

ANDY OZMENT AND  
STUART E. SCHECHTER

## the security of OpenBSD

### MILK OR WINE?




Andy Ozment is a Ph.D. student in the Computer Security Group at the University of Cambridge. He will graduate in July 2007. This article describes work performed in part when Andy was on the technical staff at MIT Lincoln Laboratory.

*andy.ozment@ieee.org*



Stuart E. Schechter is a researcher in computer security at MIT Lincoln Laboratory. He explores security problems as they relate to system design, economics, and user interfaces. Ironically, Stuart can neither digest milk nor tolerate the taste of wine.

*ses@ll.mit.edu*



**FIGURE 1: THE NUMBER OF VULNERABILITIES INTRODUCED IN EACH VERSION RELEASED DURING THE STUDY**

**PURCHASE A FINE WINE, PLACE IT IN** a cellar, and wait a few years: The aging will have resulted in a delightful beverage, a product far better than the original. Purchase a gallon of milk, place it in a cellar, and wait a few years. You will be sorry. We know how the passing of time affects milk and wine, but how does aging affect the security of software?

Many in the security research community have criticized software developers both for releasing software with so many vulnerabilities and for the lack of any apparent improvement in this software over time. However, critics have lacked quantitative evidence that applying effort over time will result in software with fewer vulnerabilities. In short, we don't know whether software security is destined to age like milk or has the potential to become wine.

We thus investigated whether or not the rate at which vulnerabilities are reported in OpenBSD is decreasing over time. For a more technical description of this work, see [1].

### Vulnerability Data

We compiled a database of 140 vulnerabilities reported in the 7.5 years between 19 May 1998 and 17 November 2005. Vulnerabilities were identified by merging data from the OpenBSD security Web page and four public vulnerability databases: NVD (formerly ICAT), Bugtraq, OSVDB, and ISS X-Force.

Figure 1 shows the number of vulnerabilities that were "introduced" in each of the fifteen versions of OpenBSD that were released during our study. A vulnerability is counted as having been introduced in a version if that version is the first to contain the vulnerability within its source code.

How do we know when a vulnerability was introduced? Vulnerability or patch reports often list the versions affected by that vulnerability; however, vendors and vulnerability hunters rarely bother to test more than two or three versions back. So vulnerability and patch reports are not a sufficiently accurate means of finding the release in which a vulnerability was introduced.

Instead, when vulnerabilities were reported, we used the patch to identify the vulnerable code in the OpenBSD CVS repository. We then tracked that code back through all the previous versions

of OpenBSD until we could identify when it had first been checked into the code base. The first release that included that vulnerable code is the release to which the vulnerability is attributed.

Our analysis covers all portions of the OpenBSD code in the OpenBSD team's primary CVS repository. This includes the X-windowing system, the Apache Web server, and many additional services not traditionally considered to be part of the core operating system. However, it excludes the "ports" collection of third-party software, which is not officially part of OpenBSD. We started our study with version 2.3, which we refer to here as the "foundation version," because it was the first source-code stable release for which all reported vulnerabilities are documented.

During the study, versions of OpenBSD were released approximately every six months. The vulnerabilities that were introduced in each version were usually checked into the CVS repository during the six months prior to that version's release. For example, the vulnerabilities attributed to version 2.4 were introduced in the six months between its release and the release of the prior version. The one exception to this rule is the foundation version (2.3): all vulnerabilities introduced before this version's release are attributed to this version. This includes more than twenty-five years since coding on Berkeley UNIX began. As a result, we see in Figure 1 that 62% of the vulnerabilities reported during the study were introduced in the foundation version.

### Source Code Evolution

The majority of vulnerabilities reported during the study were thus introduced sometime prior to the release of the foundation version. But now, 7.5 years later, does the security of the foundation version have any relevance to current versions of OpenBSD?

To answer this question, we investigated the proportion of the source code in the most recent version of OpenBSD that remains unchanged since each earlier version. Figure 2 shows the results of our analysis. Each column represents a composite version; each row represents a source version that contributes code to the composite. A line of code in a composite version of OpenBSD is said to originate in a source version if the line was last modified in that source version. The column for version 2.3 is composed of a single row: By definition, all code in this foundation version is said to originate in it. For each successive version, a new row is added to the column to represent the lines of code that were altered or introduced in that release.




FIGURE 2: THE COMPOSITION OF THE FULL SOURCE CODE

Identifying how the source code evolved over time was a difficult project. We first preprocessed each version of the source code. Only files with the suffix `.c` or `.h` were retained, and all comments were stripped. We then compared each version with each successive version. We used `diff` to compare files with the same path and filename. The `diff` tool was instructed to ignore changes in whitespace and the location of line breaks. Finally, we counted the number of lines in each version that were unchanged from the immediately prior version. By recursively repeating this process, we obtained the data in Figure 2.

The resulting estimate of code commonality is highly conservative. The `diff` tool marked code lines as changed even for trivial alterations such as variable renaming and some types of reformatting—and the OpenBSD team has been reformatting the code base. In addition, if a file from a previous version was moved or copied to a new location and if even one line of the file in the new location was changed, our analysis will treat the entire file as new. Furthermore, if the name of a file is changed, then all of the code in that file is treated as new. Our comparison results will thus understate the degree to which later releases are composed of that is substantively unchanged from earlier releases.

Despite our conservative methodology, Figure 2 shows that unchanged code from the foundation version still comprises 61% of the code in version 3.7—which was released over seven years later. The security of the source code for the foundation version is thus still pertinent to the security of the source code in current versions of OpenBSD.

However, there is another startling result that is visible in Figure 2. The number of lines of source code contributed by the foundation version to each composite version changes over time. That, in itself, is unsurprising. What is surprising is that the number of lines *increases*. How is it that the foundation version contributes more lines of code to version 3.7 than were in the foundation version itself? We discovered that the lines of code derived from the foundation version increases over time because developers reused source code files in different locations. For example, one copy of a compression library file may be used to generate a shared library while another copy of the same file may be used to compile the kernel. Code recycling via source-file replication causes a net increase in the lines of code that are present in later versions.

Several large alterations and/or introductions of code stand out in Figure 2: versions 2.6, 2.9, and 3.5. The magnitude of the changes in versions 2.6 and 3.5 is primarily due to a large number of files being renamed and slightly altered. Our current methodology thus overstates the number of new lines of code and understates the contribution of code derived from earlier versions. The changes in version 2.9 are caused in part by the renaming of files; however, they were also the result of a major upgrade of the XFree86 package.


---

### Milk or Wine?

---

Let's return to our original question: Does software security improve with age? Unfortunately, we don't have enough vulnerability reports to analyze most versions of OpenBSD. Only the foundation version provides us with enough information: 87 "foundational vulnerabilities" were reported. Our question then becomes: Is the rate of vulnerability reporting for OpenBSD version 2.3 decreasing?


We use three different approaches to answering this question. First, we divide the study into two halves and count the number of vulnerabilities reported in each half. Figure 3 shows the result, with confidence intervals calculated by assuming that vulnerability reporting is a homogeneous Poisson process. The number of vulnerabilities reported significantly declines from the first half (58 vulnerabilities) to the second (28 vulnerabilities).



**FIGURE 3: THE NUMBER OF FOUNDATIONAL VULNERABILITIES REPORTED DURING EACH HALF OF THE STUDY**

The next approach is to utilize a Laplace test, in which the discovery of vulnerabilities is assumed to be a heterogeneous Poisson process. The test assesses whether the interarrival times of vulnerability reports are decreasing. We use as our data the number of days elapsed from the identification of one foundational vulnerability to the next.

When the calculated Laplace factors are less than the lowest horizontal dotted line in Figure 4, the data indicates a decreasing rate of vulnerability reporting, with a two-tailed confidence level of 95%. The test finds evidence for a decrease in the rate of vulnerability reporting by the end of year four; by year six, the evidence for a decrease in the reporting rate is statistically significant. This test therefore also supports the conclusion that the rate at which foundational vulnerabilities are reported is declining.



**FIGURE 4: LAPLACE TEST FOR THE EXISTENCE AND DIRECTION OF A TREND IN THE RATE OF VULNERABILITY REPORTING**

In our third approach, we attempt to fit reliability growth models to our data. Although normally applied to the more random discovery of defects, these models can also be applied to the reporting of vulnerabilities. We analyzed the data with seven time-between-failures reliability growth models. One of the seven models had acceptable one-step-ahead predictive accuracy and goodness of fit for the data set: Musa's logarithmic model. According to this model, the number of vulnerabilities expected to be reported on a given day decreases from 0.051 to 0.024 over the course of the study. Furthermore, it estimates that 67.6% of the vulnerabilities in the OpenBSD 2.3 source code have now been found.

Each of our three approaches thus indicates that the rate of foundational vulnerabilities reported is decreasing.

---

### CAVEATS

The rate at which vulnerabilities are discovered and reported depends on the level of effort being expended by vulnerability hunters. To measure how much more difficult it has become to find vulnerabilities over time, we would need to normalize the rate of discovery by the effort being exerted and the skills of those exerting it. Unfortunately, vulnerability reports do not include estimates of how many individuals were involved in examining the software, the time they spent, or their relative skills. Our analysis thus can only show that, in the vulnerability hunting environment that existed during our study, the rate of vulnerability reporting decreased for the foundation version.

---

### While We're at It

Our data prompted us to consider two other questions:


- What is the median lifetime of a vulnerability?
- Do larger code changes have more vulnerabilities?

To answer the first question, we calculate the median lifetime of reported vulnerabilities for the foundation version. The median lifetime is the time elapsed between the release of that version and the death of half of the vulnerabilities reported in that version. Alas, we don't know how many vulnerabilities remain in the foundation version or when they will be found. As a result, we can provide only a lower bound on the median lifetime. The result is striking: It took at least 2.6 years to find half of all the foundational vulnerabilities that would be found during the 7.5-year study period.

The second question is related to "vulnerability densities." Software engineers have examined the defect density of code: the ratio of the number of defects in a program to the number of lines of code. Some have argued that any well-written code can be expected to have a defect density that falls within a certain range (e.g., 3–6 defects per thousand lines of code). Our second question is thus whether or not there is a linear relationship between the number of lines of code altered and/or introduced in a version of OpenBSD and the number of vulnerabilities introduced in that version.

As we cannot measure the total number of vulnerabilities present, we measure the number discovered within four years of release for each version that is at least four years old. Figure 5 illustrates the relationship between the number of lines of altered and/or introduced code and the number of vulnerabilities reported. Neither a visual examination of the fi-

gure nor Spearman's rho test finds a correlation between the number of lines of code altered and/or introduced in a version and the number of vulnerabilities introduced.



**FIGURE 5: THE NUMBER OF VULNERABILITIES INTRODUCED AND REPORTED WITHIN FOUR YEARS OF RELEASE COMPARED TO THE NUMBER OF LINES OF CODE ALTERED OR INTRODUCED, BY VERSION**

When calculated per thousand lines of code, the density of all *reported* vulnerabilities ranged from 0 to 0.033 and averaged 0.00657. There appears to be no trend with respect to the densities increasing or decreasing in newer code. These vulnerability densities are thus three orders of magnitude less than the normal range of defect densities. However, the two figures are not necessarily contradictory: Defects include both vulnerabilities and bugs that are not vulnerabilities. Furthermore, multiple identical security defects that were discovered at the same time are considered a single vulnerability in our data.

## Conclusion

Over a period of 7.5 years and 15 releases, 62% of the 140 vulnerabilities reported in OpenBSD were foundational, that is, they were present in the code at the beginning of the study. It took more than 2.6 years for the first half of these foundational vulnerabilities to be reported.

We found that 61% of the source code in the final version studied is foundational: It remains unaltered from the initial version released 7.5 years earlier. The rate of reporting of foundational vulnerabilities in OpenBSD is thus likely to continue to greatly influence the overall rate of vulnerability reporting.

We also found statistically significant evidence that the rate of foundational vulnerability reports decreased during the study period. We utilized a reliability growth model to estimate that 67.6% of the vulnerabilities in the foundation version had been found. The model's estimate of the expected number of foundational vulnerabilities reported per day decreased from 0.051 at the start of the study to 0.024. We thus conclude that the foundation version of OpenBSD is like wine: It is growing more secure with age.

---

**DISCLAIMER**

This work is sponsored by the I3P under Air Force Contract FA8721-05-0002. Opinions, interpretations, conclusions, and recommendations are those of the author(s) and are not necessarily endorsed by the United States Government.

This work was produced under the auspices of the Institute for Information Infrastructure Protection (I3P) research program. The I3P is managed by Dartmouth College and supported under Award number 2003-TK-TX-0003 from the U.S. Department of Homeland Security, Science and Technology Directorate. Points of view in this document are those of the authors and do not necessarily represent the official position of the U.S. Department of Homeland Security, the Science and Technology Directorate, the I3P, or Dartmouth College.

---

**REFERENCE**

[1] A. Ozment and S.E. Schechter, "Milk or Wine: Does Software Security Improve with Age?" *Proceedings of the 15th USENIX Security Symposium* (Berkeley, CA: USENIX Association, 2006).

Save the Date!



[www.usenix.org/nsdi07](http://www.usenix.org/nsdi07)

**4th USENIX Symposium on Networked  
Systems Design & Implementation**  
April 11–13, 2007 Cambridge, MA

Join us in Cambridge, MA, April 11–13, 2007, for NSDI '07, which will focus on the design principles of large-scale networks and distributed systems. Join researchers from across the networking and systems community—including computer networking, distributed systems, and operating systems—in fostering cross-disciplinary approaches and addressing shared research challenges.

NSDI '07 will be co-located with the following workshops, all of which will be held on April 10, 2007:

- Second Workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SysML07)  
[www.cs.duke.edu/nicl/sysml07](http://www.cs.duke.edu/nicl/sysml07)
- Third International Workshop on Networking Meets Databases (NetDB '07)  
[www.usenix.org/netdb07](http://www.usenix.org/netdb07)
- First Workshop on Hot Topics in Understanding Botnets (HotBots '07)  
[www.usenix.org/hotbots07](http://www.usenix.org/hotbots07)

**Sponsored by USENIX in cooperation with ACM SIGCOMM and ACM SIGOPS**