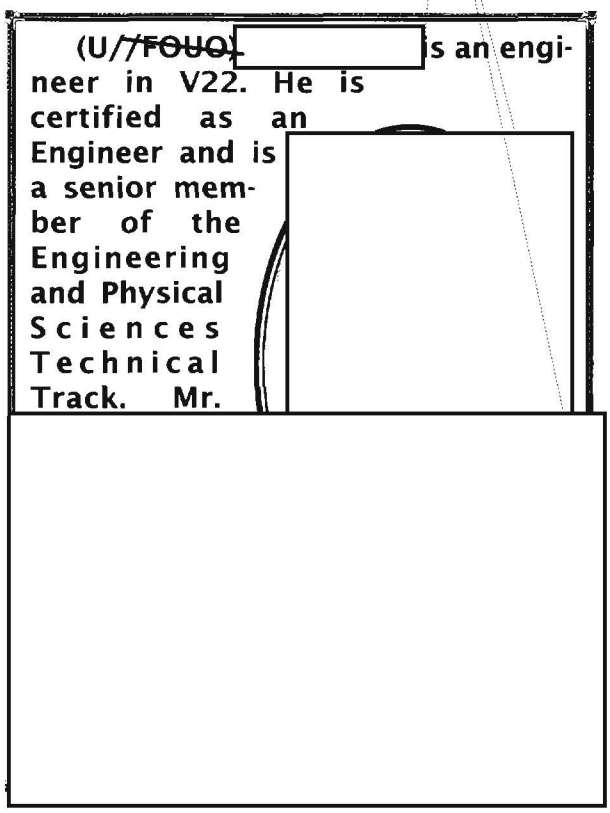

(b)(6)