NATIONAL INFRASTRUCTURE ADVISORY COUNCIL

COMMON VULNERABILITY SCORING SYSTEM

FINAL REPORT AND RECOMMENDATIONS BY THE COUNCIL

October 12, 2004

JOHN T. CHAMBERS
WORKING GROUP CO-CHAIR
CHAIRMAN AND CHIEF EXECUTIVE OFFICER
CISCO SYSTEMS, INC.

AND

JOHN W. THOMPSON
WORKING GROUP CHAIR
CHAIRMAN AND CHIEF EXECUTIVE OFFICER
SYMANTEC CORPORATION

The Common Vulnerability Scoring System

National Infrastructure Advisory Council
Vulnerability Disclosure Working Group
Vulnerability Scoring Subgroup

Abstract

Over the past several years, a number of large computer security vendors and not-for-profit organizations have developed, promoted, and implemented procedures to rank information system vulnerabilities. Unfortunately, there has been no cohesion or interoperability among these systems. Also, existing systems tend to be limited in scope as to what they cover. Finally, all of these systems tend to be Internet-centric; that is, they tend to be concerned only with vulnerabilities affecting computers connected to the worldwide Internet. The NIAC commissioned this project to propose an open and universal vulnerability scoring system to address and solve these shortcomings, with the ultimate goal of promoting a common understanding of vulnerabilities and their impact.

Participants

Working Group Members

John Chambers, Working Group Co-Chairman
John Thompson, Working Group Co-Chairman
Thomas Noonan
Maynard Webb

Study Group Members

Sharad Ahlawat, Cisco Systems
David Ahmad, Symantec Corporation
Peter Allor, Internet Security Systems, Inc.
Scott Culp, Microsoft
Jim Duncan, Cisco Systems
Gerhard Eschelbeck, Qualys, Inc.
Jack Faherty, MITRE Corporation
Shawn Hernan, CERT/CC, Carnegie Mellon University
Art Manion, CERT/CC, Carnegie Mellon University
David McKinney, Symantec Corporation
Sasha Romanosky, eBay
Howard Schmidt, eBay
Mike Schiffman, Cisco Systems
Art Wong, Symantec Corporation
Ken Watson, Cisco Systems
Andrew Wright, Cisco Systems

NIAC Staff Reviewers

Scott Blanchette, Stanford Hospital & Clinics

Doug Busch, Intel Corporation
Andy Ellis, Akamai Technologies
Rick Holmes, Union Pacific Corporation
Dale Hulvey, James Madison University
Bruce Larson, American Water
Bill Muston, TXU Corp
Eric Meyers, DuPont Company
Stan Scott, TXU Corp
Mark Stein, TXU Corp
Susan Vismor, Mellon Financial Corporation

Table of Contents

I. Introduction and Background

In July 2003, the National Infrastructure Advisory Council (NIAC) commissioned a research project to promote a common understanding of vulnerabilities and their impact through development of a common vulnerability scoring system. Current scoring systems, in use by the Computer Emergency Response Team/Coordination Center (CERT/CC), Symantec, Internet Security Systems, Cisco Systems, and others, rate vulnerabilities according to a variety of metrics and determine a single overall threat score by weighing these metrics. These systems use different, non-common metrics to characterize vulnerabilities; they are Internet-centric; they do not universally accommodate changes over time; and they do not have provisions for user operational environments with different risk profiles. The Common Vulnerability Scoring System (CVSS) is an open, comprehensive system that addresses these shortcomings and is suitable for international adoption across all types of information systems and environments.

When designing the system, researchers kept the following goals in mind:

- **Open:** Must be freely available, adoptable, and open to use by anyone. A closed standard will not be widely implemented and will not survive;
- **Comprehensive:** Must be able to describe any possible vulnerability in any type of information system;
- **Interoperable:** Should work well with existing technology and infrastructures, and not rely on proprietary technologies or formats;
- **Flexible:** Should be customized to operational environments with different risk profiles;
- **Simple:** Must be simple and straightforward to understand, implement, and use.

Vulnerability Definition and Classification

To promote a common understanding of the severity of vulnerabilities, a common vulnerability definition must be used. For CVSS, the Study Group used the definition as stated in the NIAC Vulnerability Disclosure Framework report, which was completed January 2004.

A vulnerability is defined as a set of conditions that may lead to an implicit or explicit failure of the confidentiality, integrity, or availability of an information system. Examples of the unauthorized or unexpected effects of a vulnerability may include any of the following:

- Executing commands as another user;
- Accessing data in excess of specified or expected permission;
- Posing as another user or service within a system;
- Causing an abnormal denial of service;
- Inadvertently or intentionally destroying data without permission; or
- Exploiting an encryption implementation weakness that significantly reduces the time or computation required to recover the plaintext from an encrypted message.

Common causes of vulnerabilities are design flaws in software and hardware, botched administrative processes, lack of awareness and education in information security, technological advancements or improvements to current practices, any of which may result in real threats to mission-critical information systems. For more information, see [1].

II. CVSS Overview and Scope

The Common Vulnerability Scoring System (CVSS) is designed to provide the end user with a composite score representing the overall severity and risk a vulnerability represents. Using CVSS, security professionals, executives, and end-users will have the basis for a common language with which to discuss vulnerability severity. CVSS, as shown in Figure 1, is a modular system with three distinct groups that combine the intrinsic characteristics of a vulnerability. Each of these qualities or "metrics" has a specific way of being measured and each group has a unique formula for combining and weighing each metric. CVSS also presents a very simple interface to end-users because all numeric computation is automatically performed in the background.

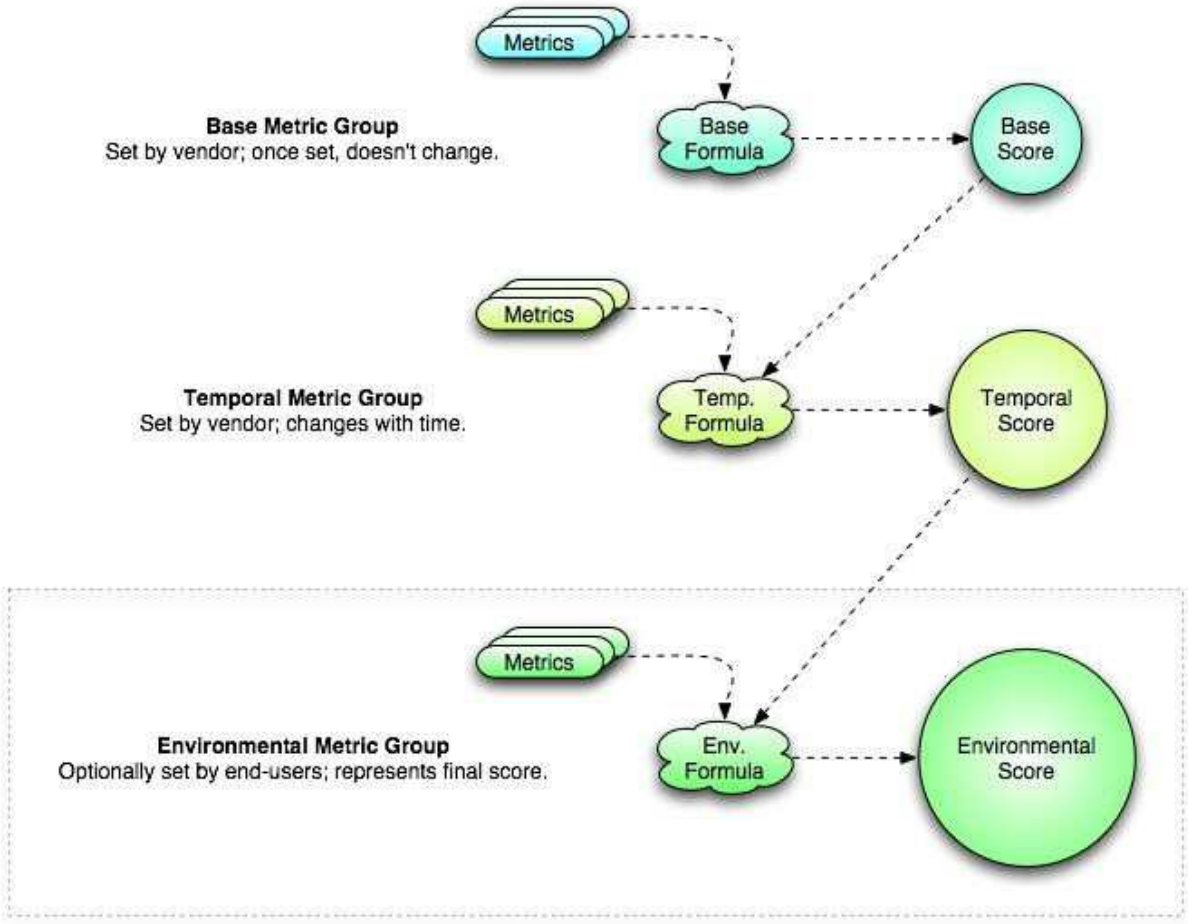

Figure 1: The CVSS Model

The scope of CVSS is constrained to contain only what is necessary to adequately describe information system vulnerabilities. CVSS does not address any external issues such as the following:

- **Potential Threat:** These are vulnerabilities that are dependent on the exploitation of other vulnerabilities before they become a risk. For example, consider a vulnerable application that resides behind a firewall, which itself is vulnerable. When scoring the

firewall vulnerability, consideration should not be made for the effects of the exploitation of the application or any other secondary vulnerabilities.

- **Real-time Attack Scoring:** CVSS does not have any capacity for tracking the threats posed by the ongoing exploitation of vulnerabilities. This is a consideration that would be tracked by other scoring systems.
- **Global Exposure Scoring:** CVSS itself is not a database and has no intrinsic capacity for collating and averaging large amounts of vulnerability scores.

Vulnerability Metrics

A metric, in CVSS parlance, is a constituent component or characteristic of a vulnerability that can be quantitatively or qualitatively measured. These values are clustered together in three groups: base, temporal, and environmental. The base group contains all of the qualities that are intrinsic to any given vulnerability that does not change over time or in different environments. The temporal group contains the characteristics of a vulnerability that are time-dependent and change as the vulnerability ages. Finally, the environmental group contains the characteristics of vulnerabilities that are tied to implementation and are specific to a user's environment. The final adjusted score represents the risk a vulnerability represents at a particular point in time for a specific environmental condition.

III. Base Metrics

Once discovered, analyzed, and catalogued, there are certain aspects of a vulnerability that do not change, assuming the initial information is complete and correct. These immutable characteristics will not change over time, nor in different environments. The base metric group captures the "access" and "impact" qualities.

The base metric group first addresses the means by which a vulnerable system may be reached. Three "access" qualities are assessed: access vector, access complexity, and authentication. Combined, these three factors describe how a target may be accessed in order to exploit the vulnerability and also whether or not there are mitigating factors that complicate the process.

In addition to "access" qualities, the base metric group also measures impact of a vulnerability on an information system. The three most widely accepted security properties for information systems are confidentiality, integrity and availability. The impact of a vulnerability on affected systems can be defined as a combination of losses to varying degrees of each of these properties. Vulnerability impact needs to be expressed in terms of the confidentiality, integrity, and availability properties: from negligible to total losses for each of the three properties as well as combinations of losses. For example, the possibility for partial loss of integrity and the partial loss of confidentiality due to a vulnerability in a logging mechanism.

Access Vector

This metric specifies whether or the vulnerability is exploitable locally or remotely.

A vulnerability exploitable with only local access ("locally exploitable") means the attacker must have either physical or authenticated login access to the target system, often either a walk-in scenario or a local account on a target computer system. Through remote access technologies, it is possible for users to log into systems from remote locations. Vulnerabilities present on the system that are exploitable by these users are still considered local, as the vulnerable components themselves are not directly exposed to any network interfaces. For example, a vulnerability in

the password management utility "passwd" is a locally exploitable vulnerability. A vulnerability in an Secure Shell (SSH) server that is exploitable by remote clients, with or without authentication, is considered remotely exploitable.

Remotely exploitable vulnerabilities are those that affect technologies which accept input from other hosts, reading directly from a network interface. If the vulnerability can be triggered by a remote host, it is considered remotely exploitable. These technologies may or may not require remote parties to authenticate with them (see authentication). Remotely exploitable vulnerabilities are also implicitly local as it is assumed that an attacker who has local access can also connect to a vulnerable technology on that system using network components.

A vulnerability that is remotely exploitable is considered to be a higher risk threat than one that is only locally exploitable, since the pool of potential attackers is greater. Therefore, if a vulnerability is only locally exploitable, its resulting CVSS score will be lower than if it were remotely exploitable.

Scoring Evaluation
Guidelines for scoring the access vector metric are as follows:
Local: The vulnerability is only exploitable locally (i.e., it requires physical access or interactive access to the target system).
Remote: The vulnerability is exploitable remotely.

Access Complexity

This metric specifies the complexity of attack required to exploit the vulnerability once an attacker has accessed the target system. In most cases, once the target system is contacted, the vulnerability can be directly exploited. The traditional example is a simple remotely exploitable buffer overflow in an Internet server program that continuously runs. Once the target system is located, there is no additional complexity in accessing the target; an attacker presumably can exploit the target at will. Other vulnerabilities require specialized access considerations or additional barriers in order to become exploitable. An example in this case would be a vulnerability in an e-mail program that is only exploitable when the user downloads and opens a tainted attachment. The additional complexity is the user-interaction required to complete a successful attack. This metric attempts to represent any potential for additional complexity in these cases.

Scoring Evaluation
Guidelines for scoring the access complexity metric are as follows:
High: Specialized access conditions exist. For example, the system is exploitable during specific windows of time (a race condition), the system is exploitable under specific circumstances (non-default configurations), or the system is exploitable with victim interaction (vulnerability exploitable only if user opens e-mail).
Low: Specialized access conditions or extenuating circumstances do not exist; the system is always exploitable.

Authentication

This metric specifies whether or not an attacker needs to be authenticated to the target system in order to exploit the vulnerability. The specific type and mechanism for authentication is not important because authentication in any form will add significant complexity to the exploit.

Additionally, this metric is an either-or consideration. Attackers without valid credentials should not be able to access the target in order to exploit the vulnerability. If any sort of authentication "Required", the final CVSS score will be considerably lower than if it were "Not Required".

It is important to note that the authentication metric is distinct from the access vector metric. The requirement for authentication represented by this metric is considered **once the system has already been accessed**. Specifically, in the case of locally exploitable vulnerabilities, this field should only be set to "Required" if authentication is needed beyond what is necessary for a user to login to the system (and thus becoming "local"). The access vector metric (local or remote) reduces the score if the vulnerability is flagged as locally exploitable, thus taking into consideration the prerequisite authentication. One example of a locally exploitable vulnerability that requires authentication is one that affects a database that is only accessible once the user has logged into the system. If the user must then login (authenticate) as a valid database user in order to exploit the vulnerability then this metric should be set to "Required".

Scoring Evaluation
Guidelines for scoring the authentication metric are as follows:
Required: Authentication is required to access and exploit the vulnerability
Not Required: Authentication is not required to access or exploit the vulnerability

Confidentiality Impact

Confidentiality refers to limiting information access and disclosure to only authorized users, as well as preventing access by or disclosure to unauthorized ones. Confidentiality is usually preserved by a system's information protection mechanisms: cryptography, data compartmentalization, identification and authentication systems, etc. Compromise of a system's information protection mechanism can negatively impact confidentiality. This metric measures the impact on confidentiality of a successful exploit of the vulnerability on the target system.

Scoring Evaluation
Guidelines for scoring the confidentiality impact metric are as follows:
None: No impact on confidentiality.
Partial: Considerable informational disclosure. Access to critical system files is possible. There is a loss of important information, but the attacker doesn't have control over what is obtainable or the scope of the loss is constrained. For example, "partial" would indicate a vulnerability that divulges bits of an encryption key or password hash information.
Complete: A complete loss of system protection resulting in all information being revealed. The attacker has sovereign control to read all of the system's data (memory, files, etc).

Integrity Impact

Integrity refers to the trustworthiness and guaranteed veracity of information. Integrity controls are meant to protect data from unauthorized modification. When the integrity of a system is sound, it is fully proof from unauthorized modification. This metric measures the impact on integrity of a successful exploit of the vulnerability on the target system.

Scoring Evaluation
Guidelines for scoring the integrity impact metric are as follows:
None: No impact on integrity.

Partial: Considerable breach in integrity. Modification of system files or information is possible, but the attacker does not have control over what can be modified, or the scope of what the attacker can affect is constrained. For example, key system or program files may be overwritten or modified, but at random or in a limited context or scope.
Complete: A total compromise of system integrity. There is a complete loss of system protection resulting in the entire system being compromised. The attacker has sovereign control to modify any system files.

Availability Impact

Availability refers to the accessibility of information resources. Almost exclusive to this domain are denial-of-service vulnerabilities. Attacks that compromise network bandwidth, processor cycles, disk space, or administrator time impact the availability of a system. This metric measures the impact on availability of a successful exploit of the vulnerability on the target system.

Scoring Evaluation
Guidelines for scoring the availability impact metric are as follows:
None: No impact on availability.
Partial: Considerable lag or interruptions in resource availability. For example, a network-based flood attack that reduces available bandwidth to a web server farm to such an extent that only a small number of connections successfully complete.
Complete: Total shutdown of the affected resource. The attacker can render the resource completely unavailable.

Impact Bias

An important consideration of the three impact metrics (confidentiality, integrity, and availability) is that the importance of the individual properties can vary among systems. For example, a vulnerability affecting the confidentiality of an encrypted file system is far more severe than one affecting its availability. To these ends, impact bias allows a score to convey greater weight to one of the three impact. The impact bias metric will have no effect if the three impact metrics are all assigned the same value.

Scoring Evaluation
Guidelines for the impact bias metric are as follows:
Normal: Confidentiality impact, integrity impact, and availability impact are all assigned the same weight.
Confidentiality: Confidentiality impact is assigned greater weight than integrity or availability impact.
Integrity: Integrity impact is assigned greater weight than confidentiality or availability impact.
Availability: Availability impact is assigned greater weight than confidentiality or integrity impact.

IV. Temporal Metrics

As a vulnerability ages, certain intrinsic characteristics will change with time. In many cases, when first discovered, a set of vulnerable systems will be at or close to its peak, while the availability of exploit and remedial information will be at its lowest point. As time progresses, patch information will become more available and more systems will be fixed as more exploits

occur, driving the need for the fix. Eventually, the set of vulnerable systems will reach its low point as remedial information reaches its high point. The CVSS temporal metrics group captures these characteristics of a vulnerability that change over time.

Exploitability

This metric measures how complex the exploitation process is once it has been accessed. As time progresses, exploit code may become available, when there previously was none. Additionally, existing exploit code may improve and work better or to be easier to implement, and in severe cases, it may be delivered as the payload of an Internet-based worm or virus.

Scoring Evaluation

Guidelines for scoring the exploit complexity metric are as follows:
Unproven: No exploit code is yet available.
Proof of Concept: Proof of concept exploit code is available. The code is not functional in all situations and may require hand tuning in order to get it to work in any situation.
Functional: Functional exploit code is available. The code works in most situations where the vulnerability is exploitable.
High: Exploitable by functional mobile autonomous code. The code works in every situation where the vulnerability is exploitable and is actively being delivered via a mobile autonomous agent (a worm or virus).

Remediation Level

This metric measures the level of solution available.

Scoring Evaluation

Guidelines for scoring the remediation complexity metric are as follows:
Official Fix: Complete vendor solution available.
Temporary Fix: There is an official, temporary fix available.
Workaround: There is an unofficial non-vendor solution available.
Unavailable: There is either no solution available or it is impossible to apply.

Report Confidence

This metric measures the degree of confidence in the existence of the vulnerability and the credibility of its report.

Scoring Evaluation

Guidelines for scoring the report confidence metric are as follows:
Unconfirmed: A single unconfirmed source or possibly several conflicting reports. There is little confidence in the validity of the report. For example, a rumor that surfaces from the hacker underground.
Uncorroborated: Multiple non-official sources; possibly including independent security companies or research organizations.
Confirmed: Vendor has reported/confirmed a problem with its own product.

V. Environmental Metrics

Different user environments can have an immense bearing on how (or if) a vulnerability affects a given information system and its stakeholders. The CVSS environmental metrics group captures characteristics of vulnerabilities that are tied to system distribution and network environment.

Collateral Damage Potential

This metric measures the potential for a loss in physical equipment and/or property damage. Collateral damage can include financial damage, physical damage, personnel casualties, or significant damage to reputation. This metric is intentionally broad to allow for individual user application to circumstances specific to each environment.

Scoring Evaluation

Guidelines for scoring the physical loss potential metric are as follows:
None: There is no potential for property damage.
Low: A successful exploit of this vulnerability may result in light property damage or loss. The system itself may be damaged or destroyed.
Medium: A successful exploit of this vulnerability may result in significant property damage or loss.
High: A successful exploit of this vulnerability may result in catastrophic property damage and loss. The range of effect may be over a wide area.

Target Distribution

This metric measures the number of target systems susceptible to the vulnerability. It is meant as an environment-specific indicator in order to approximate the percentage of systems within the environment that could be affected by the vulnerability.

Scoring Evaluation

Guidelines for scoring the target distribution metric are as follows:
None: No target systems exist, or targets are so highly specialized that they only exist in a laboratory setting. As best as can be determined, no systems currently deployed within the environment depend on target systems for business operations. Effectively 0% of the environment is considered at risk.
Low: Targets exist inside the environment, but on a small scale. Between 1% – 15% of the total environment is considered at risk.
Medium: Targets exist inside the environment, but on a medium scale. Between 16% – 49% of the total environment is considered at risk.
High: Targets exist inside the environment on a considerable scale. Between 50% – 100% of the total environment is considered at risk.

VI. Scoring

Scoring is process of combining the values of all metrics from each group into a final composite score that represents the overall risk of a given vulnerability. The CVSS scoring process is broken into three phases, one for each metric group. Scoring begins with the base metric group and then temporal and environmental metrics are applied to produce a final score. Each group has a different formula that combines its constituent metrics. The formulas operate behind the scenes of the CVSS tool, transparent to the user. See Appendix B for details on all metrics, weights, and formulas. The base metric group captures the fundamental constituent qualities of a

given vulnerability and therefore provides the foundation for the final score. The temporal and environmental metric groups apply downwards and upwards scoring modifiers to the base score.

Base Metric Scoring

The base score provides the foundation for the overall vulnerability score. The most significant metrics in the scoring process are the three impact metrics. These metrics dictate the overall effect the vulnerability will have on target systems and therefore have the strongest bearing on the final score.

Temporal Metric Scoring

The temporal score adjusts the base score by including factors that may change over time. The temporal score will be less than or equal to the base score; that is, the temporal metrics serve only to reduce the base score by a maximum of 25%. See Figure 2.

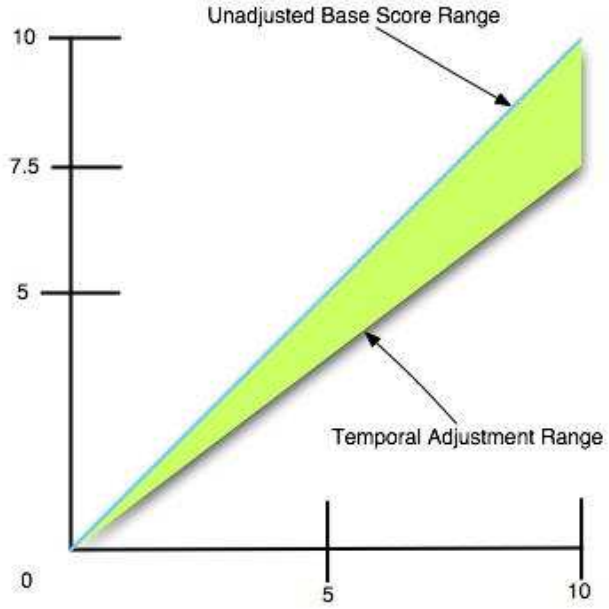

Figure 2. Temporal Adjustment Range to CVSS Base Score

Environmental Metric Scoring

The environmental score adjusts the temporal score to account for aspects of an organization's environment. The environmental score may be higher or lower than the temporal score as shown Figure 3.

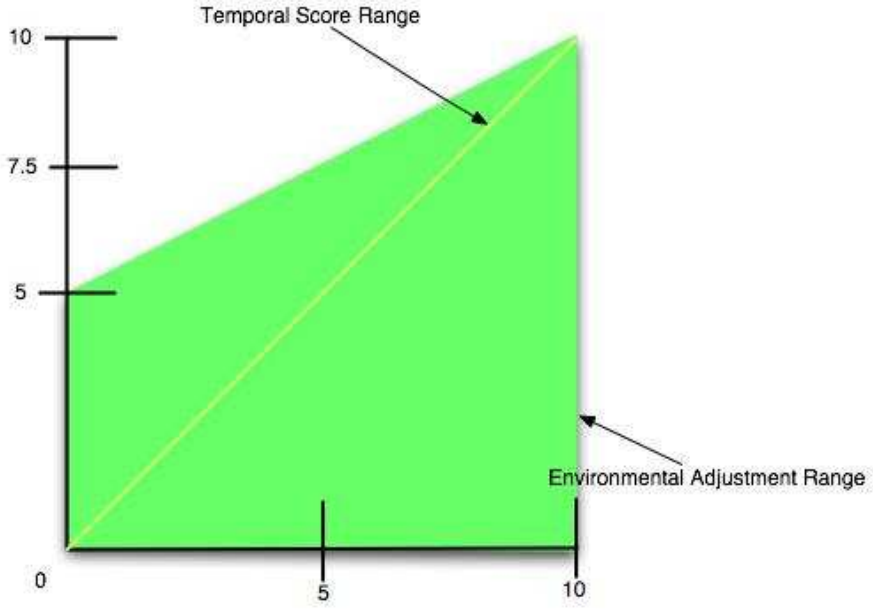

Figure 3. Environmental Adjustment Range to Temporal-Adjusted Base Score

VII. Recommendations for the President

- Support use of the CVSS by all Federal departments and agencies. Those departments and agencies involved in identifying, reporting, and scoring vulnerabilities should develop Base and Temporal scores to contribute to the worldwide body of knowledge for each vulnerability. All departments and agencies should compute Environmental (i.e., Final) scores as they become involved in remediating and resolving vulnerabilities.

- Encourage DHS to promote the use of CVSS by the global vulnerability management community, including international, state, local, and tribal governments, critical infrastructure owners and operators, and discoverers, vendors, coordinators, and users.

- Coordinate with the NIAC to identify an organization to function as the permanent home for CVSS. This organization should be responsible for maintaining and updating CVSS metrics and formulas. It should also possess significant technical expertise and experience managing vulnerabilities, and should maintain a global focus.

Appendix A: Nomenclature

Technical terms can be subjective in connotation. In order to effectively communicate, the Working Group recommends the following common definitions:

- **Vulnerability:** A set of conditions that may lead to an implicit or explicit failure of the confidentiality, integrity, or availability of an information system.
- **Vulnerability Scoring:** The process of ranking vulnerabilities in computer and information systems in order to allow parties interested in vulnerability information to easily identify, prioritize, track, and address those vulnerabilities most significant to their operations.
- **Metric:** A constituent component or characteristic of a vulnerability that can be quantitatively or qualitatively measured.
- **Vulnerability Scoring System (VSS):** The entire framework for ranking vulnerabilities, including usage guidelines and metrics.
- **Vulnerability Score (or Score):** A composite score computed for each vulnerability made up from ranking individual metrics.

Appendix B: Formula Reference

The following pseudo-code details each of the formulas included in the CVSS.

```
BaseScore =
   round to 1 digit of
     10
     * (case AccessVector of
          local:                0.7
          remote:               1.0)
     * (case AccessComplexity of
          high:                      0.8
          low:                  1.0)
     * (case Authentication of
          required:             0.6
          not-required:         1.0)
     * (   (case Confidentiality Impact of
               none:            0
               partial:         0.7
               complete:        1.0)
          * (case Impact Bias of
               normal:          0.333
               confidentiality:       0.5
               integrity:       0.25
               availability:    0.25)
        +  (case Integrity Impact of
               none:            0
               partial:         0.7
               complete:        1.0)
          * (case Impact Bias of
               normal:          0.333
               confidentiality:       0.25
               integrity:       0.5
               availability:    0.25)
        +  (case Availability Impact of
               none:            0
               partial:         0.7
               complete:        1.0)
          * (case Impact Bias of
               normal:          0.333
               confidentiality:       0.25
               integrity:       0.25
               availability:    0.5))

Temporal Score =
   round to 1 digit of
     Base Score
     * (case Exploitability of
```

unproven: 0.85
proof-of-concept: 0.9
functional: 0.95
high: 1.00)
* (case Remediation Level of
official-fix: 0.87
temporary-fix: 0.90
workaround: 0.95
unavailable: 1.00)
* (case Report Confidence of
unconfirmed: 0.90
uncorroborated: 0.95
confirmed: 1.00)

Environmental Score =
round to 1 digit of
(Temporal Score +
(10 – Temporal Score)
* (case Collateral Damage Potential of
none: 0
low: 0.1
medium: 0.3
high: 0.5))
* (case Target Distribution of
none: 0
low: 0.25
medium: 0.75
high: 1.00)

Appendix C: Example Vulnerability Scoring Worksheet

Common Vulnerability Scoring System (CVSS) Version 0.2

Vulnerability	Microsoft Outlook Express Scripting vulnerability	Microsoft LSASS vulnerability	BGP potential DOS
CVE number	CAN-2004-0 80	CAN-2004-0	CAN-2004-0
URL			

Base Metrics	Access Vector	REMOTE	REMOTE	REMOTE
	Access Complexity	HIGH	LOW	HIGH
	Authentication	NOT-REQUIRED	NOT-REQUIRED	NOT-REQUIRED
	Confidentiality Impact	COMPLETE	COMPLETE	NONE
	Integrity Impact	COMPLETE	COMPLETE	NONE
	Availability Impact	COMPLETE	COMPLETE	COMPLETE
	Impact Bias	NORMAL	NORMAL	AVAILABILITY
BASE SCORE		**8.0**	**10.0**	**4.0**
Temporal Metrics	Exploitability	FUNCTIONAL	FUNCTIONAL	UNPROVEN
	Remediation Level	OFFICIAL-FIX	OFFICIAL-FIX	UNAVAILABLE
	Report Confidence	CONFIRMED	CONFIRMED	CONFIRMED
TEMPORAL SCORE		**6.6**	**8.3**	**3.4**

	Collateral Damage Potential	NONE	NONE	NONE
	Target Distribution	HIGH	HIGH	HIGH
ENVIRONMENTAL SCORE		**6.6**	**8.3**	**3.4**

Appendix D: References

[1] “Vulnerability Disclosure Framework”, National Infrastructure Advisory Council, http://www.dhs.gov/dhspublic/interweb/assetlibrary/vdwgreport.pdf