The Development of a Common Enumeration of Vulnerabilities and Exposures

David W. Baker
bakerd@mitre.org

Steven M. Christey
coley@mitre.org

William H. Hill
bill@mitre.org

David E. Mann
damann@mitre.org

The MITRE Corporation
1820 Dolley Madison Boulevard
McLean, Virginia 22102

Presented at the Second International Workshop on Recent Advances in Intrusion Detection,

7-9 September 1999

Abstract

This paper traces the development of a Common Enumeration of Vulnerabilities and Exposures (CVE) that standardizes and lists vulnerabilities and security exposures to facilitate data sharing and comparison across computer vulnerability databases, such as those produced by security tools and academic research.

The MITRE Corporation is building a system that can integrate and manage vulnerability information from different sources (e.g., network assessment tools, intrusion detection systems [IDSs], archives) in a database for supporting enterprise security operations. However, every information security tool considered for integration has its own vulnerability database. Also, the lack of common naming conventions and a common enumeration of the vulnerabilities in the vulnerability databases hindered integration efforts. Thus, MITRE developed CVE to provide a common vocabulary for its vulnerability database system effort.

The CVE concept was first proposed in January 1999, at Purdue's Center for Education and Research for Information Assurance and Security (CERIAS) 2nd Workshop on Research with Security Vulnerability Databases, in a paper titled *Towards a Common Enumeration of Vulnerabilities*, by Steven M. Christey and David E. Mann. CVE provides a mechanism for information security community discussion on vulnerability identification and other related security issues.

CVE development was broadened by creating a CVE Editorial Board, which includes information security community representatives from tool vendors, research and educational organizations, MITRE, and others. The CVE Editorial Board is currently enumerating a large number of vulnerabilities, while simultaneously attempting to capture and codify the decision-making process. When a significant number of vulnerabilities are validated and verified, an initial version of CVE will be released to the public.

The document includes background information on MITRE's early CVE activities, a draft CVE design, CVE content and use, and lessons learned.

Presented at the Second International Workshop on
Recent Advances in Intrusion Detection,

7-9 September 1999

The Development of a Common Enumeration of Vulnerabilities and Exposures

Presented at the Second International Workshop on
Recent Advances in Intrusion Detection,

7-9 September 1999

1 INTRODUCTION

In the evolving world of computer and network security, continuous advances have been made in tool development and techniques for both compromising and protecting hosts. The recent growth in the number and quality of security products is an indicator of the perceived need for further means to protect computer systems from compromise. Effective use of these tools involves the goals of exchange, interpretation, and correlation of a large amount of information about computer vulnerabilities, yet these goals are very difficult to achieve. One of the major stumbling blocks to achieving these goals is the lack of a common listing or enumeration of computer system vulnerabilities.

The scope of this paper is to provide background information on the need for a Common Enumeration of Vulnerabilities and Exposures (CVE) and MITRE's early CVE activities. A CVE was first proposed in January 1999 at Purdue's Center for Education and Research for Information Assurance and Security (CERIAS) 2nd Workshop on Research with Security Vulnerability Databases, in a paper titled *Towards a Common Enumeration of Vulnerabilities*, by Steven M. Christey and David E. Mann. Subsequently, MITRE took the initiative to produce, implement, and evaluate a Draft CVE design. CVE development was broadened by creating a CVE Editorial Board, which includes information security community representatives from tool vendors, research and educational organizations, MITRE, and others.

CVE provides a mechanism for community-wide discussion on vulnerability identification and other related security issues.

1.1 The Need for a Common Enumeration of Vulnerabilities and Exposures

To protect a computer system, one must be able to:

- ?? Understand what various security tools can and cannot do
- ?? Identify what vulnerabilities exist in the system
- ?? Eliminate those vulnerabilities when possible
- ?? Understand and manage the risk of the remaining vulnerabilities

To convey this information, many tools include a database of security vulnerabilities. There is significant variation in the databases of these tools, and because these databases do not share a common element such as a vulnerability name, there is no way to determine when different tools are referring to the same vulnerability. The results are as follows:

- ?? Security tools are difficult to evaluate in a sufficiently detailed fashion with respect to individual vulnerabilities.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

?? Analysts face a substantial learning curve each time they begin to use a new tool.

?? Correlation among tools made by separate vendors can be difficult to achieve.

?? Difficulties can be encountered when trying to relate defensive information, such as CERT[1] advisories, to the vulnerabilities and systems in question.

?? The lack of common naming conventions for vulnerabilities deters productive discussions on countermeasures and other defensive actions.

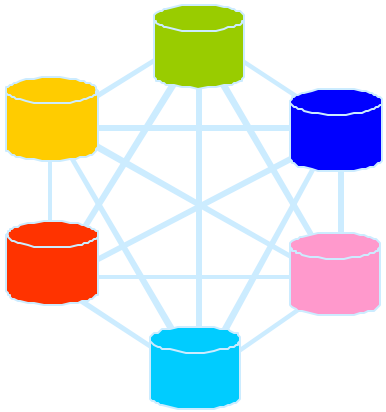

Figure 1-1. Mapping each database to every other database

To achieve interoperability, one could develop individual mappings among products. This effort would quickly become unwieldy, as it would require on the order of N^2 mappings (N=number of vulnerability database sources), as depicted in Figure 1-1.

Alternately, one could develop a comprehensive list of vulnerabilities to use as an intermediary and perform only N database mappings, as depicted in Figure 1-2. CVE could be this intermediary. CVE is defined as a standardized list that enumerates and discriminates among all publicly known vulnerabilities; assigns a standard, unique name to each vulnerability; exists independently of the multiple perspectives that define a vulnerability; and is publicly open and sharable, without distribution restrictions. CVE could help address these problems, because it would perform the following functions:

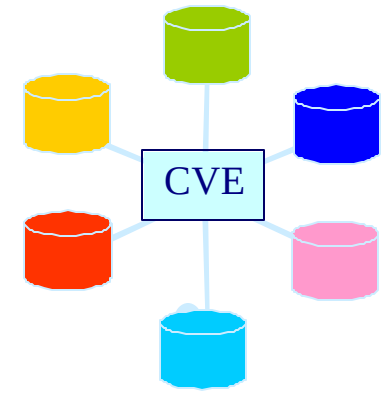

Figure 1-2. Mapping each database to the CVE

?? Facilitate interactions among humans by providing standardized names for reference purposes and a comprehensive vulnerability listing

?? Facilitate tool interoperability, both immediately as early adopters develop private tool integration strategies and inherently as vendors adopt CVE

---

[1] CERT is a registered trademark of the Software Engineering Institute, at Carnegie-Mellon University, Pittsburgh, Pennsylvania

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

?? Allow for evaluation of tools to determine those vulnerabilities the tools recognized

However, several aspects of tool comparison and evaluation would not be addressed using CVE:

?? The severity of the vulnerabilities addressed

?? Local configurations, security policies, and risk postures, and the issues that arise from them

?? How well a particular tool addressed a vulnerability

1.2 Related Work

Two groups are involved in related work, specifically oriented to security tool interoperability. The first group is known as the Common Intrusion Detection Framework (CIDF) working group[2], which includes researchers funded by the Defense Advanced Research Projects Agency (DARPA)[3]. The CIDF working group has defined a protocol and language for exchanging information among IDSs using this information in its research projects.

Another group is the Intrusion Detection Working Group (IDWG)[4]. Its charter, as designated by the Internet Engineering Task Force (IETF)[5], is to define data formats and exchange procedures for sharing information of interest regarding intrusion detection and response systems.

The primary focus of each group has been the development of protocols and languages to facilitate communication among intrusion detection system (IDS) elements. The CIDF has produced a sample list of attack identifiers, but the intent of the list was merely for experimentation rather than a serious attempt at a complete enumeration. Neither group has progressed to the point of developing a list of vulnerabilities or attacks, but any developed protocol or language will require such a list. CVE could fulfill that requirement.

1.3 Why a Public CVE?

The goals of having CVE are to provide a complete enumeration and discriminate among all publicly known vulnerabilities; assign a standard, unique name to each vulnerability; and exist independently of the multiple perspectives of what constitutes a vulnerability. For CVE to have any impact on interoperability and communication, CVE must be openly available to the

---

[2] CIDF information at: http://gost.isi.edu/cidf

[3] DARPA information at: http://www.darpa.mil

[4] IDWG charter at: http://www.ietf.org/html.charters/idwg-charter.html

[5] IETF information at: http://www.ietf.org

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

public, without restrictions on distribution. Some extensive commercially available vulnerability databases now exist; however, most are protected from open use by copyright restrictions.

Certainly, any public discussion of vulnerability information may help a hacker to compromise a system. However, for the following reasons the benefits of a publicly available CVE would outweigh its risks:

- CVE is restricted to publicly known vulnerabilities.
- Sharing information is more difficult within the information security community than it is for hackers (e.g., commercially funded databases have copyright issues precluding their use).
- It takes much more work for an organization to protect its networks and eliminate all possible security holes than for a hacker to find a single vulnerability, exploit it, and compromise the network.
- Community opinion is shifting towards sharing information, as reflected in the fact that the CVE Editorial Board includes key organizations in the community.

A widely accepted common enumeration would enhance tool interoperability, enable effective communication, and create a more cohesive security posture, thus reducing the overall risk of compromise.

1.4 MITRE's Early CVE Work

1.4.1 Corporate Introduction

In partnership with government clients, MITRE is a not-for-profit corporation working in the public interest. It addresses issues of critical national importance, combining systems engineering and information technology. Our work in the protection of government systems and critical infrastructure assets has clearly illustrated the need for improved defensive information. By improving the ability of system administrators and security staff to communicate clearly and effectively concerning vulnerabilities, and to effectively select tools that suit their needs, the ability to defend against attack is improved.

1.4.2 Identifying the need for CVE

The work discussed in this paper was begun by MITRE's Information Security Committee, which oversees internal corporate security activities. MITRE uses a variety of tools to assess its computer and network security posture. These tools include several network assessment tools and multiple IDSs from a number of different vendors. Each tool represents, identifies, or processes vulnerabilities in one form or another; and each tool uses its own implicit or explicit database, which defines the vulnerabilities the tool

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

recognizes. MITRE is building a system to integrate and manage vulnerability information from different sources in a centralized database, which can be linked to internal databases for supporting enterprise security operations.

For this purpose, the following capabilities were considered essential:

- When vulnerability information is received from multiple assessment and/or IDS tools, a determination must be made if that information identified the same vulnerability.
- When an IDS alarm is received, a determination must be made of the status for that host in terms of recent assessment scans for the related vulnerability.
- When a network assessment reveals a host is susceptible to a particular vulnerability, collation and distribution of all information pertaining to that vulnerability must take place.
- Tools must be compared in terms of their completeness, how many vulnerabilities identified in CERT advisories are detected, and what gaps exist in vulnerability detection by the tools utilized.

Four separate challenges can be identified in trying to achieve commonality among all security tools and other vulnerability information sources:

- Inconsistent naming conventions
- Management of similar information from diverse sources
- Management of multiple evolving perspectives of the same vulnerability
- Complexity of mapping among databases

MITRE conducted an effort to represent vulnerability information by studying taxonomies and current vulnerability databases. Each source of information related to vulnerabilities was based on a different perspective. These disparate perspectives make it difficult to design schemas that fully capture all of this information. Further, each of these perspectives can evolve over time, requiring modifications to any representation that attempts to unify these concepts. An approach that mitigates these representational discrepancies was desired.

Classification schemas and taxonomies strive to organize sets of information. Several methods attempt to define sharable schemas, universal primitives, and robust taxonomies for vulnerability information. However, no consensus on these methods has been selected. Although this remains a critical research area, the MITRE team decided to meet immediate and operational objectives with a simple, unstructured listing of known vulnerabilities and security exposures: a one-dimensional enumeration of the publicly known vulnerabilities. An enumeration would move MITRE toward achieving a capability that provides immediate interoperability among different databases

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

and tools, and it would foster a greater degree of data exchange across the information security community by implicitly defining a common list of vulnerabilities.

A crucial factor in the effort was the accommodation of multiple perspectives of vulnerabilities and exposures. The perspective of a vulnerability is not based solely upon hardware or software. System configuration, the security policy, and the acceptable risk of an organization all play a role in determining what defines a vulnerability or a security exposure. Identifying a common enumeration that will be useful regardless of perspective was critical. Independence from multiple competing perspectives should be a goal of CVE, such that it has application across multiple views of vulnerabilities and exposures.

The remaining sections address the following topics: a Draft CVE, moving beyond the Draft CVE, and conclusions.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

2 A DRAFT CVE

This section presents a description of the Draft CVE[6], including CERIAS presentation highlights as they pertain to CVE development, design considerations, and a Draft CVE overview.

2.1 CERIAS Presentation Highlights

The concept of CVE was introduced at the Second Workshop on Research with Security Vulnerability Databases, hosted by CERIAS on 20-22 January, 1999. In his opening remarks for the workshop, Dr. Eugene Spafford urged the information security community to adopt a standard for identifying vulnerabilities in a manner similar to the way the virus detection community adopts a common name for viruses. Several vendors were represented at the workshop. The primary focus of the workshop centered on the question of how to construct a vulnerability database that would be sharable across the trusted information security community. Such a database may implicitly include what we are calling CVE.

As a result of the discussions at the CERIAS workshop, additional topics were brought forward to enable the development of CVE. Perhaps the most significant was the need for the establishment of a CVE Editorial Board. This Editorial Board would represent multiple perspectives of vulnerabilities, provide a broad level of experience in the different aspects of vulnerability identification and enumeration, and serve as a forum for discussion and development of a CVE that would fulfill expectations. Another topic was the desire to capture consensus where possible among the differing perspectives of the various views of vulnerabilities. Also, the need for consistency in the decision-making process for identifying and enumerating vulnerabilities and exposures was addressed. To bring CVE to an operational state and be able to maintain its currency with constantly emerging vulnerabilities, CVE must be manageable. Last, an operational CVE must be achievable within a reasonable period of time.

2.2 Design Considerations

The following CVE design considerations were identified:

?? Simplicity

?? Completeness

---

[6] We differentiate here between the Draft CVE (the initial attempt at developing CVE, produced by MITRE), and CVE, the current enumeration project being developed by the CVE Editorial Board.

?? Public availability

?? Enabling interoperability

?? Independence from multiple competing perspectives

To function as a cross-reference, an entry listed in CVE does not need to include any attributes beyond a name and a textual description, as long as the name is unique and the description includes enough information for a human reader to distinguish one vulnerability from other vulnerabilities. Interoperability can be achieved by mapping multiple vulnerability databases to the vulnerabilities specified in CVE, with CVE serving as a logical bridge. Once mappings are created, information contained in CVE is not actually needed. The power of CVE lies in simplicity, comprehensiveness with respect to enumerating vulnerabilities, and an implicit naming convention.

One of the primary initial decisions was to avoid creating a taxonomy or other formal categorization of vulnerabilities. While developing a taxonomy is normally an essential component in categorizing a topic, it was believed to be a difficult problem and beyond the scope of this effort. Further, since CVE is a simple list, entries would only need to be added to the list. By maintaining this simple approach, most issues about categorizing vulnerabilities would be avoided. The facts that a particular vulnerability was known, validated, and determined to be different from the others would allow it to be assigned a name and entered into CVE. Vulnerability databases can be used to determine how a particular vulnerability needs to be categorized in accordance with the perspective of that database.

To reiterate, CVE was designed to be merely a concordance or index, listing all of the vulnerabilities by a name and then providing enough information to distinguish one from another. While CVE should be comprehensive, it was recognized that lack of consensus and the sheer number of vulnerabilities may preclude CVE from containing every known vulnerability, at least in the early stages of CVE development. However, the intent of CVE is to be complete with regard to all publicly known vulnerabilities.

The decision regarding naming was to keep the name a simple number with a “CVE” prefix. More detailed naming conventions would not add to the usefulness or effectiveness of CVE. Additionally, a more complex naming convention might increase debate or disagreement over issues that have no real relevance to a simple listing, such as CVE.

2.3 Draft CVE Overview

The Draft CVE was compiled from data obtained from a variety of publicly available sources, including a vulnerability database, the lists of vulnerabilities assessed by six commercial and freeware network assessment tools, attack signatures from a commercial IDS, and an exploit database from a well-known hacker site. Also, several vendors contributed to this effort by providing

feedback and vulnerability data. The Draft CVE listed 663 distinct vulnerabilities as of 28 April 1999.

2.3.1 Content

As shown in Table 2-1, CVE uses a simple representation to identify vulnerabilities: a name and a short description. The name must be unique, and the description must have enough information to allow a human reader to distinguish among vulnerabilities.

Table 2-1. Sample listing of a CVE Entry:

Name	Description
CVE-1999-0003	Execute commands as root via buffer overflow in ToolTalk database server (rpc.ttdbserverd)

The CVE name is defined as a simple number with a prefix of "CVE" and the year the vulnerability was listed in CVE. The name should not be assumed to imply a relative order in which vulnerabilities are discovered or to indicate the severity of the vulnerability; it only reflects the order in which the name was placed in CVE.

The description should provide enough information to allow a security expert to distinguish between vulnerabilities and search for the name of a specific vulnerability. Therefore, the description does not need to include all information that is normally associated with a vulnerability in other databases. It is not intended to provide educational value for an end user, such as a system administrator. Other information sources (such as traditional vulnerability databases and security tools) are useful for that purpose.

Optimally, a description should only contain the most specific and distinctive information. To support a search, the description may contain more information than is necessary to distinguish among vulnerabilities.

The nature of CVE descriptions also affects the usefulness of CVE to some types of users. For example, the descriptions themselves do not provide enough information to the system administrator who needs to apply patches for vulnerabilities or understand them better. There is an inherent dependence on other more complete vulnerability databases to provide such information.

2.3.2 CVE Maintenance Extension

An additional set of data closely related to CVE, referred to as the CVE Maintenance Extension (CMEX), was created to provide maintenance information to a human reader and support looking up a specific vulnerability. In general, CMEX would be useful to individuals who wish to develop a map from a database to a CVE entry. As CVE is adopted as a naming convention, most end users will not need to use CMEX. Their primary concern will be the

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

standardized name that links vulnerability information from various sources. The current CMEX contains the following data extensions:

- **Administrative data** - The only metadata associated with CVE at this time is a version number. The version number reflects the date of the most recent modification to any part of CVE. Individual CVE vulnerabilities also have their creation date and modification date.
- **Category** - The category has a direct impact on the content decisions that have been made with respect to that entry. Other than the role these categories play with respect to determining CVE content, they do not connote any commitment to a formal taxonomy on the part of CVE. Entries were organized for ease of initial assessment into five categories:
  - SF - software faults/bugs
  - CF - configuration problems
  - SA - presence of a service (e.g., Finger service is running)
  - MP - evidence of a malicious presence or system compromise
  - AN - anomalous state (e.g., integrity checking violation)
- **Reference** - CMEX also includes references for most CVE entries. References are freely and publicly available documents that describe vulnerabilities. Similar to the descriptive text, a reference helps the human reader more easily distinguish among entries.
- **Thesaurus** - CMEX includes a thesaurus that expands some domain-specific terminology referring to the same concept (e.g., “buffer overflow” and “buffer overrun”). The thesaurus may also include alternate spellings of the same term (e.g., “Win95” and “Windows 95” or “DoS” and “denial of service”).
- **Keywords** - The keywords are derived from the descriptive text. Keywords are useful for searching CVE.

The contents of CMEX do not infer a taxonomy. CMEX was designed purely for maintenance of the CVE, and may be extended as required. In addition, the CMEX was intentionally designed to minimize overlap with vulnerability databases as much as possible.

2.3.3 Early Results

Since CVE was intended to serve as a standard concordance for vulnerabilities, vulnerability databases need to be mapped to CVE names. CVE was not intended to supplant the particular nomenclature used by that database; it is only for use as a translation mechanism. The mappings link database entries with corresponding CVE entry names.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

Mapping to CVE can be a laborious process, especially if the vulnerability database being mapped contains very little information. Some information retrieval techniques can reduce the effort significantly, such as a keyword search. The process should become more efficient as the mapper becomes familiar with the contents and content decisions of CVE.

The mappings from these vulnerability databases to CVE facilitated the creation of detailed comparisons among tools, databases, and published exploits. For example, one comparison explicitly differentiated among assessment tools with respect to coverage of the entries identified in CVE. A different comparison revealed gaps between the tools and the set of exploit scripts available from a well-known hacker site.

Finally, the creation of CVE identified two broad classes of problems for which no clear enumeration or taxonomy presently exists, specifically, configuration problems and usage of unauthorized network services. CVE may provide a mechanism for community-wide discussion on these areas.

During the mapping of various data sources to the Draft CVE, a number of expected issues arose. An example of some of the results of tool mappings might provide some insight into the problems encountered. For example, consider NFS vulnerabilities.

The mappings determined there was inconsistent enumeration of NFS vulnerabilities. Depending on which data source was used, there were different numbers of Network File System (NFS) vulnerabilities, which resulted from the perspective of the originator of the data source. CERT had issued 6 advisories pertaining to NFS vulnerabilities. CyberCop[7] Scanner conducted 13 different checks for NFS vulnerabilities. The X-Force[8] database had 20 entries for NFS vulnerabilities. The INFILSEC[9] vulnerability database had 4 entries for NFS vulnerabilities. Additionally, different names existed for the same vulnerability. One NFS vulnerability was called by three different names: “NFS file guessing check” in CyberCop; “nfs-guess” in X-Force Database; and “SunOS NFS Jumbo and fsirand patches” in a CERT advisory[10].

---

[7] Network Associates Incorporated. 1999. Proprietary Vulnerability Database for CyberCop Scanner 2.4 (CyberCop is a trademark of Network Associates, Incorporated, Santa Clara, California)

[8] Internet Security Services. 1999. Online Database X-Force. Published electronically at http://www.iss.net/xforce

[9] Infilsec Systems Security. 1997. Online Database. Published electronically at http://www.infilsec.com

[10] CERT Coordination Center, 1991. CERT Advisory CA-91:21. Published electronically at http://www.cert.org/ftp/cert_advisories/CA-91%3a21.SunOS.NFS.Jumbo.and.fsirand

It was possible to make some initial comparisons among different tools. For example, a comparison between two different security tools, from a well-known hacker exploit site, and the publicly available CERT advisories, the following information was determined:

?? There are differences in the level of abstraction by the different tools.

?? Gaps were found in tools, in comparison to other tools, hacker sites, and CVE.

?? Gaps were found in CVE itself (tools addressed vulnerabilities not listed in CVE).

?? Mappings provide a filter for objective comparison between tools (e.g., some tools have high counts due to checks for configuration problems).

?? CERT advisories addressed a small portion of known vulnerabilities, which is a reflection of the CERT's approach of releasing advisories only for the most serious and widespread problems.

Table 2-2 shows a sample mapping summary of tools, CERT advisories and a hacker site. The first column displays the name of the source mapped to CVE. The second and third columns display the number of entries in the data source that are mapped or not mapped to a CVE entry. The fourth column displays the number of unique CVE entries to which the data source mapped. The last column indicates the percentage of CVE covered by the data source.

Table 2-2. Sample of CVE Mapping Results

Name	# Mapped to CVE	# Not mapped to CVE	# Unique	CVE Coverage (%)
Tool 1	240	100	197	30
Tool 2	280	40	270	41
CERT	104	93	90	14
Hacker Site	47	26	43	7

These detailed comparisons were previously extremely difficult and time-consuming to perform.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

3 MOVING BEYOND THE DRAFT CVE

This section discusses expected activities in moving beyond the Draft CVE, including the motivation for the project, the CVE collaborative process, the pursuit of CVE content decisions, the role of CVE in the IDS community, and lessons learned.

3.1 Motivation

The motivation of this effort was to achieve the broadest perspective, highest validation, and widespread acceptance for CVE. To support that motivation, MITRE formed a CVE Editorial Board, consisting of members of academia, security tool vendors, security experts, and researchers, to review and validate the contents of CVE. A more detailed explanation of the processes involved in this review and validation is contained below. With members of the information security community involved in the validation, maintenance, and updating of CVE, and the added benefit of a sense of community ownership, CVE could quickly become a standard feature in tools and databases created by various vendors. By making CVE a publicly available, freely distributable document, and allowing public disclosure of the discussions and voting by the CVE Editorial Board, CVE will be a true public interest asset.

3.2 The CVE Collaborative Process

The CVE Editorial Board was established to validate the draft CVE and discuss issues related to its adoption, distribution, and maintenance. The first meeting of Editorial Board members was held on 9 May 1999 at the SANS-99 conference in Baltimore, Maryland, and the board meets periodically. This validation has produced a collaborative process for bringing new information into CVE.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

3.2.1 Editorial Board Members

At the time of this writing, the Editorial Board is comprised of the following representatives[11]:

David Baker (MITRE)
Andy Balinsky (Cisco Systems)
Matt Bishop (University of California-Davis)
Steven Christey (MITRE)
Russ Cooper (NTBugtraq)
Marc Dacier (IBM Zurich Research Laboratory)
Bill Fithen (CERT)
Andre Frech (ISS)
Bill Hill (MITRE)
Kent Landfield (Network Flight Recorder, Inc.)
Elias Levy (Bugtraq/Security Focus)
Dave Mann (MITRE)
Pascal Meunier (Purdue University CERIAS)
Stephen Northcutt (OSD/BMDO)
Craig Ozancin (AXENT)
Alan Paller (SANS)
Paul Proctor (CyberSafe)
Mike Prosser (L-3 Network Security)
Adam Shostack (Bindview)
Steve Snapp (CyberSafe)
Eugene H. Spafford (Purdue University CERIAS)
Stuart Staniford-Chen (Silicon Defense)
Bill Wall (STAT Ops/Harris Corporation)

3.2.2 Roles

The CVE Editorial Board has been validating the Draft CVE by a voting process. The Board primarily communicates via an electronic mailing list on which proposals, discussions, and votes are conducted. Some members of the Editorial Board vote and others perform an observer role. The Board is currently moderated by Steven Christey of MITRE, to provide a central channel for the Board's activities.

3.2.3 Phases of Vulnerability Consideration

The Board operates in five primary phases when considering entries for CVE: Assignment, Announcement, Interim Decision, Final Decision, and Publication.

3.2.3.1 Assignment phase

A Candidate Naming Authority (CNA) assigns a entry candidate a name equivalent to the next CVE number by replacing CVE with CAN (e.g., CAN-1999-0345), once the CNA is satisfied the entry is not yet in CVE and it appears to meet the requirements for entry. This numbering scheme allows a name to be assigned to a problem in the early stages of its discovery, but makes it clear the problem has not been validated as a entry by the Board.

---

[11] Listed in alphabetical order.

3.2.3.2 Proposal Phase

The CNA proposes the candidate to the Board. The proposal includes the following data: candidate name; identification of the CNA making the proposal; date the candidate name was assigned; date the candidate name was publicly announced; the category; the reference(s); and the proposed description. During this phase, the Board discusses, reviews, and votes upon the candidate.

Currently, due to the effort required to validate the numerous entries listed in the Draft CVE, voting on entry candidates is often done in clusters or groups of entries sharing some common characteristics. Each board member votes according to one of the following responses for each entry in a cluster:

- **Reviewing** - member is still reviewing/researching the candidate
- **Accept** - member accepts the entry as proposed
- **No Opinion** - member expresses no opinion and does not vote on a candidate (useful if member is not expert in that type of vulnerability)
- **Reject** - member rejects the candidate (e.g., not a vulnerability; unconfirmed; not a CVE entry; duplicate of an existing CVE entry; it subsumes another CVE entry; it is subsumed by another CVE entry)
- **Modify** - entry is generally acceptable, but some details are incorrect or missing (e.g., the description needs slight modification; the reference is incorrect)
- **Recast** - member does not believe the entry should be entered into CVE without being heavily modified (e.g., the entry should be merged with another entry; the entry should be split into multiple entries)

3.2.3.3 Interim Decision Phase

After a significant period without discussion or after receiving sufficient Accept votes, the moderator posts a decision based on the votes and/or discussion. Members have a few days to post objections. If significant discussion ensues, the entry remains at Interim Decision Phase.

3.2.3.4 Final Decision Phase

When discussion has terminated or the moderator believes making a decision is in the best interest of the community, the Final Decision is issued. The moderator makes a final decision and announces it to the Board.

3.2.3.5 Publication Phase

If accepted, a candidate is published (announced to the public). If rejected, the decision is recorded in the candidate database.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

3.2.4 Other Decision-Making Activities

Voting is also conducted on content decisions for CVE. The issues surrounding content decisions will be discussed in detail in the Section 3.3. The process is similar to that for individual candidates, but refers to the process for deciding how entries should be added to CVE. There are four primary phases to the content decision process.

3.2.4.1 Proposal Phase

The moderator names the content decision and proposes it to the Board, indicating candidates affected by the proposed decision. A time limit for discussion is proposed, which should be no less than a week. The Board discusses the content decision and each member votes according to one of the following categories:

- **Accept** - member accepts the content decision as stated
- **No Opinion** - member expresses no opinion and does not vote on a content decision
- **Modify** - member wishes to modify the content decision
- **Reject** - member rejects the content decision

For content decisions, a lack of a vote from a member is considered a No Opinion vote. There was discussion to interpret a lack of a vote to mean consent, as a lack of response could be interpreted as acceptance. However, such a decision can lead to flawed or erroneous assumptions, since many topics have two or more differing viewpoints and there would be no way to establish which view represented the opinions of non-voters.

3.2.4.2 Modification Phase

In some cases, the moderator may make changes to the content decision and re-propose them to the Board, based on Board feedback. The Board would then vote again, by a deadline as imposed by the Moderator.

3.2.4.3 Interim Decision Phase

Once the moderator decides there has been sufficient agreement and discussion to resolve a content decision, and no new issues related to that decision have been raised, the moderator makes an Interim Decision to Accept or Reject the content decision. In cases where most or all of the Board members vote No Opinion, the moderator has discretion to determine whether to Accept or Reject the decision. A deadline for final comments or discussion allows at least four days.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

3.2.4.4 Final Decision Phase

The moderator makes the Final Decision to ACCEPT or REJECT the content decision. Entry candidates affected by the content decision are re-evaluated, which may require a modification of some candidates.

3.2.5 Formalizing the Process

The CVE collaborative processes of the Board began somewhat informally. As CVE matures, the processes will adapt and formalize to the extent necessary to ensure the validity and reliability of CVE.

A part of that process is the maintenance of the publicly available Board discussions, which provides a documented basis for decisions made. To ensure openness in CVE, the mailing list will be publicly mirrored on a CVE Web site. Also available will be the entire CVE, references concerning vulnerability enumeration, a Frequently Asked Questions (FAQ) document, and a basic search capability of CVE.

3.2.6 Protecting CVE

CVE content will be copyrighted to protect it from unauthorized alteration, but CVE will be freely available for download to the public and freely distributable. Full public distribution is an essential requirement of an acceptable common enumeration. MITRE is committed to participate and keep CVE available in the public interest.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection,
7-9 September 1999

3.3 Pursuing CVE Content Decisions

3.3.1 Introduction

Probably the most important aspect of populating CVE is determining the basis for deciding when a entry candidate is a "CVE entry." This determination requires proposing a set of definitions of entries for the purpose of CVE. The next requirement is a means to evaluate candidate entries against some set of criteria to determine if the entry candidate meets the definition, if it is at the appropriate level of abstraction, and whether or not the entry is already listed in CVE. These criteria are referred to as content decisions.

The approach taken has been first to identify those vulnerabilities that most of the community would accept as such, although some may require modification of the description. Considering the initial Draft CVE contained over 650 entries, this identification process required sorting or organizing them into some form of groups or clusters by operating system and basic category of the vulnerability. Once the easy ones were identified, the harder ones that have complex or controversial content decision requirements were brought forward. As of this writing, the Editorial Board is currently engaged in review and validation of the more difficult entries identified in the Draft CVE. New content decisions are being proposed and evaluated. Much consideration is being given to the outcome of these decisions, as the future of CVE and its usefulness depend on the decisions made now.

Once the initial set of entries has been determined and the content decisions made, emerging vulnerabilities will be proposed as candidates and incorporated into CVE. Older vulnerabilities that are rarely encountered will then be incorporated to ensure the comprehensiveness of CVE.

Some of the most critical decisions facing the board are described below.

3.3.2 Defining Vulnerability Amid Multiple Perspectives

Defining the meaning of vulnerability is very difficult. Although there may be broad agreement that many specific conditions are in fact vulnerabilities, for many other conditions whether or not they are vulnerabilities is a matter of debate. Often, this evaluation depends upon perspective and security posture. Certainly, many organizations would not consider running finger, rexec, or HTTP services, to be a vulnerability. On the other hand, some organizations may consider the use of such services to be a severe violation of their security policy, not only because of the inherent vulnerabilities of some of the services, but simply because of the amount of information that could be gleaned by entities wishing to compromise their system. Use of these services amounts to an "exposure" by the organization. An exposure would be a state in a system

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

that is not a vulnerability in itself, but could lead to compromise through vulnerabilities accessible through that exposure.

Defining vulnerabilities for the purpose of CVE has prompted many discussions among Board members. The Board is currently working on the difficult process of defining the term "vulnerability," at least for CVE and well enough to make progress. The exposure terminology has helped the Board move in this direction.

3.3.3 Level of Abstraction

A key issue in deciding the content of CVE has been the determination of what level of abstraction the vulnerabilities should be defined. Some types of vulnerabilities can easily be described at one or two levels, while others are much more complex.

In some cases, the level of abstraction may be too high for some types of data sharing. Consider the case in which a tool discovers a configuration problem involving a default password. When presenting results to an end user (e.g., a system administrator who needs to fix the problem), the end user needs to know the precise account or password that needs to be changed. A tool may have multiple methods for exploiting the same kind of vulnerability; from the CVE perspective, they are all "instances" of the same CVE entry.

The level of abstraction may be too low for other applications. For example, from the perspective of educating programmers to avoid or correct common programming errors that result in vulnerabilities, the specific instances of "buffer overflows" may be at too low a level except to point to instances of the general problem.

In the cases where there is a mismatch between levels of abstraction, CVE can still serve as a facilitator for data sharing. For example, in the cases where the level of abstraction is too high, CVE could provide a base language that could then be extended to provide lower level descriptions. In cases where the level of abstraction may be too low, sublists of CVE entries could be used to "package" related entries under the higher level concept.

3.3.4 Configuration Problems

Vulnerabilities related to configuration problems are being openly discussed for the first time. Little research has been done in this area, and there is not yet a useful language or nomenclature for such research and discussion. One difficulty in discussion of configuration problems and the presence of potentially dangerous services is that they have a very high cardinality. For example, consider the number of well-known default accounts with no passwords (dozens if not hundreds), or worse, the number of well-known port numbers identified by the Internet Assigned Numbers Authority (thousands).

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

In general, many types of configuration problems have a high cardinality. High cardinality vulnerabilities are impractical or impossible to enumerate individually within CVE. Enumerating such vulnerabilities would significantly increase the cost of maintaining and searching CVE, and reduce the effectiveness of CVE as a means of bridging or comparing vulnerability databases.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

4 Role of CVE in the IDS Community

CVE could serve a critical role in the IDS community, as well as other security subcommunities. The simplicity of CVE and its uses for information correlation could make IDSs more interoperable, thereby enhancing the communication and security of organizations using IDSs.

4.1 Interoperability

An enterprise could utilize an array of CVE-compatible IDSs protecting its network and assets. The inclusion of CVE in these systems would allow them to serve the following functions:

- ?? Identify potential gaps in coverage
- ?? Select the best tools to provide coverage without reliance on a single vendor for a "suite" solution
- ?? Correlate assessment tool results to IDS coverage to eliminate false positives or reduce alarm severity for "covered" vulnerabilities
- ?? Share incident/attack data between tools, reducing "double counting" of vulnerabilities
- ?? Compare IDS performance to enterprise needs/priorities

4.2 Reporting Consistency

IDSs that were extended to generate incident information using CVE nomenclature would generate a standard description of attacks. This standard description could be rapidly disseminated to incident response organizations, such as CERT, the Federal Bureau of Investigation's InfraGard, or another trusted entity. Such standardization could allow more effective analysis for determining macro-level attack patterns.

4.3 IDS Comparisons

The CVE could be useful when conducting quantitative comparisons of various IDSs (e.g., how accurately they detect attacks, how quickly they process attack signatures, the percentage of known attacks each tool detects, how quickly vendors provide updates). CVE, or a subset of CVE, could provide a standard list against which all IDSs could be compared.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

4.4 Common Attack List

CVE names could be used as part of a regularly updated list of common attacks. This list would be similar to the WildList[12] used in the anti-virus community. An enterprise could use this list to focus security efforts on the most typical attacks, or allow the enterprise to give an IDS vendor a list of attacks to use as a focus for providing signature capabilities.

4.5 CIDF and IDWG Uses

CIDF and IDWG could benefit from CVE. CVE would allow these groups to deal with naming attacks that are vulnerability related. Ongoing discussions within both communities relate to developing a standard attack name or event name listing, with issues nearly identical to the issues addressed in the formation of CVE. The existence of CVE could save the time and effort of developing another list of names, provided CVE maintains a broad view that encompasses the needs of the CIDF and IDWG.

[12] WildList information at http://www.wildlist.org/

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

5 Lessons Learned

Lessons learned from the CVE project can be grouped into four categories discussed below: restrict the scope, use a formal decision-making process, encourage timely decisions, and promote community participation.

5.1 Restrict the Scope

One of the primary reasons for the continual progress of the CVE project has been the decision to restrict the scope of the effort. By excluding the use of any sort of categorization scheme or taxonomy and restricting the representation of a CVE entry to a name and a description, many potential issues raised by participants were avoided. One consideration in this decision is that the lack of any categorization mechanism and a robust schema drastically reduces the direct utility of the enumeration. To achieve agreement among participants, certain desired capabilities were not included, focusing on the primary goal of providing a mechanism to facilitate data sharing. For this reason, CVE is defined as a dictionary, not a database.

Another aspect of the restricted scope of CVE has been to focus only on a narrow aspect of information security concerns. Specifically, CVE enumerates vulnerabilities within computer systems. It does not enumerate threat agents or attack tools, such as exploit scripts. Neither does it enumerate security-related facts about the operation of a computer system, such as physical access controls and back-up procedures. CVE exists within a broader context and represents only one area among other critical security concerns. This restriction in scope has made the goals of CVE achievable, and has limited the scope of controversy.

5.2 Use a Formal Decision-making Process

Use of a formal decision-making process has helped the moderator determine the need for a period of time for discussion, followed by a voting time frame. Also, deliberations would benefit from a mechanism beyond a Yes or No vote, to allow for modifications or qualifications to be made. Additionally, the vote should be an interim decision, allowing a time frame for other members to determine whether the modifications or qualifications raise issues not yet addressed that may alter their vote.

5.3 Encouraging Timely Decisions

Timeliness of decision making is a critical lesson learned. A concern was raised that having a democratic voting process might delay processing of new entries. While timeliness is a major consideration, it should not be the overriding concern. Accuracy and completeness must be paramount. Additionally, most of the real use for CVE is involved in interoperability and

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

information sharing. As CVE moves from validation of a large initial startup mode, to a smaller, steady stream of new candidates, the amount of time required to validate entries will diminish. As new candidates are proposed and discussed, they will be assigned a candidate number so that discussions and information sharing on a new vulnerability can continue while the decision making is underway.

Also, where timeliness has been identified as an issue, the CVE Editorial Board has been responsive enough to accommodate this need. The democratic nature of CVE is preferable in an effort like this, where the real benefit is to gain broad consensus and acceptance. Should the process be an individual effort, without community participation or ownership, then there would be little probability of the CVE ever gaining the necessary acceptance to be useful.

5.4 Promote Community Participation

Community participation is essential to the success of the CVE project. Because of the nature of such an enumeration and the potential for wide usage, it is critical to seek and obtain support from as many groups as possible within the affected community. Commercial vendors, members of academia, research institutions, government agencies, independent experts, as well as the expected end users, must all be considered. For purposes of attempting tool integration alone, it is critical to seek membership from as many vendors with a significant market share as possible. To ensure that the effort remains true to its intent, maintains its integrity, and considers new research, it is crucial to seek support from academia, research institutions, operational network analysts, response teams, and other experts.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection, 7-9 September 1999

6 CONCLUSIONS

Currently, the CVE Editorial Board is validating the Draft CVE. CVE does not now include all known vulnerabilities, but this is not a real limitation of CVE. It merely reflects the incompleteness of the identification of vulnerabilities used in populating the Draft CVE and the time constraints involved in preparing CVE for entry into the public arena.

The Board has shown a capacity for productive discussion on a number of issues related to vulnerability identification and enumeration. The Board has come to accommodate a wide view of vulnerabilities and awareness of the multiple perspective issue has expanded. The Board has continually added new members, becoming more representative of the community. The decisions made thus far by the Board have verified the inclusive intent of CVE. As the Board continues to grow and encourage debate on the issues, the concepts necessary for future developments in vulnerability research will become clear.

MITRE will maintain, update, and publish CVE on a public Web site for the foreseeable future. CVE will be as reflective of community opinion as possible through open discussions among members of the Board. Other information security experts will be invited to participate on the Board on an as-needed basis through recommendations from the Board members. MITRE will continue to work with intrusion detection standardization groups such as the CIDF and IDWG on matters pertaining to tool evaluations and interoperability standards.

Presented at the Second International Workshop on Recent Advances in Intrusion Detection,
7-9 September 1999

Author Biographies

Steven M. Christey has been at The MITRE Corporation since 1989, initially conducting research in artificial intelligence (AI) and moving into the information security arena in 1993. He has been MITRE's primary network security auditor for the last five years; consequently, his operational expertise is in network-based risk assessment and risk management. In the last two years he has conducted research which blends his experience in AI and security, in topics such as automated vulnerability analysis of source code and distributed security assessment. Mr. Christey holds a B.S. in Computer Science from Hobart College.

William H. Hill is a Principal INFOSEC Engineer in the Security and Information Operations Division at The MITRE Corporation. He has been involved in computer networking and information security since 1990, working in network programming, design, operations, and security, and most recently in vulnerability analysis, intrusion detection and response, and incident investigation. Mr. Hill coordinated a Department of Defense multi-agency product evaluation of government and commercially developed intrusion detection systems (IDSs), and has worked extensively with the U.S. Army Land Information Warfare Activity to develop and deploy a large IDS for the Army. Mr. Hill has been working with researchers funded by the Defense Advanced Research Projects Agency on the Common Intrusion Detection Framework, a set of protocols to facilitate communication and analysis between IDSs. Prior to joining MITRE, Mr. Hill worked for Bell Atlantic, managing network operations for its Internet Center. Mr. Hill holds a B.S. from Florida State University, and a Master of Science in Computer Science from George Mason University.

Dr. David E. Mann is a Senior INFOSEC Engineer/Scientist in the Security and Information Operations Division at The MITRE Corporation. Since joining MITRE in 1997, he has been involved in the development of various vulnerability databases, the security analysis of a routing protocol, and the development of a security policy management tool for IPsec. Prior to joining MITRE, he worked for the Quinsoft Corporation developing database application development software and the Naval Postgraduate School as a Visiting Professor of Mathematics. Dr. Mann holds a B.S. in Mathematics from Eastern Nazarene College, an M.S. in Mathematics from the University of Vermont, and a Ph.D. in Mathematics from Northeastern University.

Presented at the Second International Workshop on the Recent Advances in Intrusion Detection, 7-9 September 1999

David W. Baker is an INFOSEC Engineer in the Security and Information Operations Division at The MITRE Corporation. He joined MITRE in 1998, and has worked extensively with the U.S. Army Land Information Warfare Activity in the deployment and operation of a large scale intrusion detection system for the Army. He has also been active in other projects involving threat and vulnerability analyses for critical infrastructures. Prior to joining MITRE, he worked as a Special Agent for the U.S. Army Criminal Investigation Command, last serving as the command's principal forensic science advisor and a member of the Department of Defense Forensic Science Advisory Committee. He has worked with information security and operational security concerns since 1990. Mr. Baker holds a B.S. from The State University of New York, and a Master of Forensic Science degree from The George Washington University.

Presented at the Second International Workshop on the Recent Advances in Intrusion Detection, 7-9 September 1999