

# Determining the Utility of Cyber Vulnerability Implantation

## The Heartbleed Bug as a Cyber Operation

Johan Sigholm

Department of Military Studies  
Swedish National Defence College  
Stockholm, Sweden  
johan.sigholm@fhs.se

Emil Larsson

Svenska Dagbladet  
Stockholm, Sweden  
emil.larsson@svd.se

**Abstract**—Flaws in computer software or hardware that are as yet unknown to the public, known as zero-day vulnerabilities, are an increasingly sought-after resource by actors conducting cyber operations. While the objective pursued is commonly defensive, as in protecting own systems and networks, cyber operations may also involve exploiting identified vulnerabilities for intelligence collection or to produce military effects. The weaponizing and stockpiling of such vulnerabilities by various actors, or even the intentional implantation into cyberspace infrastructure, is a trend that currently resembles an arms race. An open question is how to measure the utility that access to these exploitable vulnerabilities provides for military purposes, and how to contrast and compare this to the possible adverse societal consequences that withholding disclosure of them may result in, such as loss of privacy or impeded freedom of the press.

This paper presents a case study focusing on the Heartbleed bug, used as a tool in an offensive cyber operation. We introduce a model to estimate the adoption rate of an implanted flaw in OpenSSL, derived by fitting collected real-world data. Our calculations show that reaching a global adoption of at least 50 % would take approximately three years from the time of release, given that the vulnerability remains undiscovered, while surpassing 75 % adoption would take an estimated four years. The paper concludes that while exploiting zero-day vulnerabilities may indeed be of significant military utility, such operations take time. They may also incur non-negligible risks of collateral damage and other societal costs.

**Keywords**—cyber operations; computer network operations; vulnerabilities; exploitation; intelligence

### I. INTRODUCTION

The increasing use of the cyber domain by individuals as well as businesses and government organizations has made it a welcome target for broad ranges of actors with malicious intent [1]. While hackers, hacktivists and organized cyber criminals, motivated by financial gain, ideology, or personal goals, are the most active actor categories, data breaches as a result of offensive cyber operations have increased rapidly over the past few years [2]. Due to the sensitive nature of such operations, attaining comprehensive, reliable and publishable data is inherently challenging, which hampers the ability to study cyber operations from a scientific or journalistic perspective. While the dissemination of classified information is

undisputedly a serious matter, the disclosures made by former defense contractor Edward Snowden [3] have contributed to shedding light on several aspects of offensive cyber operations, carried out not only by the United States, but also by government agencies of such countries as Sweden and the UK.

The use of highly advanced cyber operations for military purposes, and to reach strategic goals of national interest, is relatively new. Up until a few years ago, intelligence collection in cyberspace relied mainly on the same methods employed by scammers and cyber criminals. Gaining access to a target computer was achieved by so-called “phishing,” sending out malware-infected emails commonly disguised as spam, which compromised the receiver’s computer when opening an attachment or clicking a link. However, a plethora of new methods have been developed during the last few years, focusing instead on such approaches as injecting manipulated data packets straight into the infrastructure used by the target, accomplished through the assistance of network operators [4]. This has dramatically improved the mission success rates. According to an internal U.S. National Security Agency (NSA) memo [5], up to 80 % of all such missions now succeed, in comparison to 1 % of those previously relying on spam [6].

However, there are still some targets that are hard to access: those which employ advanced security methods, those which are moving, or those which are as yet unknown. In these cases, other tools from the offensive cyber operation toolbox may be required. One such “tailored access” tool involves leveraging so-called zero day vulnerabilities, flaws in computer software or hardware that are unknown to everyone but those that have discovered them [7] or in some cases, intentionally created them [8]. This attack vector is hard to defend against, and may thus be effective against targets where other methods fail to yield desired results. Conducting cyber operations employing these methods, especially if they involve compromising services, protocols or infrastructure critical to society, or that are used by a substantial amount of users globally, nevertheless carries a risk of resulting in negative outcomes in the form of adverse collateral consequences [9]. Simultaneously it is important to emphasize that the possible positive outcomes (expected gain) of an operation need to be assessed, since risk can only provide half of the knowledge needed in the decision-making process [10].

Being able to evaluate the utility of an offensive cyber operation, i.e. how well it supports the achievement of overall goals in relation to its costs, is an important metric in order to understand the effectiveness and impacts of exploiting cyber vulnerabilities. It is also pertinent in making decisions on whether to employ a specific attack vector in order to achieve a certain effect, or whether to opt for a different approach. The research question of this paper can be formulated as: *How can the utility of vulnerability implantation be determined?* We focus on a specific aspect of cyber operations involving the exploitation of undisclosed software flaws, namely the time it takes for an implanted vulnerability to reach global adoption. We study this question with a scenario based on the OpenSSL Heartbleed bug that was reported in April 2014. In our reasoning, we will assume that it was intentionally implanted, as part of an offensive cyber effects operations.

The rest of the paper is structured as follows; Section II gives a background to the development of cyberspace during the last few years, including use patterns and costs for the necessary infrastructure. Section III describes the central case study, presents the collected data, and derives a model for predicting global adoption of an implanted vulnerability by data fitting. Section IV discusses the results, and Section V offers some conclusions and suggests possible future work in the area.

## II. CONTEMPORARY CYBERSPACE

Traditionally, the concept of military defense in the cyber domain has included such actions as reviewing code, keeping clients and servers up-to-date with the latest patches, controlling and monitoring data traffic between different network segments, and generally adhering to best-practice information security guidelines, similar to those employed by the industry. However, protecting own systems, networks and resources from emerging antagonistic cyber threats, especially those emanating from actors with large resources at their disposal, requires new offensive methods [7]. Offensive cyber operations may also be used in order to achieve operational effects, or as a support to conventional military operations. These operations are conducted so as to achieve maximal *military utility*. Technical artifacts can be said to have military utility if they, or the effects of their use, allow the goals of a military operation to be reached at a lower cost [11]. Costs can be measured in economic terms, but also in such quantities as saved lives or political risks. In a prescient talk, the NSA goals for a hypothetical, contemporary program to facilitate intelligence gathering would be to “eliminate, prevent or weaken encryption, try to enable access to information, frustrate players who make it harder for us [12].”

If resources are limited, a solution with a broader area of use will be seen as having a greater military utility than a specialized solution. However, if greater economical and personnel resources are available, a number of specialized solutions may, in combination, provide a greater military utility. One such example is within the area of fighter jets. For smaller countries with limited resources, a multirole fighter aircraft may result in greater utility than several platforms specialized for air-to-air, air-to-surface, and reconnaissance missions. A country with greater resources may, on the other

TABLE I. NEGATIVE CONSEQUENCES OF OFFENSIVE CYBER OPERATIONS

| Examples of unintended or collateral consequences of cyber vulnerability exploitation |  |
|---|--|
| Consequence   | Example  |
| Integrity and privacy issues  | Discouragement of the exercise of natural and legal rights |
| Public safety issues  | Risk of cyber attacks against critical infrastructure      |
| Loss of government credibility  | Reduced reputation nationally and internationally          |
| Reduced freedom of the press  | Journalistic sources are not protected                     |
| Direct costs to the industry  | Maintenance and cleanup costs                              |
| Indirect costs to the industry  | Loss of trade secrets, lost business or revenue            |

hand, get more military utility by investing in several specialized aircraft, each one of which will yield a greater mission effect than a multirole aircraft could. The benefits in terms of increased efficiency and reduced cost of a broad, undirected approach may be offset not only in terms of effectiveness, but also by increased risk of collateral consequences, such as those described in Table I.

Standard economic models, like the Net Present Value model which compares the values of opportunities over time, are not very well suited to determining the value of preserving democracy and the security of states. Studies that assess the intelligence factor within economic models are scarce [13]. Military budgets are, by necessity, determined on the basis of probable risk mitigation. Some of the values that are at risk from insufficient protection against direct enemy action – government credibility, transparency, freedom of the press – are ironically also potentially lost from improper handling of intelligence resources. In particular, the ability of the press to protect sources is jeopardized by the kind of omnipresent collection capabilities described in this paper.

When considering the lesser democratic function of physical security, a different way of estimating the value can be found by comparing the cost incurred by security in the reconstruction of Iraq to that of a physically secure, working democracy. In this treacherous environment, security represented as much as 25 % of the total project cost [14]. In other words, the potential cost increase for any project when considering the loss of rule of law is 33 %.

Turning to direct costs to the industry, Fig. 1 [15] shows the worldwide spending on servers and administration. During the last five years, the spending on physical hardware, power and cooling has stabilized on historically low levels. During the same time period, the yearly cost of management has continued to increase, and the number of servers under management has skyrocketed. The cost of management and administration is heavily dependent on a constant need to upgrade and update software in order to mitigate or eliminate the effects of vulnerabilities discovered. A significant portion of the mitigation cost of any security bug is thus based on the rate of adoption.

Large organizations typically do not upgrade to the newest version of any software until it has been thoroughly vetted.

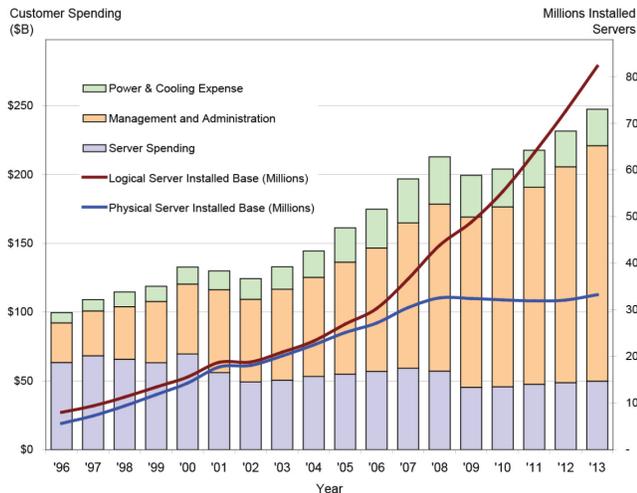


Fig. 1. Worldwide spending on servers, power/cooling and management/administration.

Conversely, there is increased pressure to use a recent, rather than older, version from developers and software vendors who require new features, as well as by reduced or ended support after new versions are released. Thus there exists a window of opportunity during which disclosure of a vulnerability will greatly reduce the cost of mitigation. After this window, adoption will swiftly grow to the majority of the installed base. In order to better understand the shape and duration of this effect, we turn to a case study.

### III. CASE STUDY: HEARTBLEED

#### A. Background

In order to determine the utility of exploit implantation we will make use of a scenario. The vulnerability of the scenario will be the so-called Heartbleed bug in OpenSSL, discovered in March 2014 [16]. OpenSSL, the software containing the vulnerability, is an open-source implementation of the cryptographic protocols Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), for UNIX and other POSIX-compliant operating systems. In development since 1998, it is used by an estimated [17] two-thirds of all web servers globally. Despite the massive installation base and complex, critical nature of the software, OpenSSL is mostly developed by volunteers. The verifiability and quality of the code is low [18], and “reading the source code [...] is close to impossible” [12]. Despite a notable vulnerability in the random numbers generated on the Debian Linux variant between 2006 and 2008, no large scale code review has since been done.

The Heartbleed vulnerability was reported to have been known and exploited by the NSA “for years” [19], a claim that has since been officially denied [20]. Let us now assume that the Heartbleed bug was indeed not only known, but also designed and clandestinely implanted in the OpenSSL software base by an unspecified intelligence organization. As the specific details of the implantation phase are not relevant for our study, they are omitted from the scenario. However, in order to determine if this operation will contribute to reaching our goals (at a lower cost), what we would like to determine is

how quickly nodes implement the compromised code (adoption rate), and thus how long it would take to achieve an acceptable probability of compromise of a given target. This value can then be used as a variable in a traditional Probability of Kill ( $P_k$ ) calculation, a metric that is commonly used to compare the effectiveness of weapons systems. In combination with the cost of the operation, this will give us a way to estimate the military utility of the operation.

It is important to keep in mind that the Heartbleed bug is in itself not a cryptographic flaw or an effect of a weakness in SSL/TLS protocol specification. It is rather the result of a fairly trivial programming mistake (or in our scenario: an intentionally inserted flaw) in a subsystem of OpenSSL providing keep-alive functionality, a so-called Heartbeat. Using Heartbeats is an efficient way for a client or a server to find out if their peer is still active without having to go through a costly session renegotiation, requiring generation and exchange of new keys. An attacker could exploit the flawed heartbeat subsystem by sending a short or empty heartbeat request to the target machine, while incorrectly declaring that it also provided up to 64 kilobytes of data as a payload. As the payload variable is not validated by the target, although mandated by the protocol specification [21], it will reply with a heartbeat including whatever happens to reside in the 64 kilobytes of memory adjacent to the request buffer. By sending many consecutive requests, an attacker could, slowly and over time, collect large portions of sensitive memory content from the target system, without leaving any trace of the attack, as these requests are not logged. This could allow the attacker to come by such information as usernames, passwords, and private cryptographic keys, potentially highly useful for subsequent attacks.

Our scenario starts on New Year’s Eve 2011 when an OpenSSL core developer committed a patch into the application’s source code repository, a piece of software submitted by a German Ph.D. student contributing on a voluntary basis. The patch containing the vulnerable code does not start spreading until the next major release of OpenSSL (version 1.0.1) is released in March of 2012, with the vulnerable code enabled by default. Given that the vulnerability was introduced by just two people, and passed all code reviews for more than two years, the cost of a covert operation aiming to introduce a vulnerability in this way can be assumed to be relatively low. In order to mitigate the collateral cost of such an introduction, we would also introduce honeypots to detect independent exploitation of the vulnerability. In the event of enemy probing, we would utilize back channels to inform the greater community and minimize the potential impact. When the Heartbleed vulnerability was discovered in April of 2014, it was found independently by several sources (Finnish private security firm Codenomicon and Google Security) within a short time of each other, looking much the same as if one of them was given the information to ensure disclosure [22].

#### B. Empirical data

To further understand the value of introducing a software vulnerability like the Heartbleed bug in OpenSSL, we utilized data on the adoption and uptake rate of TLS v1.1 and v1.2. These TLS versions were introduced simultaneously with the

Heartbeat mechanism in OpenSSL v. 1.0.1, released in March 2012. We consider this time to be  $t = 0$  for the purposes of this analysis. Since April 2012, the SSL Pulse project [23], a non-commercial research effort conducted by the SSL Labs of California-based security company Qualys, performs a monthly scan of notable SSL-capable web sites. The sites selected for scanning, some 180-200,000, are drawn from the Alexa list of the most popular websites worldwide. In Table II, we present the data collected each month for the adoption rate of TLS v1.2.

We present the number for TLS v1.2, which is consistently on the order of 10 % higher than that for v1.1. It seems likely that systems which support only the most recent version have adopted software which supports both, but have been configured to allow only the more secure v1.2. The authors are not aware of any encryption stacks which support only v1.1.

As a comparison, TLS v1.0 was specified in January of 1999, but adoption in April of 2012, after some 150 months, was only 99.4 % [23]. Certain systems will be unmanaged or abandoned, others will be prohibitively expensive or complicated to upgrade. From this measurement, and based on the discussion in Section II regarding management costs, it seems reasonable to believe that adoption speed will decrease as cumulative adoption approaches 100 %. In other words, we believe that the true adoption rate can be expressed by a sigmoidal function. Unfortunately, we are not yet at a point in the available data set where the adoption rate drops off, and thus cannot make accurate predictions as to the curve shape when it approaches its horizontal asymptote. It is also uncertain (and perhaps unlikely) whether the true adoption curve will be symmetric in the manner of standard sigmoid functions.

### C. Modeling

In order to extend the collected two years' worth of adoption data into the future, we selected three approaches [24]:

1. Linear extension of the average rate of change.

TABLE II. TLS v1.2 ADOPTION

| Observed global adoption rates of TLS v1.2 <sup>a</sup> |               |         |        |
|---|---------------|---------|--------|
| Survey month  | Adoption rate | (cont.) |        |
| 2012-04   | 1.0 %         | 2013-05 | 14.1 % |
| 2012-05   | 1.4 %         | 2013-06 | 15.1 % |
| 2012-06   | 3.8 %         | 2013-07 | 15.8 % |
| 2012-07   | 4.1 %         | 2013-08 | 17.0 % |
| 2012-08   | 4.5 %         | 2013-09 | 17.8 % |
| 2012-09   | 5.0 %         | 2013-10 | 19.7 % |
| 2012-10   | 5.8 %         | 2013-11 | 20.7 % |
| 2012-11   | 7.6 %         | 2013-12 | 22.5 % |
| 2012-12   | 8.4 %         | 2014-01 | 25.7 % |
| 2013-01   | 11.4 %        | 2014-02 | 28.2 % |
| 2013-02   | 11.9 %        | 2014-03 | 30.2 % |
| 2013-03   | 12.7 %        | 2014-04 | 32.3 % |
| 2013-04   | 13.4 %        | 2014-05 | 35.8 % |

<sup>a</sup>. Data provided by SSL Pulse [23].

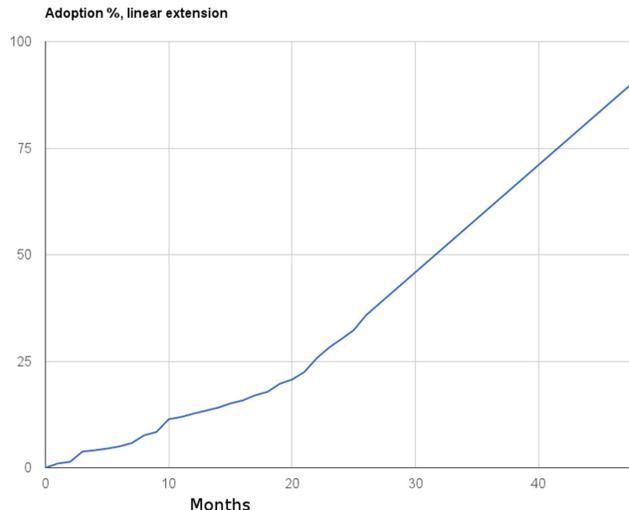


Fig. 2. Linear extension of the average rate of change.

2. Automated curve-fitting to explore 572 different candidate functions to the available data, and selecting the best-fitting curve (least root-mean-square error) which is continuous up to 100 % adoption.
3. Assumption that adoption can be described by a sigmoid function, where TLS v1.0 data  $\{t = 150; y = 99.4\}$  is appended near the horizontal asymptote. As in approach 2, we utilize automated curve fitting to find the best-fitting sigmoid by exploring 46 different candidate functions.

Approach 1, linear extension, had an average rate of change across the entire 27-month period of 1.25 % per month. The rate of adoption increased significantly after 22 months, corresponding to software maturity and mainstream adoption, with a monthly average rate of change of 2.52 %. Extending linearly to  $t = 36$  and  $t = 48$  gives us 61 % and 91 % adoption, respectively. We visualize these results in Fig. 2.

Approach 2, best-fit, returned the common third-order polynomial

$$y = a + bt + ct^2 + dt^3 \quad (1)$$

with coefficients  $a = -0.68697318$ ,  $b = 1.44066888$ ,  $c = -0.06210496$ ,  $d = 0.00230425$ , and a root-mean-square error of 0.75222. We visualize this function and its fit in Fig. 3.

Extending this polynomial we find a high adoption rate of 78 % for  $t = 36$ . The model breaks down at  $t = 39$  when 100 % adoption is reached. We can conclude that the model is likely to be useful only in the short term, i.e. while mainstream adoption is ongoing.

Approach 3, best fitting sigmoidal function with long-term adoption based on TLS v1.0, gave us the logistic sigmoid function

$$y = \frac{a}{1 + e^{-\frac{t-b}{c}}} + d \quad (2)$$

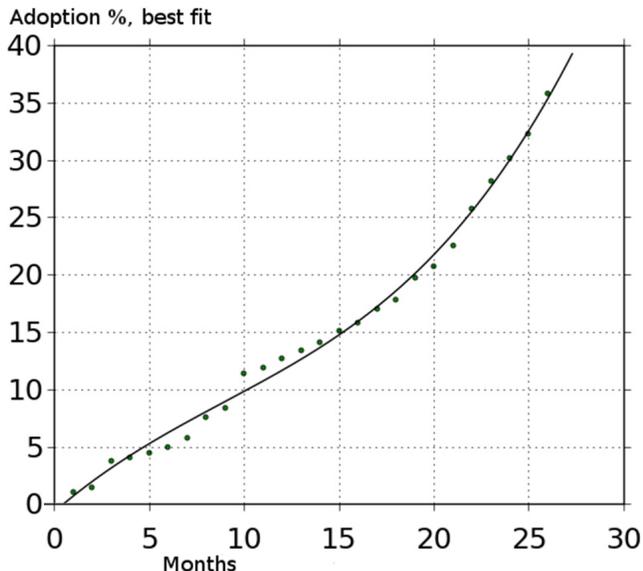


Fig. 3. Best fitting third-order polynomial function.

with coefficients  $a = 105.571648647$ ,  $b = 32.26740092$ ,  $c = 12.46453332$ ,  $d = -6.1106070649$ , but with a significantly higher root-mean-square error of 1.10325.

This model gives us adoption rates of 55 % and 76 % at  $t = 36$  and  $t = 48$ , respectively, as visualized in Fig. 4.

We summarize the results in Table III. We can conclude that with all models, more than half of the servers will have adopted the implanted vulnerability after three years, and more than three quarters after four years.

#### D. Evaluation

After constructing the models using data available until May 2014, we compared the most realistic predictive model, best-fit sigmoidal (Fig. 4), to later collected data from June through August 2014 (Table IV) in order to help validate it. Although it seems clear that the model is not precise enough to be useful in predicting month-by-month change, the trend is not towards model and reality diverging. We believe that a plausible explanation for the modest slowdown in adoption compared to the model could be summer change freezes during the sampled period. An alternative explanation could be a “return to the mean” after increased activity following the discovery and widespread publicity of the Heartbleed bug.

### IV. DISCUSSION

As seen in the scenario above, even though OpenSSL is a very popular open-source implementation of globally employed cryptographic protocols, reaching widespread

TABLE IV. ADOPTION OVER TIME

| Adoption over time for each function |        |             |         |
|--------------------------------------|--------|-------------|---------|
| Function                             | Linear | Cubic       | Sigmoid |
| t = 36 months                        | 61 %   | 78 %        | 55 %    |
| t = 48 months                        | 91 %   | N/A (100 %) | 76 %    |

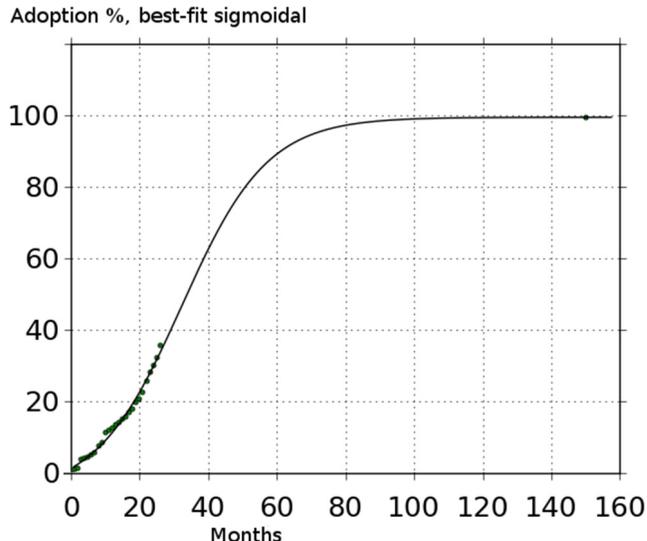


Fig. 4. Best fitting logistic sigmoid function.

adoption of a new version is not achieved overnight. Implanting a vulnerability in software like OpenSSL, even though it might be both feasible and relatively inexpensive, requires a mission endurance of several years, and is thus unlikely to be useful in operations that are time-sensitive.

Nevertheless, having access to a vulnerability in a system actively in use by a target of interest provides a unique opportunity to collect valuable intelligence, to thwart a future attack that is being prepared, or to disrupt or disable cyber infrastructure – denying adversarial access to weapons platforms or systems for command and control. The capability to conduct successful offensive cyber effects operations is thus quite naturally something that a state would pursue. Whereas access to weakly protected targets may be attained through a range of less sophisticated attack vectors, such as sending out malware-infected spam, engaging adequately protected targets may require tailored development of novel tools and methods. While such operations, like the one described in this paper, would require considerable patience, it is conceivable that implanting a flaw in popular software like OpenSSL could be seen as a sort of “investment”, to be used against hardened or air-gapped targets, those that are moving, or those that are yet unknown and will need to be engaged rapidly as they emerge. This of course means that a significant amount of users are left unprotected and potentially exposed while the implanter is waiting for a specific target to adopt the flaw, if or when a third party independently gains knowledge of the vulnerability.

TABLE III. MODEL EVALUATION

| Predicted vs. observed TLS v1.2 adoption rates |                         |                        |        |
|--|-------------------------|------------------------|--------|
| Survey month                                   | Predicted adoption rate | Measured adoption rate | Change |
| 2014-05  | 35.7 %                  | 35.8 %                 | +0.1   |
| 2014-06  | 37.7 %                  | 37.0 %                 | -0.7   |
| 2014-07  | 39.8 %                  | 38.5 %                 | -1.3   |
| 2014-08  | 41.9 %                  | 40.8 %                 | -1.1   |

It is also important to realize that the Heartbleed vulnerability discussed in this paper affects not only web servers, but also clients. A reverse version of the attack could be designed so that a user connecting to a malicious server is made to leak information located in the web browser process memory to the server. As an example, at the time of disclosure more than 50 million handheld devices running version 4.1.1 of the Android operating system were vulnerable to such an attack [25]. This paper does not consider the risks and costs of client-side attacks and mitigation, but a fair assumption is that they are far from insignificant.

The server-side cleanup cost of the Heartbleed bug can be estimated either by top-down or bottom-up approaches. A top-down approach [26] based on figures from the cleanup of the worm W.32 nimda estimated the cost to \$500 million. A bottom-up approach [27] scanned the public internet to find some 28 million publicly available SSL-capable servers. Of these, 615,000 were found to be vulnerable, and 330,000 had already been patched. Tatu Ylönen, creator of the SSH protocol, has estimated [28] the cost of patching and cleanup to hundreds or thousands of dollars per server. In other words, this method gives us a cost of between \$189 million and \$1.9 billion.

As with most other work carried out by military and intelligence services, the goal of cyber operations is fundamentally to uphold national security, and to safeguard citizens from such threats as terrorism, organized crime and military aggression. Doing so requires strict information and operational security, acting covertly and clandestinely, and functioning behind several layers of secrecy. Assuming there will not be another major insider leak like the one carried out by Edward Snowden, it is unlikely that we will find out if the Heartbleed bug was in reality something other than a programming mistake. While secrecy is commonly put in place for good reasons, it may nevertheless be problematic from a democratic perspective if it comes in conflict with other societal values. Securing and maintaining public approval of state activities is not limited to efficiency and providing good value for tax money, but also requires that the state and its agencies can be held to a reasonable degree of public accountability through internal auditing mechanisms as well as investigative journalism. Losing this public consent will likely be costly in terms of social capital.

A central question is therefore if offensive cyber operations are worth their cost. There are obviously many issues to consider when deciding if a cyber operation should be given the green light. Not only should the military utility be considered, e.g. if the operation contributes to reaching military goals to a lower cost and within the allotted timeframe, or if a conventional military approach would be more suitable, but also if the employed methods are proportionate and in accordance with national and international law. While there are established practices for evaluating the utility of a military operation and generating and ranking alternative methods of approach, such as the analytic hierarchy process and war gaming, it is much harder to evaluate the societal utility of such an operation. A problem incurred by cyber operation secrecy is that while a certain method of approach may appear very effective to reach a certain defined military or intelligence goal,

lack of transparency may prevent an analysis from a broader, national or global perspective, which could result in a negative contribution to reaching overarching strategic goals, such as upholding international relations or safeguarding values that are at the core of a democratic society.

## V. CONCLUSIONS AND FUTURE WORK

Offensive cyber effects operations can be seen as a natural extension of traditional military and intelligence operations to the cyber arena, as a means to project state power. They may offer rare and effective means to collect intelligence, create effects or to advance national interests in an unconventional manner, and may be the only option of gaining access to certain targets of high value. However, it is still an immature domain and many countries are still scrambling to develop cyber doctrines and guidelines. As is also the case with conventional military interventions, offensive cyber operations carry risks of collateral damage and unintended adverse effects. The development of specifically tailored cyber vulnerabilities, and the time it takes for them to gain effect, may require significant time and effort, and it is reasonable that review and approval for these kinds of operations is required on the strategic level, even though the mission itself is purely tactical. If the value (expected gain) of reaching the goals of the operation is sufficiently high, it may be a viable option to carry out such an operation, even though the risks are high.

The utility of implanting an exploitable cyber vulnerability can likely be significantly increased through the introduction of some self-destruct or target-identification functionality, which in theory could lessen or mitigate adverse collateral effects. The potential costs of leakage of friendly information, along with conceivable civilian and civic impact, are frankly quite staggering. A sound approach to stockpiling or implantation of cyber vulnerabilities, certainly one as significant as Heartbleed, would thus be to fully take into consideration the collateral cost of the operation, and to apply a risk management policy and procedure approved at the highest executive or political level.

In this paper we have demonstrated an approach towards defining the utility of cyber vulnerability exploitation in a specific scenario. Future work includes devising a generalized proof of implant adoption and exploitation probabilities using additional vulnerability data sets. Further statistical analysis of known Heartbleed exploitation incident data would provide a deeper understanding of the utility of such cyber operations, and would help in rigorously quantifying the probable loss of security to both targeted and collateral entities. We have also pointed out areas where unforeseen costs could arise, which must be taken into consideration when calculating the societal utility of a given operation. Further study includes developing new models and tools to minimize military and societal cost, mitigate collateral damage, and to compare the utility of different cyber operations.

### *Concluding remarks*

During the writing of this paper, the White House put out a press release [29] reinforcing the NSA claim that the U.S. government was not aware of the Heartbleed vulnerability in advance, but noted that “disclosing a vulnerability can mean that we forego an opportunity to collect crucial intelligence.” It

also described nine different guiding questions when determining whether to disclose or stockpile a vulnerability. The first four of these are concerned with collateral damage and civilian cost when the vulnerability is made public or used by the enemy. The last five of these concern the military utility of exploiting the vulnerability. It is notable that the utility is described in terms of intelligence gathering (“an opportunity to collect crucial intelligence,” “How badly do we need the intelligence [...]”) and does not acknowledge other, e.g. effects-based use cases.

It also pertinent to reflect on the sources used in this paper. As mentioned in the introduction, the sensitive nature of cyber operations makes them difficult to study through first-hand sources. Researchers as well as journalists are thus frequently left to an array of secondary, oftentimes online sources, lacking proper peer-review. The risk of impact to research validity and reliability is apparent. In this paper we have taken care not to exaggerate generalizability, and stress that our results should only be considered for the scenario that we present.

#### REFERENCES

- [1] J. Sigholm, “Non-State Actors in Cyberspace Operations,” *Journal of Military Studies*, vol. 4, no. 1, March 2013.
- [2] Verizon, “2014 Data Breach Investigations report,” April 2014 [Online]. Available: <http://www.verizonenterprise.com/DBIR/2014>
- [3] G. Greenwald, *No Place to Hide: Edward Snowden, the NSA, and the U.S. Surveillance State*. Metropolitan Books, 2014. ISBN 978-16-27790-73-4
- [4] A. Guarino, “The State vs the People,” *IEEE Engineering & Technology*, vol. 8, no. 10, pp. 43-45, November 2013.
- [5] Cryptome archive, “NSA Phishing and MTM Attacks,” March 12, 2014 [Online]. Available: <http://cryptome.org/2014/03/nsa-phishing-mtm.pdf>
- [6] Der Spiegel, “Inside TAO: Documents Reveal Top NSA Hacking Unit,” December 29, 2013 [Online]. Available: <http://www.spiegel.de/international/world/the-nsa-uses-powerful-toolbox-in-effort-to-spy-on-global-networks-a-940969.html>
- [7] J. Sigholm and M. Bang, “Towards Offensive Cyber Counterintelligence: Adopting a Target-Centric View on Advanced Persistent Threats,” in *Proc. IEEE European Conference in Intelligence Security Informatics (EISIC 2013)*, Uppsala, Sweden, August 2013.
- [8] S. Egelman, C. Herley, and P. C. van Oorschot, “Markets for Zero-Day Exploits: Ethics and Implications,” in *Proc. 2013 New Security Paradigms Workshop (NSPW’13)*, pp. 41-46, September 2013.
- [9] M. Dunn Cavelti, “Breaking the Cyber-Security Dilemma: Aligning Security Needs and Removing Vulnerabilities,” *Journal of Science and Engineering Ethics*, April 2014.
- [10] H. Liwång, M. Ericson, and M. Bang, “An Examination of the Implementation of Risk Based Approaches in Military Operations,” *Journal of Military Studies*, vol. 5, no. 1, May 2014.
- [11] S. Axberg et al, *Lärobok i Militärteknik*, vol. 9: Teori och Metod (Textbook on Military Technology: vol. 9: Theory and Method). Swedish National Defence College, 2013. ISBN 978-91-86137-23-6.
- [12] P-H. Kamp, “NSA operation ORCHESTRA: Annual Status Report,” keynote speech FOSDEM 2014, February 2014, [Online]. Available: <https://www.youtube.com/watch?v=fwcl17Q0bpk>
- [13] E. Pechta and A. Tishlerb, “The value of military intelligence,” in *Proc. 17<sup>th</sup> Annual International Conference on Economics and Security (ICES2013)*, Stockholm, Sweden, June 2013.
- [14] J. Scahill, *Blackwater: The Rise of the World’s Most Powerful Mercenary Army*. Nation Books 2007. ISBN 1-56025-979-5.
- [15] R. Perry, J. S. Bozman, J. C. Pucciarelli, and J. Scaramella, “The Cost of Retaining Aging IT Infrastructure,” IDC Whitepaper, February 2012.
- [16] B. Grubb, “Heartbleed disclosure timeline,” *The Sydney Morning Herald*, April 15, 2014 [Online]. Available: <http://www.smh.com.au/it-pro/security-it/heartbleed-disclosure-timeline-who-knew-what-and-when-20140415-zqrk.html>
- [17] D. Yadron, “Internet security relies on very few,” *Wall Street Journal*, April 11, 2014 [Online]. Available: <http://online.wsj.com/news/articles/SB20001424052702303873604579495362672447986>
- [18] P-H Kamp, “Please Put OpenSSL Out of Its Misery,” *ACM Queue*, vol. 12, no. 3, pp. 20-23, April 2014.
- [19] M. Riley, “NSA Said to Exploit Heartbleed Bug for Intelligence for Years,” *Bloomberg News*, April 12, 2014 [Online]. Available: <http://www.bloomberg.com/news/2014-04-11/nsa-said-to-have-used-heartbleed-bug-exposing-consumers.html>
- [20] Office of the Director of National Intelligence, “Statement on Bloomberg News story that NSA knew about the ‘Heartbleed bug’ flaw and regularly used it to gather critical intelligence,” *Intelligence Community on the Record, Official Statement*, April 11, 2014 [Online]. Available: <http://icontherecord.tumblr.com/post/82416436703/statement-on-bloomberg-news-story-that-nsa-knew>
- [21] R. Seggelmann, M. Tuexen, and M. Williams, “Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) Heartbeat Extension,” *Request for Comments: 6520, ISSN 2070-1721*, February 2012.
- [22] E. Markowitz, “Behind the Scenes: The Crazy 72 Hours Leading Up to the Heartbleed Discovery,” *Vocativ.com*, April 10, 2014 [Online]. Available: <http://www.vocativ.com/tech/hacking/behind-scenes-crazy-72-hours-leading-heartbleed-discovery/>
- [23] SSL Pulse, a non-commercial research effort provided by Qualys, Inc. [Online]. Available: <https://www.trustworthyinternet.org/ssl-pulse/>
- [24] A. Christopoulos and M. Lew, “Beyond Eyeballing: Fitting Models to Experimental Data,” *Critical Reviews in Biochemistry and Molecular Biology*, vol. 35, no. 5, pp. 359-391, 2000.
- [25] C. Arthur, “Heartbleed makes 50M Android phones vulnerable, data shows,” *The Guardian*, April 15, 2014 [Online]. Available: <http://www.theguardian.com/technology/2014/apr/15/heartbleed-android-phones-vulnerable-data-shows>
- [26] S. Kerner, “Heartbleed SSL Flaw’s True Cost Will Take Time to Tally,” *eWeek*, April 19, 2014 [Online]. Available: <http://www.eweek.com/security/heartbleed-ssl-flaws-true-cost-will-take-time-to-tally.html>
- [27] R. Graham, “600,000 servers vulnerable to Heartbleed,” *Erratasec Blog*, 9 April 2014 [Online]. Available: <http://blog.erratasec.com/2014/04/600000-servers-vulnerable-to-heartbleed.html>
- [28] K. Jackson, “More Than A Half-Million Servers Exposed To Heartbleed Flaw,” *InformationWeek DARKReading*, April 9, 2014 [Online]. Available: <http://www.darkreading.com/informationweek-home/more-than-a-half-million-servers-exposed-to-heartbleed-flaw/d/d-id/1204318>
- [29] M. Daniel, “Heartbleed: Understanding When We Disclose Cyber Vulnerabilities,” *White House blog*, April 28, 2014 [Online]. Available: <http://www.whitehouse.gov/blog/2014/04/28/heartbleed-understanding-when-we-disclose-cyber-vulnerabilities>