

DISI - Via Sommarive 14 - 38123 Povo - Trento (Italy)
<http://www.disi.unitn.it>

SECURITY TREND ANALYSIS WITH CVE TOPIC MODELS

Stephan Neuhaus and Thomas Zimmermann

May 2010

Technical Report # DISI-10-034

Security Trend Analysis with CVE Topic Models

Stephan Neuhaus
Università degli Studi di Trento
Trento, Italy
Stephan.Neuhaus@disi.unitn.it

Thomas Zimmermann
Microsoft Research
Redmond, WA, USA
tzimmer@microsoft.com

Abstract—We study the vulnerability reports in the Common Vulnerability and Exposures (CVE) database by using topic models on their description texts to find prevalent vulnerability types and new trends semi-automatically. In our study of the 39,393 unique CVEs until the end of 2009, we identify the following trends, given here in the form of a weather forecast: **PHP: declining, with occasional SQL injection. Buffer Overflows: flattening out after decline. Format Strings: in steep decline. SQL Injection and XSS: remaining strong, and rising. Cross-Site Request Forgery: a sleeping giant perhaps, stirring. Application Servers: rising steeply.**

Keywords—security; trends; machine learning

I. INTRODUCTION

Security is often viewed as an arms race between crackers, who try to exploit flaws in deployed systems, and security people, who wish to make that impossible. It is therefore desirable to know of emerging trends in security in order to be able to think about countermeasures before these emerging trends become large-scale problems.

One large and publicly available source of vulnerability reports is the Common Vulnerabilities and Exposures database (CVE) hosted at MITRE¹. According to the CVE's FAQ², "CVE is a list of information security vulnerabilities and exposures that aims to provide common names for publicly known problems. The goal of CVE is to make it easier to share data across separate vulnerability capabilities (tools, repositories, and services) with this 'common enumeration.'" This seems largely to be true: other databases often contain CVE identifiers whenever they reference a vulnerability that is also in the CVE.

The only large study on publicly available data has been done on the CVE in 2007 [1]. This study relied on a manual classification of entries, helped by a classification system called Common Weakness Enumeration (CWE)³. The CWE aims to be a complete dictionary for software weaknesses.

One problem with the CVE is that the CWE classification system is far too detailed: there are simply too many CWEs to choose from. Therefore, a small set of 19 CWEs was chosen for the user interface to avoid overwhelming the

person making the CVE entry. Steven Christey from MITRE told us: "[MITRE] worked with the NVD team at NIST to come up with a set of identifiers that was fairly small but still gave broad coverage for most vulnerabilities. (You don't want to have to train people to map 700+ CWE names, nor do you want to write the user interface for building up a menu of that size ;-)".

One consequence of this is that it will lead to coarse-grained classification and consequently to loss of information. Also, keeping the list fixed will cause emerging trends to be buried in related but irrelevant classifiers. The same problem arises when using supervised learning techniques, since the topics are fixed in advance and are not allowed to emerge by themselves. In addition, even though the classification system is already coarse-grained, many CVE entries do not have any CWE classification at all. That holds especially for the earlier entries. Therefore, to analyze trends for CVE data, there is a lot of manual work involved; work that will include the exercise of discretion and good taste.

In this paper, we follow a different approach, which is not based on CWEs. Instead we use Latent Dirichlet Allocation (LDA), an *unsupervised learning technique*, on the description texts of CVE entries in order to come up with our own classification system, called a *topic model*. This allows us to identify prevalent topics, but also *emerging trends*, in an automated fashion.

CVE entries also often contain score vectors from the Common Vulnerability Scoring System (CVSS) [2]. We looked at the evolution of these scores over time, which can reveal trends that are captured through topic analysis.

Our contributions to the analysis of large corpora of vulnerability data in general and to the CVE in particular are:

- We have conducted the *first independent study* on the whole body of the CVE database outside of MITRE.
- Since we used a completely different methodology than the original MITRE study, we offer an independent way to *validate* the results periodically published by MITRE. While most of our findings will agree with MITRE's, the independence in methodology gives a greater confidence in the results.
- We offer a mostly *automated approach* to analysing large corpora of vulnerability texts such as the CVE.

¹<http://cve.mitre.org/>

²<http://cve.mitre.org/about/faqs.html>

³<http://cwe.mitre.org/>

Topic models will automatically find the prevalent topics in the CVE and researchers are not limited to manual and potentially error-prone labeling, for example by CWEs or supervised learning techniques. Overall, the topics identified by our automated analysis are in high agreement with the manual CWE classification in the CVE database.

- Related to the previous point, this work can be the basis of a *recommendation system* that suggests possible topics to people entering new CVEs into the CVE database. This could solve the problem of the relative coarse-grainedness of CVE classifications.
- We show how topic analysis helps to identify emerging trends. We found one trend (vulnerabilities involving application servers) for which we offer evidence of growing importance, but which has been judged unproblematic in the original MITRE study.
- We show how well the topics found on the CVE data agree with the CWE classifiers in use today.
- We analyse the CVSS CIA scores for independence.
- We find at least one trend that has gone unnoticed in the original MITRE study, namely that application servers seem to have become a massive target for attackers.

The remainder of this paper is organised as follows: First, we describe our methodology (Section II). Then we present our results (Section III) and describe potential threats to validity of this study (Section IV). We finish the paper with a discussion of related work (Section V), as well as conclusions and future work (Section VI). We also provide a step-by-step guide on how to replicate our study (Appendix A).

This paper contains a large amount of large graphics and tables, which we present at the end of the paper rather than in-line.

II. METHODOLOGY

A. Data Gathering and Corpus Preparation

An annotated version of the CVE database is offered by NIST, and there known as the National Vulnerability Database (NVD)⁴. MITRE creates the CVE descriptions and adds the relevant references. The NVD team at NIST receives this basic information from MITRE (and through the public CVE site). NIST then adds other information such as structured product names and versions, and also maps the entries to CWE names.

NIST offers files named *nvdCVE-2.0-year.xml* for download, where *year* is a number from 2002 to 2010 inclusive. The file for 2002 contains entries from 1988 through 2002, and the file for 2010 is necessarily incomplete. Since there is usually a period of consolidation when a new CVE entry (called a *candidate*) might be changed or removed, we only looked at entries that were published up to 2009, inclusive.

⁴<http://nvd.nist.gov>

&apos	and	does	into	or	then	whether
"	are	e	is	p	there	which
+	as	ff	it	s	these	while
++	at	for	k	some	they	who
?	attack	from	kei	ss	thi	whose
??	be	g	later	such	to	with
a	been	get	mc	t	up	within
about	by	go	no	than	us	without
all	can	ha	not	that	v	x
allow	dn	has	o	the	via	
also	do	if	of	their	wa	
an	doe	in	on	them	when	

Table I
LIST OF STOP WORDS.

aix 1	execut 1	ibm 1
arbitrari 1	file_comp 1	overflow 1
buffer 1	for 1	rcp 1
code 1	function 1	remot 1

Figure 1. Stemmed word list for CVE 2002-1621

I order to build timelines of topics, we need to know when the problem in a CVE was discovered, disclosed, or published. For this, we use the disclosure date of the corresponding OSVDB⁵ entry. The OSVDB is an open vulnerability database that is comparable to CVE in scope. While a CVE entry contains a date field called `published-datetime`, this field cannot be used to date a CVE reliably, according to information from MITRE.

From each CVE entry, we extracted the CVE ID (a unique identifier) and the summary text. We then subjected the words in the summary text to the Porter stemming algorithm [3], which attempts to find word stems so that for example the words ‘programming’, ‘programmed’, and ‘programs’ would all be mapped to their stem ‘program’. Next, we removed a number of stop words, i.e., words that are so common that they do not help differentiate between CVEs; see Table I.⁶ Finally we counted the occurrence of each stem in the summary text. After these operations, we therefore had, for each CVE, a list of word stems and the number of times they occurred in that CVE. For example, Figure 1 shows the list for CVE 2002-1621, clearly a buffer overflow.

B. Topic Models

We subjected these stemmed CVEs to Latent Dirichlet Allocation (LDA). This is an unsupervised machine learning technique that treats CVEs as bags of words, where each CVE is characterised by a distribution over topics, and each topic in turn is characterised by a distribution over words [4]. We let LDA identify 40 topics, starting from random topic

⁵<http://www.osvdb.org>

⁶Seemingly nonsensical stop words like ‘a’, ‘e’, ‘g’, ‘k’, or ‘mc’ come about by breaking words like “e.g.”, “a.k.a.”, or “MC/ServiceGuard” at word separators like the period ‘.’ or the slash ‘/’.

assignments.⁷ If there are m unique stemmed words in the entire CVE corpus, and if there are n CVEs, the result of LDA contains, among other things, a file that assigns the word with index k ($1 \leq k \leq m$) in document c ($1 \leq c \leq n$) a topic z ($1 \leq z \leq 40$), meaning that word k in document c is about topic z . Note that the same word can be assigned to different topics in different documents. Another result of the LDA software is a list of the most frequent words w in a topic z , together with their probabilities $\hat{p}(w | z)$.

First, we define the probability that a given document c is about a given topic z :

$$\hat{p}(z | c) = \sum_{\substack{w \in c \\ w \text{ is about } z}} \frac{\# \text{occurrences of } w \text{ in } c}{\# \text{words in } c}. \quad (1)$$

This formula can lead to fractional (or fuzzy) assignments. For example, CVE 1999-1471 contains 18 words, eight of which are assigned to topic 21 (manually labeled “Buffer Overflow”), and ten to topic 22 (“Credentials Management”). This CVE is about the passwd program; the description reads “Buffer overflow in passwd in BSD based operating systems 4.3 and earlier allows local users to gain root privileges by specifying a long shell or GECOS field.”, so the assignment (about 44% buffer overflow, 56% credentials management) is quite accurate.

We also give an indication of the *overall importance* of a topic, which is given by

$$\hat{p}(z) = \frac{1}{\# \text{CVEs}} \sum_c \hat{p}(z | c). \quad (2)$$

When LDA is run on a corpus as large as the CVE data with a comparatively large number of topics, it can happen that some of the topics are really the same and should be combined. This will be the case when there is a large overlap in the set of words that appear in different topics. When k topics z_1, \dots, z_k are combined into a topic z , a that topic will have the joint probability

$$\hat{p}(z) = \sum_{1 \leq j \leq k} \hat{p}(z_j), \quad (3)$$

where $\hat{p}(z_j)$ is given by Equation (2).

One result of the LDA software is a list of the most frequent words in a topic, together with their probabilities. Again, when k topics are combined, a word w will have the joint probability

$$\hat{p}(w | z) = \sum_{1 \leq j \leq k} \frac{\hat{p}(w | z_j)}{\hat{p}(z_j)} \hat{p}(z), \quad (4)$$

where $\hat{p}(z)$ is given by Equation (3).

⁷Automatically finding the best number of topics when doing unsupervised learning is at this point an unsolved problem. It has been solved for supervised learning [5], and there are some ideas for unsupervised learning, but at the time of writing, there is no solid theory. The number 40 seemed to be a good value with good topic separation.

C. Trend Analysis

LDA does not measure changes over time. One alternative is to use Dynamic Topic Models [6], which would change the topic distributions and the word distributions within a topic from one year to the next. Topics over Time [7] works the other way around, assuming that a CVE has a time stamp depending on a topic-specific beta distribution.

Both approaches have restrictions, however: dynamic topic models penalise large changes in word distributions for any particular topic from one year to the next, whereas topics-over-time distributions are rather inflexible.

Therefore, we follow Hall et al. [8] and study only the post hoc empirical probabilities that a randomly selected CVE in a given year y is about topic z .

$$\begin{aligned} \hat{p}(z | y) &= \sum_{c:t(c)=y} \hat{p}(z | c) \hat{p}(c | y) \\ &= \sum_{c:t(c)=y} \hat{p}(z | c) \frac{1}{\# \text{CVEs in year } y} \\ &= \frac{1}{\# \text{CVEs in year } y} \sum_{c:t(c)=y} \hat{p}(z | c), \end{aligned} \quad (5)$$

where $\hat{p}(z | c)$ is given by Equation (1) and $t(c)$ is the year in which CVE c was published. This formula is simple to compute, since $\hat{p}(z | c)$ is a result of the LDA model estimation process.

Plotting $\hat{p}(z | y)$ against y will then give an indication about the *relative importance* of topic z in year y : if $\hat{p}(z_1 | y)$ is twice as high a $\hat{p}(z_2 | y)$, then twice as many CVEs were about z_1 than about z_2 in year y .

In addition to a graphical display, it is also interesting to know how much a topic has changed in importance, both since measurements began and in the last year. In general, we compute the change from year y to year $y + n$ as the *average annual change*:

$$\text{change} = \left(\frac{\hat{p}(z | y + n)}{\hat{p}(z | y)} \right)^{1/n} - 1, \quad (6)$$

The intuition behind this is that this number will give the relative change that the initial level of importance will have to undergo every year to arrive at the final level of importance. For example, if this number is 0.1, then that means that the importance of the topic has risen by 10% per year on average.

D. Causes and Impacts

During exploratory analysis of the CVE corpus we noticed a curious feature, namely that most (72.9%) were of the form “ x allows remote attackers to y ” or similar forms. For example, CVE 2008-0895 reads, “BEA WebLogic Portal 10.0 and 9.2 through MP1, when an administrator deletes a single instance of a content portlet, removes entitlement policies for other content portlets, which *allows attackers*

	Actually is x	Is not x
Classified as x	true positive (TP)	false positive (FP)
Classified as not x	false negative (FN)	true negative (TN)

$$precision = \frac{\#TP}{\#TP + \#FP} \quad recall = \frac{\#TP}{\#TP + \#FN}$$

Figure 2. Assessing the quality of a classifier.

to bypass intended access restrictions.” This separates the description of the CVE into two parts, the first of which describes the *cause* of the problem and the second the *impact*. For those CVEs for which we found such a division, we subjected them to the same LDA process in the hope of finding out trends in both causes and impacts.

E. Alignment with CWEs

One problem is that topic models, being an unsupervised learning technique, do not give indications of goodness-of-fit, so we usually have no quantitative indication on how well the classification works. In this case, however, many CVEs do have an independent classification in the form of CWE assignments, so we can compare the CWE value with the topic that was assigned by running LDA.

Usually, the quality of a classifier is determined by using the false-positive and false-negative rates, or equivalently the related measures of *precision* and *recall*. High precision values mean that classifications are often correct, and high recall values mean that many documents will be correctly classified; see Figure 2. We cannot use these values directly, however, since our topic assignment based on LDA is *probabilistic*. Remember that a document can be about multiple topics, which is captured by the probability $\hat{p}(z | c)$. We therefore need to transform our probabilistic topic assignment into a crisp one, where each document is assigned exactly one topic. We do this for a given document c by taking the topic z for which $\hat{p}(z | c)$ is largest.

We compute precision and recall values separately for each topic. This helps us identify topics that are aligned particularly well with CWEs. Let $cwe(c)$ be the CWE value for a document c and let $\hat{z}(v)$ be a mapping that maps a CWE v onto a topic. (For simplicity, we assume that \hat{z} is single-valued.) Then we count for each document c with topic z :

- a *true positive* for z if $z = \hat{z}(cwe(c))$;
- a *false positive* for z **and** a *false negative* for $\hat{z}(cwe(c))$ if $z \neq \hat{z}(cwe(c))$.

The intuition is that if the CWE topic and the LDA topic are the same for a document, i.e., $z = \hat{z}(cwe(c))$, then that should count as a true positive. If they are not, however, it should be counted as a false positive for the LDA topic z and as a false negative for the true CWE topic $\hat{z}(cwe(c))$. We do not count true negatives because they are not needed for the computation of precision and recall (see Figure 2).

F. CVSS Scores

CVE entries have sections containing scores from the the Common Vulnerability Scoring System (CVSS), a framework to assess and prioritise vulnerability information [2]. According to its creators, the main benefits of the CVSS are:

- **Standardized Vulnerability Scores.** The idea is to enable organisations to leverage a single vulnerability management policy across all of its platforms.
- **Open Framework.** The idea is to be able to answer the question “Which properties gave it that score? How does it differ from the one released yesterday?”
- **Prioritized Risk.** Here the idea seems to be that when the vulnerability score is computed, it is somehow indicative of actual risk.

CVE data contain six pieces of CVSS information:

- **CIA Impact.** This is an assessment of the impact of the vulnerability on *confidentiality*, *integrity*, and *availability* (CIA). Each impact assessment comes in three levels, “none” (no impact), “partial” (partially compromised) and “complete” (completely compromised).
- **Authentication.** This is the number of instances the attacker must *authenticate* in order for the attack to work. Values are “none” (no authentication required), “single instance” (must authenticate once), or “multiple instances” (must authenticate twice or more).
- **Access.** This says *how close* the attacker must be to the attacked host. Values are “local” (local access required), “adjacent network” (access to broadcast or collision domain required), or “network” (access to any network suffices).

In addition, the *complexity* of the access is captured. Values are “low”, “medium”, and “high”.

Given the requirements of standardisation and risk prioritisation, the individual CVSS scores should be *independent* and *uncorrelated*, that is, it should not be possible to predict the value of one CVSS score from the value of another.

III. RESULTS

A. Overview

The earliest CVE has an OSVDB disclosure date of August 1, 1982 (CVE 1999-0531, a candidate that is now rejected). By contrast, the CVE *with the earliest published-datetime field in the CVE database* is CVE 1999-0095, reporting the sendmail DEBUG hole that was exploited by the Morris worm [9]). The latest are from December 31, 2009, comprising a total of 39,749 entries. After removing duplicates, 39,393 unique CVEs remained.⁸

After removing stop words and stemming, the CVE summaries could be as short as a single word (for example

⁸We have checked 53 duplicate CVE IDs manually to confirm that they were indeed duplicate entries having identical fields and not different entries with the same CVE ID.

for CVE 1999-0657, for which the entire summary reads “WinGate is being used”, which after removing stop words and stemming reduces simply to “wingat”), but also as long as 99 words (for CVE 2007-0018, which describes fairly completely a stack-based buffer overflow in an ActiveX control). The distribution is shown in Figure 3 (left). It is clearly a unimodal distribution, having a median of 18 words, a mean of 18.85 words, and a standard deviation of 6.49 words. But it is not a symmetric distribution (skewness 1.51, kurtosis 5.91)⁹.

The right part of Figure 3 shows how the CVE database grew over the years. Apparently, the number of CVE entries peaked in 2006 with 6,885 submissions in that year, and has been slightly declining in 2007 (to 6,393) and more drastically in 2009 (by a massive third to 4,446). We assume that over time, more CVE entries that are published in 2010 will turn out to have been disclosed in 2009, so that this number will probably rise.

B. Topic Models on the Entire Corpus

Combining topics, we arrived at 28 unique topics. Topics for which we could not find a good name are simply labeled “Topic x ”, where x is the topic number produced by LDA.

The first main result is a table with all 28 unique topics (Table VII). The first column contains the topic name z . The second column contains $\hat{p}(z)$, the proportion of words that have been assigned that topic by LDA, according to Equation (3); the table is sorted by the $\hat{p}(z)$ column. The remaining columns contain the most frequent words from that topic and their frequency $\hat{p}(w | z)$, according to Equation (4).

The table contains all the expected topics such as buffer overflows, format string vulnerabilities, SQL injection, cross-site scripting, cross-site request forgery and so on, and roughly in the order that we would intuitively expect them to be from following the security press.

However, there are also a few surprises:

- There are topics just for Linux kernel issues, and one for Internet Explorer. The reason is that with partial/fuzzy assignments, LDA will not try to find mutually exclusive topics, since a document can be partly about one topic and partly about another.
- The proportion of resource management issues is very high. This is due to the combination of denial-of-service topics with more traditional resource management topics such as memory leaks. This was done because manual inspection of respective CVEs showed that these topics were not really separable: a CVE that had words assigned to resource management issues usually also had words assigned to a denial-of-service topic.
- CVE Management issues (disputed entries and rejects) make up almost 3% of the CVE corpus. We will return

to this topic when discussing the changes in topic importance below.

- The word “PHP” appears in four of the 28 topics, SQL injection and cross-site scripting, implicating PHP in these vulnerabilities. However, the proportion of PHP-specific vulnerabilities has been declining since 2007, especially for arbitrary code executions due to PHP. One reason might be that support for PHP 4 was finally discontinued in August 2007, forcing web site operators to upgrade to the more secure PHP 5. See also Section III-C below.
- For the topic labeled “Cross-Site Request Forgery”, the strings “cross-sit” and “forgeri” appear only beyond the 10 most frequent words (stemming reduces “cross-site” to “cross-sit” and “forgery” to “forgeri”).

C. Trend Analysis

The second main result is a graphical display of $\hat{p}(z|y)$ for the 28 topics since 2000 (Table IV). The main findings are:

- Well-known vulnerabilities like buffer overflows and format strings are declining, though perhaps not as steeply as one would wish. It seems that buffer overflows are harder to exploit now than they were in the past, and therefore that effective protective measures have finally found their way into operating systems, compilers, and libraries.
- Resource management issues are also generally declining. This puts the high percentage of this topic in Table VII into perspective: the situation was bad, but it is getting better.
- Privilege escalation and link resolution issues are also going down, perhaps indicating better compartmentalization and secure-by-default configurations.
- Perhaps surprisingly, the importance of exploits allowing arbitrary code execution is also going down slightly in recent years (but see the section on causes and impacts below).
- SQL injection and cross-site scripting have dents in their growth curves of the last few years.

The third main result is Table II, which shows more quantitatively the change in importance for the 28 unique topics identified by LDA from 2000 to 2009 and from 2008 to 2009, according to Equation (6).

The table is sorted in decreasing order of combined change¹⁰; the “since 2000” column can be used to see whether the change from 2008 to 2009 was in accordance with or against the general trend since 2000. For example, application server-related vulnerabilities have risen by 20.6% annually since 2000, and by 83.2% from 2008 to 2009.

⁹Skewness and kurtosis have been computed with methods compatible with SAS and SPSS.

¹⁰We used a weighted geometric average, weighing the change since 2000 ten times as high as the change since 2008.

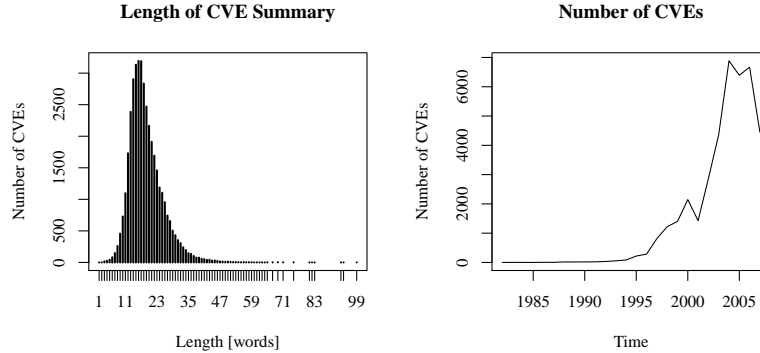


Figure 3. Length of CVE summary text after stemming and stop word removal (left) and number of CVEs issued per year between 1988 and 2009 (right).

Name	since 2000	since 2008	Name	since 2000	since 2008
Cross-Site Scripting	+46.6	+3.8	Firewalls	-0.7	+44.1
SQL Injection	+45.7	-39.6	Resource Management	-1.4	+36.1
Arbitrary Code (PHP)	+37.7	-23.0	Linux	-2.6	+50.0
PHP	+29.9	-8.7	Arbitrary Code	+2.6	-16.1
Application Servers	+20.6	+83.2	Directory Traversal	+1.9	-19.6
Topic 35	+15.6	-8.7	Format String	+1.9	-23.1
Microsoft Office	+7.9	+37.6	Buffer Overflow	-4.3	+37.7
Mozilla	+5.3	+68.8	Message Boards	+0.4	-22.7
Information Leak	+10.6	-5.7	Topic 17	-10.4	+45.2
Microsoft Windows	-2.6	+139.0	Credentials Management	-7.4	-12.8
Topic 7	+1.0	+65.6	Arbitrary Code (IE)	-11.2	+5.1
Java	+7.0	-13.1	Cryptography	-10.6	-2.4
CVE Issues	+5.7	-9.5	Privilege Escalation	-17.8	-8.5
Cross-Site Request Forgery	+5.5	-17.6	Link Resolution	-14.5	-51.3

Table II

TOPICS SORTED ACCORDING TO THEIR CHANGE IN IMPORTANCE SINCE 2008. THE COLUMN MARKED “SINCE 2000” CONTAINS THE AVERAGE CHANGE PER YEAR SINCE 2000 IN PERCENT, AND THE COLUMN MARKED “SINCE 2008” CONTAINS THE CHANGE FROM 2008 TO 2009 IN PERCENT. LONG-TERM CHANGES WEIGH MORE THAN SHORT-TERM CHANGES.

The comparatively unimportant Topic 35 appears relatively high up the list because it rose from a mere 0.33% in 2008 to 1.5% in 2009. Since link resolution has been stagnating or declining on average since 2007, we believe this to be a random fluctuation, not a new trend.

Above we mentioned that incomplete or disputed CVE entries make up about 3% of CVE entries. Looking now at the changes in importance, we can see that this quality has been eroding since 2000, but apparently some quality measures have taken effect. In fact, when one looks at the plot of that topic in Figure IV, one can see that after a sharp, almost exponential increase until 2006, the topic has been falling.

D. Causes and Impacts

The CVE corpus contains 28,699 unique entries of the form “*x*, allows attackers to *y*” or similar forms, 72.9% of all unique entries.

Tables VIII and IX show causes impacts respectively, analogous to Table VII. What can be seen from the table is that the four most frequent causes are responsible for two thirds of all CVEs; after the top four, topics fall off

sharply in importance. PHP is identified as one of the four major causes, making it a prime candidate for improvement. For the impact data, the situation is even more clear-cut: arbitrary script and code executions make up almost 58% of all impacts. This suggests that research on the prevention of such vulnerabilities could have major impact (no pun intended). While it is true that adding denial of service and information leaks would make the figure rise to 92%, preventing such attacks is in our opinion more difficult than preventing arbitrary code execution.

One disappointing element in Tables VIII is the appearance of the word “unspecifie[d]” in the table entries for various topics. Looking at the corresponding CVEs, one finds that the phrase “unspecified vulnerabilities” or “unspecified vectors” is responsible for this. For example, CVE 2007-0114 reads, “Sun Java System Content Delivery Server 5.0 and 5.0 PU1 allows remote attackers to obtain sensitive information regarding ‘content details’ via unspecified vectors.” Entries like these simply mean that the source cannot or will not disclose the actual cause of the vulnerability. That will in turn mean that the CVE entry

will be incomplete, rendering it less useful.

There is an apparent discrepancy in the cause/impact data versus the entire data as seen in Table VII. For example, in Table VIII, we see that 19% of all CVEs are about cross-site scripting, whereas the corresponding column sum in Table VII gives only about 9.7%. The reason for this effect is that sometimes a given cause may allow different impacts. For example, a buffer overflow may at one time allow the execution of arbitrary code, at another it may lead to a crash and therefore to a denial of service. So LDA might assign a document to either the cause (buffer overflow) or to the impact topic (arbitrary code or denial of service), or partly to both, depending on which part dominates in the description. As soon as the document is split into cause and impact, however, assignments can be made more clearly.

Figures V and VI show graphs of the unique cause and impact topics analogous to Figure IV. For both data sets, the important information is that the graphs follow the same general curves as the corresponding ones in Figure IV. For example, cross-site scripting and SQL injection are rising, whereas buffer overflows, denial of service and privilege escalation are falling and PHP seems to have had its peak in 2006 or 2007.

E. Alignment with CWEs

Next, we looked at how well the 28 topics found by LDA aligned with the 19 pre-assigned CWEs that are available to someone who enters a new CVE. To do that, we first needed to map LDA topics to CWEs. We ended up with the assignment shown in Table III, but not all topics were assigned a CWE or vice versa. The reason for this is that an LDA topic might not coincide naturally with an available CWE or only with an CWE that is not used in the NVD. For example, Topic 23 was labeled “Privilege Escalation”, and CWE 269 (improper privilege management) exists for just such a case. But this CWE is not one of the 19 CWEs offered for CVE entries. This is not a problem for LDA, since partial assignments are possible (“this document is 50% about topic 1 and 50% about topic 2”), but the CVE does not allow such partial assignments. That said, we were able to map 12 of the 24 LDA topics directly to CWEs.¹¹

Table III shows the precision and recall values according to Figure 2. These show that the performance of LDA is very good when it comes to standard categories like cross-site scripting, directory traversals, link resolution or SQL injection (precision of 80% or more and recall of 80% or more). Other categories fare much worse, among them the buffer overflow. This is surprising at first glance, since buffer overflow reports ought to have the two words

¹¹Of the 19 CWEs used in the NVD, we could not map seven CWEs to LDA topics; they are: CWE 16 (configuration/insecure defaults), CWE 20 (improper input validation), CWE 78 (OS command injection), CWE 189 (numeric errors), CWE 255 (credentials management), CWE 310 (cryptographic issues), and CWE 362 (race conditions).

“buffer” and “overflow” occurring somewhere. However, this need not be the case; for example, CVE 2008-0090 reads, “A certain ActiveX control in npUpload.dll in DivX Player 6.6.0 allows remote attackers to cause a denial of service (Internet Explorer 7 crash) via a long argument to the SetPassword method.”, so possible CWEs would include 20 (input validation), 399 (resource management), or 255 (credentials management), yet the CVE was assigned CWE 119 (buffer overflow), something that is not apparent from the description.

False positives and false negatives in the assignment of LDA topics to CWEs can exist for multiple reasons.

- 1) In practice a CVE can be about multiple CWEs; while LDA accounts for this with its probabilistic assignment, the actual CWE assignment of in the NVD database does not.
- 2) Ambiguous entries in the CVE database. For example if a buffer overflow allows the injection of arbitrary code, should the CVE entry be classified as CWE 119 (buffer overflow) or CWE 94 (code injection)?
- 3) The assignment of a CWE to a CVE is to some degree arbitrary, too, as the example for CVE 2008-0090 and buffer overflow shows.

F. CVSS

Figure 4 shows the development of CIA impact scores for confidentiality, integrity, and availability (top row) and for authentication and access proximity and complexity (bottom row). What is striking is that for the CIA plots the individual lines track each other extremely well across figures.¹² This suggests that CIA scores are correlated and not independent in practice.

Figure 4 shows several trends. The fraction of complete compromises for confidentiality, integrity, and availability have decreased substantially until 2006; but since 2007 the fractions are increasing again (see top row). The plot “Access” shows that local access has declined dramatically over the past years and that most vulnerabilities can now be exposed through the network; adjacent networks are almost never used for the access score. The plot “Access Complexity” shows that the complexity of attacks is increasing (fewer Low, more Medium). The plot “Authentication” shows that in recent years a small fraction of vulnerabilities required users to authenticate once; however the majority (over 90%) required still no authentication at all.

IV. THREATS TO VALIDITY

We let LDA seed topics randomly. On the one hand, this is precisely the point of using an unsupervised learning technique. On the other hand, random topic seeding could

¹²So well that we spent a few hours tracking a non-existent bug in the software that produced the graph. In fact, the Pearson correlation coefficients are very high (above 0.75) but do not reach 1.00, which means that the lines do not match perfectly.

P[%]	R[%]	LDA Topic Name	CWE	CWE Name
97.8	94.6	SQL Injection	89	Improper Sanitization of Special Elements used in an SQL Command ('SQL Injection')
98.1	85.4	Cross-Site Scripting	79	Failure to Preserve Web Page Structure ('Cross-site Scripting')
93.1	85.6	Directory Traversal	22	Path Traversal
57.6	80.1	Link Resolution	59	Improper Link Resolution Before File Access ('Link Following')
51.8	75.3	Format String	134	Uncontrolled Format String
60.1	57.6	Buffer Overflow	119	Failure to Constrain Operations within the Bounds of a Memory Buffer
29.7	49.3	Resource Management	399	Resource Management Errors
24.9	54.5	Cross-Site Request Forgery	352	Cross-Site Request Forgery (CSRF)
33.1	18.6	Information Leak	200	Information Leak (Information Disclosure)
28.0	18.0	Cryptography	310	Cryptographic Issues
12.1	38.7	Credentials Management	255	Credentials Management
14.2	8.7	Arbitrary Code	94	Failure to Control Generation of Code ('Code Injection')

Table III
PRECISION AND RECALL FOR MAPPABLE CVEs. THE FIRST COLUMNS ARE FOR PRECISION ('P') AND RECALL ('R').

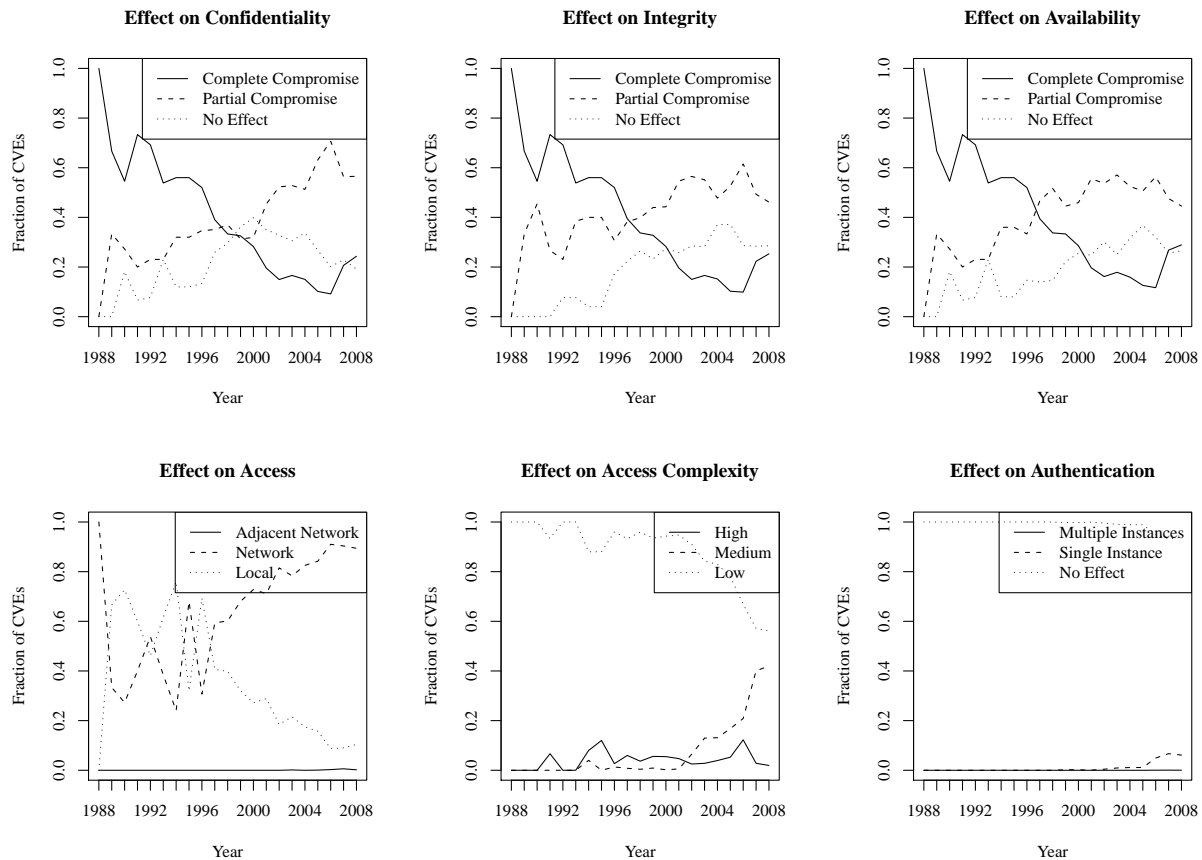


Figure 4. Development of CVSS impact vectors until 2009.

yield different topics on different runs. While not having conducted a systematic investigation, we ran LDA a number of times with different random seeds and, apart from different topic numbers, got the same topics that we report in this study.

The CVE database is very uneven when it comes to the quality of the vulnerability descriptions. They range from almost meaningless entries such as “WinGate is being used” (CVE 1999-0657; see above) to fairly complete analyses of the problem, including the root cause. This could lead LDA to skew its analysis towards those CVEs that are better reported.

Different vendors may have different terminologies when describing vulnerabilities. This may lead to essentially identical vulnerabilities being assigned different topics. This threat is to some extent mitigated by manually assigning labels to topics and combining topics.

Different vendors may also have different disclosure strategies, which could skew the trends we publish. There is evidence for this because December 31 has an unusually high number of CVEs: 7.7% of all CVEs (3051) were published on that day, where publishing CVEs more regularly would give only 1/365, or 0.28%. (Other end-of-month dates do not have such a disproportionately high publishing frequency). This threat is mitigated by aggregating CVEs by year, while at the same time introducing the threat that trends within a year or trends spanning adjacent years will be obscured. On the other hand, there are concerns with the publication date. For example, CVE 2001-1586 has a publication date of February 12, 2010 (sic!). There are 340 such entries (CVE ID having a year that is less than 2010, but the `published_datetime` field being in 2010), but an explanation for this is still being sought.

V. RELATED WORK

The study that is closest to this work is of course the paper by Christey et al., also on CVE trends [1]. The main differences to our study are:

- 1) They analysed the data up to and including 2007, whereas we analysed the data up to 2009.
- 2) They used manual classification of CVE entries, also using information that is not public, even though it is not clear from the report text *how* they categorised the CVEs. The categories appear associated with CWE numbers, but some of the CVE numbers used in the report are not in the 19 CWE numbers that are provided for CVE entry, such as CWE 415 (double-free vulnerability). A request for clarification by email seems to indicate largely manual classification based on indicator words, phrases in the text (like manual topic models), or non-public information.
- 3) We did not differentiate between open vs. closed source or between different operating systems.

Mainly, our findings agree with this report: buffer overflows are still high on the list, but in decline when compared to web application attacks like cross-site scripting or SQL injection; a low number of causes are responsible for the majority of CVE entries; and decline in link following and directory traversal; regular appearance of information leaks.

However, we disagree on the interpretation of cross-site request forgery. Christey et al. state that it is a sleeping giant (emphasis on “sleeping”) and use the low prevalence of 0.1% in 2006 as an argument. We measure 1.8% prevalence in 2006, and also the growth rate, while negative in the last year, is by no means negligible. Given the ease with which CSRF might be exploited, we should definitely keep an eye on this sleeping giant, lest he wake up.

We seem also to have uncovered a category of attacks on application servers that seems to be rising faster than any other attack category. This category does not appear in the CWEs and hence is missing from the MITRE study.

Another large-scale study is regularly done by Microsoft; the most recent example is the Microsoft Security Intelligence Report for H1 2009 [10]. The report uses a vast array of data sources, ranging from online sources like Bing and Windows Live Hotmail to programs collecting data such as the Malicious Software Removal Tool or various filters in Internet Explorer. Overall, the report tends to focus more on malware and browser exploits.

Both Li et al. [11] and Ozment et al. [12] studied in 2006 how the numbers of defect and security issues evolved over time. While Li et al. reported an increase, Ozment reported a decrease in the rate at which new vulnerabilities are reported; for our data, we observed an increase in vulnerabilities until 2006 and since then a decrease. Li et al. used supervised techniques such as Support Vector Machines and Naive Bayes to classify software defects; for security vulnerabilities they used manual classification. In contrast our work uses LDA, an automated and unsupervised learning technique for vulnerability data.

Coverity used its static analysis tools to scan a large number of open-source programs and concluded that the quality of open source software is increasing [13]. We have not yet tried to separate our CVE data into entries pertaining either to open source or to closed source software.

Eric Rescorla looked at the related problem of vulnerability discovery and, by modeling the discovery process and estimating model parameters, concludes that finding vulnerabilities is often without a clear effect on its lifetime [14]. Also, he noticed that vulnerability data is very heterogeneous. His conclusion was therefore that vulnerability discovery should be de-emphasized the quality of gathered data be increased.

Scarfone and Mell compared the scores for CVSS version 1 with CVSS version 2 for over 11,000 CVEs [15]. They found that v2 has higher scores than v1 and that v2 has greater score diversity than v1. Thus they believe that

v2 successfully addressed deficiencies of the original version v1. Fruehwirth and Männistö analyzed 720 CVEs from the NVD and found that adding context information to CVSS scores significantly improved the prioritization and selection of the vulnerability response process [16].

VI. CONCLUSIONS AND FUTURE WORK

We studied the Common Vulnerability and Exposures from the National Vulnerability Database by using topic models on their description texts. Our results include the following.

- Eliminating cross-site scripting, SQL injection, and buffer overflows, and making PHP more secure will eliminate the majority of all CVEs.
- Application server vulnerabilities have massive growth rates and are probably the “next big thing” for years to come.
- Cross-site request forgery is indeed a sleeping giant, and it is probably already stirring.

Our contribution is not intended to replace manual methods, but rather to complement them: our findings by and large support the analysis done by MITRE [1]. By providing a methodically completely independent way, we gain mutually supporting evidence that the trends that we are seeing are actually real. Also, while the analysis in this paper is only a snapshot in time, this analysis is mostly automated and can thus be easily repeated. To facilitate replication, we provide an archive file with all the data and scripts (see Appendix VII).

Another important point is that this work is not another “Top- n vulnerabilities”-type study such as the CWE/SANS Top 25 Programming Errors¹³ or the OWASP Top 10 Project¹⁴. The main difference is that Top- n lists are awareness instruments based on consensus, not necessarily on actual data, like our study.

Our future work will include replication on other corpora like the SANS Consensus Security Alerts [17]. Many of these corpora have fewer but longer documents, so one question that could be investigated is whether the same trends emerge. In addition, we plan to distinguish between open-source and closed-source projects in our topics. For open-source projects CVEs often link to specific bug report, which could help to obtain the actual change that fixed the vulnerability. Thus another data source that we plan to tap into are bug databases and version archives.

VII. ACKNOWLEDGMENTS

Stephan Neuhaus was supported by the European Commission under contract No. 216917, project EU-FP7-IST-IP-MASTER.

¹³<http://www.sans.org/top25-programming-errors/>

¹⁴http://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project

REFERENCES

- [1] S. M. Christey and R. A. Martin, “Vulnerability type distributions in CVE,” <http://cwe.mitre.org/documents/vuln-trends/index.html>, May 2007.
- [2] P. Mell, K. Scarfone, and S. Romanosky, “The complete guide to the common vulnerability scoring system 2.0,” National Institute of Standards and Technology, Tech. Rep., Jun. 2007.
- [3] M. F. Porter, “An algorithm for suffix stripping,” *Program*, vol. 14, no. 3, pp. 130–137, 1980.
- [4] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent dirichlet allocation,” *Journal of Machine Learning Research*, vol. 2003, no. 3, pp. 993–1022, January 2003.
- [5] D. Blei and J. McAuliffe, “Supervised topic models,” in *Advances in Neural Information Processing Systems 20*, J. Platt, D. Koller, Y. Singer, and S. Roweis, Eds. Cambridge, MA: MIT Press, 2008, pp. 121–128.
- [6] D. M. Blei and J. D. Lafferty, “Dynamic topic models,” in *ICML '06: Proceedings of the 23rd international conference on Machine learning*. New York, NY, USA: ACM, 2006, pp. 113–120.
- [7] X. Wang and A. McCallum, “Topics over time: a non-markov continuous-time model of topical trends,” in *KDD '06: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, NY, USA: ACM, 2006, pp. 424–433.
- [8] D. Hall, D. Jurafsky, and C. Manning, “Studying the history of ideas using topic models,” in *Proceedings from the EMNLP 2008: Conference on Empirical Methods in Natural Language Processing*, October 2008, pp. 363–371.
- [9] E. H. Spafford, “The internet worm program: An analysis,” Purdue University, West Lafayette, IN 47907-2004, Purdue Technical Report CSD-TR-823, 1988.
- [10] R. Boscovich, T. J. Campana, D. Canavor, B. Dang, J. Faulhaber, V. Gullotto, Y. Huang, J. Jones, J. Lambert, T. Lee, Z. Mador, R. Mordani, B. Neerumalla, J. Ness, H. O’Dea, S. Parthasarathy, A. Penta, P. Pottorff, S. Reasor, C. Seifert, A. Shostack, G. Stathakopoulos, A. Stone, S. Wu, and T. Zink, “Microsoft security intelligence report volume 7: January through june 2009,” Microsoft, Inc., Tech. Rep., 2009.
- [11] Z. Li, L. Tan, X. Wang, S. Lu, Y. Zhou, and C. Zhai, “Have things changed now?: an empirical study of bug characteristics in modern open source software,” in *ASIS '06: Proceedings of the Workshop on Architectural and System Support for Improving Software Dependability*. ACM, 2006, pp. 25–33.
- [12] A. Ozment and S. E. Schechter, “Milk or wine: does software security improve with age?” in *USENIX-SS'06: Proceedings of the 15th USENIX Security Symposium*. Berkeley, CA, USA: USENIX Association, 2006, pp. 93–104.
- [13] Coverity, Inc., “Coverity scan open source report 2009,” Coverity, Inc., Tech. Rep., 2009.

- [14] E. Rescorla, “Is finding security holes a good idea?” in *Third Annual Workshop on Economics of Information Security*, Jul. 2006, <http://www.dtc.umn.edu/weis2004/rescorla.pdf>.
- [15] K. Scarfone and P. Mell, “An analysis of cvss version 2 vulnerability scoring,” in *MetriSec’09: Proceedings of the International Workshop on Security Measurements and Metrics*, 2009.
- [16] C. Fruehwirth and T. Männistö, “Improving cvss-based vulnerability prioritization and response with context information,” in *MetriSec’09: Proceedings of the International Workshop on Security Measurements and Metrics*, 2009.
- [17] SANS, “@Risk: The consensus security alert,” <http://www.sans.org/newsletters/risk/>, September 2009.
- [18] (2009, Nov.). [Online]. Available: <http://www.cs.princeton.edu/~blei/lda-c/>
- [19] R Development Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2008.
- [20] E. Dimitriadou, K. Hornik, F. Leisch, D. Meyer, and A. Weingessel, *e1071: Misc Functions of the Department of Statistics (e1071)*, TU Wien, 2009.

APPENDIX A. REPLICATION GUIDE

LDA is a general technique that can be used to analyse topics for all kinds of corpora. Also, the techniques to extract graphs like Figure IV are more or less the same for all corpora. We therefore give a step-by-step guide how to replicate our results on the CVE data or to use the same techniques on other corpora.

- 1) *Gather data*. Gather a corpus of n documents describing vulnerabilities. Each document must carry a time stamp. Also choose the number of topics you wish to estimate, N .
- 2) *Stem words, apply stop words (optional)*. Subject the words in the documents to a stemming algorithm and to a list of stop words.
- 3) *Compute vocabulary and word counts*. The result of the previous step(s) is a set of n documents containing a total of m unique words. Enumerate the unique words from 1 to m and build a matrix $\langle m_{jk} \rangle$ where m_{jk} is the number of occurrences of word k in document j . This matrix will generally be sparse.
- 4) *Run LDA*. Now run the LDA software on $\langle m_{jk} \rangle$. We used David Blei’s LDA implementation for C [18]. The result will be a matrix $\langle z_{jl} \rangle$ ($1 \leq j \leq n$ and $1 \leq l \leq N$) where z_{jl} is the topic that was assigned to word l in document j .
- 5) *Compute post-hoc probabilities*. Compute the post-hoc probabilities according to Equations (1), (2), (5) and (6).
- 6) *Join topics (optional)*. If you find that some of the topics produced by LDA are really equal, join them

using Equations (3) and (4). The order of this and the previous step can be interchanged; in this case, you will have to devise a mapping from original topics to joint topics and adjust $\langle z_{jl} \rangle$ accordingly. The result will however be the same.

- 7) *Plot results*. At this point, you have for each unique topic its overall importance, its average annual change, its development in importance during the relevant time period, and a list of the most frequent words in the topic. We used R [19] and the e1071 package [20] to make plots like Figure 3, and a number of custom Perl scripts to create tables like Table VII. We used the sparklines feature of Microsoft Excel 2010 to create the sparklines of Tables IV–VI

To facilitate replication of this study, an archive file with all the data and scripts is available from <http://www.disi.unitn.it/~neuhaus/archive.tar.gz>. The archive does not contain third-party software such as LDA or R.

Table IV
RELATIVE IMPORTANCE, ALL 28 TOPICS. TOPICS ARE ORDERED ALPHABETICALLY, FOLLOWED BY UNNAMED TOPICS. THE y AXES ARE ALL EQUAL, TO FACILITATE COMPARISON.

Topic	Trend	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
Application Servers		1%	1%	1%	1%	1%	1%	3%	3%	3%	5%
Arbitrary Code		5%	4%	5%	5%	5%	5%	6%	7%	7%	6%
Arbitrary Code (IE)		3%	2%	2%	2%	2%	1%	1%	1%	1%	1%
Arbitrary Code (PHP)		0%	1%	1%	2%	2%	3%	11%	8%	3%	2%
Buffer Overflow		19%	18%	17%	18%	15%	10%	8%	10%	8%	11%
CVE Issues		1%	2%	2%	2%	2%	3%	4%	4%	2%	2%
Credentials Management		5%	6%	6%	5%	4%	4%	3%	3%	4%	3%
Cross-Site Request Forgery		2%	2%	2%	2%	2%	3%	2%	3%	3%	3%
Cross-Site Scripting		0%	1%	6%	6%	8%	13%	13%	9%	10%	10%
Cryptography		2%	2%	2%	2%	1%	2%	1%	1%	1%	1%
Directory Traversal		6%	8%	5%	4%	4%	4%	4%	4%	5%	4%
Firewalls		1%	2%	2%	2%	2%	1%	1%	2%	2%	2%
Format String		2%	3%	2%	2%	3%	2%	1%	1%	1%	1%
Information Leak		2%	2%	3%	3%	3%	3%	5%	4%	4%	3%
Java		3%	1%	2%	2%	2%	2%	2%	2%	3%	3%
Link Resolution		6%	6%	4%	4%	4%	3%	1%	2%	3%	1%
Linux		2%	2%	2%	3%	3%	3%	2%	3%	2%	3%
Message Boards		2%	1%	2%	1%	2%	2%	2%	1%	1%	1%
Microsoft Office		1%	1%	1%	1%	1%	1%	1%	1%	2%	2%
Microsoft Windows		4%	3%	3%	3%	2%	1%	1%	1%	1%	2%
Mozilla		1%	2%	2%	2%	3%	2%	2%	2%	2%	3%
PHP		0%	0%	1%	1%	1%	3%	4%	3%	4%	4%
Privilege Escalation		12%	10%	8%	8%	5%	4%	3%	3%	3%	2%
Resource Management		14%	12%	12%	13%	13%	9%	6%	8%	7%	10%
SQL Injection		1%	1%	2%	2%	4%	11%	10%	8%	17%	10%
N/A		1%	1%	1%	1%	1%	1%	1%	1%	1%	1%
N/A		4%	5%	5%	3%	3%	2%	2%	2%	1%	2%
N/A		0%	0%	1%	1%	2%	1%	2%	2%	2%	2%

Table V
RELATIVE IMPORTANCE OF DISCOVERED CAUSES.

Topic	Trend	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
ActiveX		1%	2%	2%	2%	2%	2%	1%	4%	3%	2%
Application Servers		3%	2%	4%	3%	2%	1%	1%	1%	1%	3%
Buffer Overflow		17%	14%	14%	17%	13%	8%	6%	8%	7%	10%
Cisco IOS		3%	4%	5%	3%	3%	2%	1%	2%	1%	2%
Credentials Management		4%	4%	3%	2%	2%	1%	1%	2%	2%	2%
Cross-Site Request Forgery		1%	1%	1%	1%	1%	1%	1%	2%	3%	3%
Cross-Site Scripting		11%	9%	15%	14%	17%	21%	22%	16%	18%	17%
Directory Traversal		6%	12%	7%	6%	5%	5%	6%	6%	7%	5%
Format String		2%	2%	2%	3%	3%	2%	1%	2%	1%	2%
Internet Explorer		2%	3%	3%	4%	3%	3%	4%	3%	2%	2%
Java		1%	1%	1%	2%	2%	2%	2%	2%	2%	2%
Microsoft Office		2%	2%	2%	1%	1%	1%	1%	1%	1%	2%
Mozilla		2%	2%	2%	2%	2%	2%	1%	1%	2%	4%
PHP		5%	6%	6%	8%	8%	10%	19%	17%	10%	8%
SQL Injection		10%	8%	11%	12%	15%	23%	21%	19%	28%	21%
Security Appliances		1%	3%	2%	3%	2%	1%	1%	1%	1%	1%
N/A		4%	2%	2%	1%	2%	2%	1%	2%	1%	1%
N/A		5%	5%	4%	3%	3%	2%	1%	2%	1%	2%
N/A		4%	3%	2%	3%	3%	1%	1%	1%	1%	2%
N/A		2%	3%	3%	2%	2%	2%	1%	2%	2%	2%
N/A		2%	1%	1%	1%	2%	1%	1%	1%	1%	2%
N/A		4%	3%	3%	3%	3%	2%	2%	2%	2%	3%
N/A		4%	2%	3%	2%	2%	2%	2%	2%	2%	2%
N/A		3%	4%	2%	2%	2%	2%	1%	2%	1%	1%

Table VI
RELATIVE IMPORTANCE OF DISCOVERED IMPACTS.

Topic	Trend	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
Arbitrary Code		15%	18%	20%	25%	22%	17%	23%	24%	17%	24%
Arbitrary Script		17%	17%	20%	21%	24%	39%	37%	31%	43%	35%
CVE Dispute		1%	2%	2%	1%	2%	3%	5%	5%	3%	5%
Denial of Service		30%	23%	21%	22%	21%	13%	8%	12%	9%	11%
Information Leak		22%	23%	18%	13%	14%	13%	14%	16%	17%	15%
Information leak		1%	1%	3%	4%	3%	4%	3%	2%	1%	1%
Privilege Escalation		11%	12%	11%	10%	10%	8%	6%	7%	8%	7%
Resource Abuse		2%	2%	2%	1%	1%	1%	1%	1%	1%	1%
N/A		1%	2%	2%	2%	1%	1%	1%	1%	1%	0%
N/A		0%	0%	0%	0%	0%	0%	1%	1%	0%	0%
N/A		0%	0%	0%	0%	0%	1%	1%	1%	1%	1%
N/A		0%	0%	0%	0%	0%	0%	1%	1%	0%	0%

Name	%						
Buffer Overflow	11.33	overflow 6.29 remot 4.29	execut 6.15 command 3.05	buffer 5.74 long 2.83	arbitrari 5.55 craft 1.31	code 4.77 file 1.30	
Cross-Site Scripting	9.20	script 14.31 inject 6.62	html 7.26 vulner 6.48	cross-sit 7.15 arbitrari 6.34	xss 6.92 remot 5.73	web 6.73 paramet 4.47	
Resource Management	9.13	caus 10.33 applic 1.41	servic 9.96 packet 1.26	denial 9.44 os 1.24	remot 4.94 larg 1.17	crash 4.25 server 1.10	
SQL Injection	8.91	sql 14.53 execut 6.68	inject 7.65 arbitrari 6.58	paramet 7.31 remot 6.23	command 6.99 php 5.14	vulner 6.69 index 1.65	
Arbitrary Code	5.97	php 3.40 control 1.98	remot 3.24 code 1.56	arbitrari 2.57 vulner 1.51	execut 2.31 dll 1.43	modul 2.11 argument 1.33	
Privilege Escalation	4.99	user 6.60 access 1.64	local 5.76 file 1.26	privileg 5.74 program 1.25	gain 5.46 bypass 1.21	root 2.05 system 1.20	
Arbitrary Code (PHP)	4.36	php 20.14 includ 5.15	remot 10.17 arbitrari 5.10	code 5.99 paramet 5.10	file 5.88 vulner 4.95	execut 5.21 url 4.94	
Directory Traversal	4.17	director 9.93 remot 6.28	dot 9.31 vulner 5.87	file 8.63 read 4.14	travers 6.71 paramet 3.16	arbitrari 6.51 sequenc 1.94	
Credentials Management	3.73	authent 7.53 access 3.31	password 5.21 gain 2.71	user 5.21 bypass 2.43	administr 5.09 privileg 2.38	remot 4.58 account 2.34	
Information Leak	3.47	inform 10.51 third 4.30	obtain 7.29 error 3.89	detail 4.63 remot 2.96	note 4.39 unknown 2.69	parti 4.31 messag 2.64	
CVE Issues	2.96	note 7.54 might 1.22	issu 4.96 vulner 0.98	disput 2.17 but 0.91	cve 1.52 refer 0.82	candid 1.50 report 0.81	
Link Resolution	2.84	file 16.18 creat 4.22	user 7.23 symlink 3.41	local 6.60 temporari 2.87	arbitrari 5.05 tmp 1.88	overwrit 4.28 read 1.64	
PHP	2.64	php 21.48 inc 2.50	paramet 6.38 multipl 2.01	admin 3.38 remot 1.87	action 2.70 includ 1.49	index 2.65 class 1.08	
Topic 17	2.44	remot 4.21 configur 1.93	address 3.46 access 1.85	connect 2.35 port 1.40	server 2.29 bypass 1.20	ip 2.21 firmwar 1.18	
Linux	2.44	kernel 4.81 linux 3.54	function 4.21 servic 3.10	caus 3.93 denial 2.86	user 3.82 null 2.20	local 3.67 memori 1.92	
Cross-Site Request Forgery	2.43	request 6.30 direct 3.95	inform 5.73 web 3.31	remot 5.32 access 3.12	sensit 4.86 store 2.91	obtain 4.37 password 2.46	
Application Servers	2.27	oracl 7.21 unknown 4.25	unspecifi 6.64 applic 3.69	vector 5.45 server 3.40	compon 5.39 remot 3.31	vulner 5.37 ibm 2.85	
Java	2.00	file 7.08 arbitrari 2.98	upload 5.26 earlier 2.93	java 3.86 execut 2.78	remot 3.79 vulner 2.27	extens 3.79 jre 1.92	
Mozilla	1.95	remot 3.99 web 2.08	mozilla 3.12 url 1.75	firefox 3.03 arbitrari 1.68	javascript 2.11 domain 1.62	site 2.10 document 1.59	
Firewalls	1.79	cisco 4.58 through 1.68	remot 3.10 aka 1.35	secur 2.49 commun 1.30	devic 1.96 io 1.25	packet 1.79 servic 1.24	
Microsoft Windows	1.57	server 8.22 aka 2.07	window 7.28 xp 2.03	microsoft 3.53 request 1.95	remot 3.04 vulner 1.94	servic 2.27 name 1.63	
Message Boards	1.55	field 8.26 form 2.29	asp 7.89 name 2.00	remot 3.46 paramet 1.96	user 2.59 email 1.85	messag 2.42 board 1.57	
Topic 35	1.48	vector 10.25 relat 5.57	unspecifi 9.82 involv 4.78	unknown 8.24 multipl 2.52	vulner 7.73 possibl 1.80	impact 7.45 probabl 1.00	
Format String	1.41	string 8.31 code 4.58	format 8.24 function 3.85	arbitrari 5.45 specifi 2.77	execut 5.15 remot 2.65	vulner 4.97 earlier 2.12	
Arbitrary Code (IE)	1.33	internet 8.41 vulner 2.10	explor 7.14 aka 2.06	microsoft 4.48 execut 1.91	remot 4.02 arbitrari 1.80	window 2.74 object 1.78	
Microsoft Office	1.29	file 4.25 vulner 3.45	microsoft 4.07 code 3.37	offic 3.96 execut 3.12	corrupt 3.76 arbitrari 2.80	memori 3.53 aka 2.58	
Cryptography	1.24	file 3.33 remot 2.04	password 3.01 forc 1.96	encrypt 2.83 flash 1.63	make 2.16 gener 1.28	easier 2.07 adob 1.25	
Topic 7	1.12	ftp 7.68 other 2.24	server 4.78 adob 2.19	command 4.09 arbitrari 1.82	remot 3.10 reader 1.68	possibl 2.82 vector 1.56	

Table VII
THE TEN MOST FREQUENT WORD FOR THE 28 TOPICS IDENTIFIED BY LDA; SEE SECTION III-B.

Name	%					
SQL Injection	20.05	vulner 9.91 multipl 2.53	inject 7.95 asp 1.78	sql 7.67 script 1.20	php 5.25 index 1.08	earlier 2.77 cross-sit 0.84
Cross-Site Scripting	17.84	vulner 11.19 multipl 3.54	script 10.94 earlier 2.41	cross-sit 10.42 index 1.17	xss 10.09 cgi 1.06	php 4.31 modul 0.86
PHP	11.82	php 12.62 earlier 3.53	vulner 8.68 multipl 1.97	file 6.26 inject 1.68	remot 5.45 sql 1.67	includ 5.22 index 1.64
Buffer Overflow	9.25	overflow 14.09 appl 2.05	buffer 13.60 multipl 1.92	stack-bas 3.29 os 1.70	server 2.70 heap-bas 1.68	function 2.35 mac 1.57
Directory Traversal	5.92	director 12.70 index 1.67	vulner 10.96 multipl 1.52	travers 10.94 server 1.11	php 6.73 file 0.83	earlier 4.36 enabl 0.82
Internet Explorer	2.84	earlier 7.93 microsoft 3.44	vulner 6.89 file 2.59	internet 6.37 remot 2.19	explor 6.20 includ 2.04	php 5.61 inject 1.86
Topic 38	2.27	access 8.88 under 4.66	store 5.63 insuffici 4.37	web 5.49 inform 3.81	control 5.28 password 3.52	root 4.99 sensit 3.28
ActiveX	2.22	control 12.62 dissector 2.13	activex 10.54 ether 1.97	dll 6.46 certain 1.67	vulner 3.48 multipl 1.66	ocx 2.60 unspecifi 1.38
Topic 27	2.14	window 8.71 weblog 2.80	server 6.83 properli 2.67	microsoft 5.71 explor 2.65	internet 4.19 through 2.61	xp 2.82 bea 2.30
Cisco IOS	2.01	cisco 8.21 run 2.36	apach 3.96 router 2.29	through 3.26 configur 1.61	io 2.97 seri 1.50	firmwar 2.57 modul 1.48
Topic 8	2.00	server 5.32 servic 1.87	client 3.26 web 1.78	window 2.41 exchang 1.64	connect 2.13 novel 1.61	microsoft 1.95 agent 1.32
Format String	1.91	vulner 4.58 differ 1.61	string 4.00 php 1.53	format 3.88 messag 1.40	earlier 2.92 gener 1.34	function 2.28 through 1.25
Topic 19	1.86	server 2.30 manag 1.32	commun 1.85 vulner 1.23	http 1.79 unspecifi 1.22	session 1.68 through 1.18	cooki 1.57 web 1.18
Mozilla	1.85	mozilla 4.28 acrobat 2.77	firefox 4.09 seamonkei 2.04	adob 3.88 vulner 1.99	reader 2.91 thunderbird 1.92	engin 2.86 overflow 1.87
Credentials Management	1.83	build 5.39 administr 2.61	manag 3.79 hp 2.42	authent 3.42 network 2.06	configur 2.74 requir 1.61	server 2.74 default 1.50
Java	1.81	earlier 7.24 updat 3.30	vulner 5.06 jdk 2.63	jre 5.04 sun 1.94	php 4.68 sdk 1.65	java 3.91 environ 1.63
Application Servers	1.72	applic 7.58 webspher 2.12	server 5.22 vulner 1.96	oracl 4.52 unspecifi 1.85	compon 3.19 web 1.58	ibm 2.65 through 1.19
Cross-Site Request Forgery	1.69	vulner 8.45 iphon 2.46	cross-sit 7.01 script 2.43	request 6.13 appl 2.13	forgeri 5.60 xss 2.11	csrf 5.60 through 1.95
Topic 11	1.66	check 5.11 certain 1.70	properli 4.51 file 1.58	function 2.45 could 1.51	earlier 2.44 error 1.47	handl 1.87 vulner 1.35
Topic 40	1.63	server 5.90 ftp 1.75	earlier 4.89 hitachi 1.38	vulner 2.92 director 1.14	through 2.66 manag 1.05	web 2.36 enabl 0.97
Topic 2	1.63	ca 3.39 comput 2.10	backup 3.18 unspecifi 1.79	earlier 2.79 certif 1.70	associ 2.29 vulner 1.59	through 2.21 manag 1.41
Topic 23	1.47	server 3.87 system 2.79	earlier 3.84 function 1.54	java 3.45 decod 1.53	vulner 3.29 multipl 1.46	sun 3.10 web 1.43
Security Appliances	1.31	applianc 3.44 earlier 1.84	secur 3.00 other 1.77	cisco 2.38 vulner 1.54	version 1.99 through 1.37	possibl 1.86 server 1.36
Microsoft Office	1.20	offic 10.34 gold 3.02	microsoft 7.44 word 2.75	excel 3.50 format 2.04	viewer 3.17 powerpoint 2.03	file 3.05 xp 1.96

Table VIII
THE TEN MOST FREQUENT WORDS FOR THE 20 TOPICS IDENTIFIED BY LDA ON THE "CAUSE" DATA.

Name	%						
Arbitrary Script	33.59	arbitrari 10.75 script 5.06	paramet 8.94 html 4.91	command 6.32 inject 4.86	execut 6.28 web 4.73	sql 5.14 php 2.42	
Arbitrary Code	20.80	code 10.56 url 3.02	execut 10.37 long 2.63	arbitrari 10.13 file 2.13	php 5.06 craft 1.67	paramet 3.43 possibl 1.03	
Information Leak	15.49	file 7.52 read 2.97	dot 5.69 obtain 2.93	arbitrari 4.76 detail 1.99	inform 3.64 note 1.85	paramet 3.21 parti 1.84	
Denial of Service	13.05	caus 11.07 request 1.54	servic 10.54 long 1.44	denial 10.10 applic 1.27	crash 4.68 possibl 1.19	packet 1.90 consumpt 1.12	
Privilege Escalation	7.79	bypass 5.82 administr 2.96	access 4.66 privileg 1.97	authent 3.45 paramet 1.82	gain 3.35 cooki 1.76	asp 3.05 restrict 1.63	
CVE Dispute	3.64	note 7.98 affect 1.62	issu 4.80 disput 1.51	vector 1.65 cve 1.46	vulner 1.62 report 1.45	might 1.62 bug 1.06	
Information leak	2.27	obtain 9.26 messag 6.38	inform 8.92 request 4.49	sensit 8.22 reveal 3.88	error 7.68 invalid 2.97	path 6.97 paramet 2.95	
Resource Abuse	1.01	extens 7.16 document 3.12	file 6.44 exe 3.10	filenam 4.01 upload 2.91	attach 3.73 demonstr 2.49	contain 3.57 messag 2.40	
Topic 23	0.90	domain 2.76 php 1.66	log 2.21 server 1.62	demonstr 1.84 user 1.55	other 1.80 window 1.52	charact 1.70 script 1.45	
Topic 34	0.59	php 13.43 modul 2.28	paramet 8.05 search 2.01	comment 2.84 form 1.80	index 2.54 languag 1.36	archiv 2.42 help 1.17	
Topic 25	0.45	php 31.95 vector 1.77	cover 3.49 includ 1.54	alredi 3.24 paramet 1.16	lib 3.20 auth 1.07	note 2.74 common 0.83	
Topic 36	0.42	php 34.41 index 1.16	inc 15.44 footer 0.90	includ 1.82 header 0.81	paramet 1.72 modul 0.75	admin 1.67 plugin 0.72	

Table IX
THE TEN MOST FREQUENT WORDS FOR THE 12 TOPICS IDENTIFIED BY LDA ON THE "IMPACT" DATA.