A MODEL FOR HOW TO DISCLOSE PHYSICAL SECURITY VULNERABILIES[*]

Roger G. Johnston, Ph.D., CPP
Vulnerability Assessment Team
Argonne National Laboratory

ABSTRACT

When security vulnerabilities are discovered, it is often unclear how much public disclosure of the vulnerabilities is prudent. This is especially true for physical security vis a vis cyber security. We never want to help the "bad guys" more than the "good guys", but if the good guys aren't made aware of the problems, they are unlikely to fix them. This paper presents a unique semi-quantitative tool, called the "Vulnerability Disclosure Index" (VDI), to help determine how much disclosure of vulnerabilities is warranted and in what forum. The VDI certainly does not represent the final, definitive answer to this complex issue. It does, however, provide a starting point for thinking about some of the factors that must go into making such a decision. Moreover, anyone using the VDI tool can at least claim to have shown some degree of responsibility in contemplating disclosure issues.

INTRODUCTION

Vulnerability Assessors and others who discover vulnerabilities in physical security devices, systems, measures, or programs often face difficult decisions about whom to warn, when, and in how much detail. When a formal vulnerability assessment (VA) has been chartered, the sponsor of the VA often owns the findings. Proprietary ownership of a VA study, however, doesn't automatically end the matter, it just brings additional people into the conundrum. Furthermore, it doesn't even necessarily relieve the vulnerability assessors of their responsibility to society to warn of clear and present danger.

When a particular vulnerability is unique and isolated within a single, small organization, a public disclosure is probably unwise. Many security vulnerabilities, however, are very extensive and global. The Vulnerability Assessment Team[1] (VAT) at Argonne National Laboratory, for example, has discovered fundamental vulnerabilities in a number of different physical security devices, systems, measures, and programs that could potentially have wide ranging implications for many individuals and organizations. The VAT has demonstrated serious vulnerabilities (as well as potential countermeasures) associated with the use of tamper-indicating seals[2,3,4], radio frequency identification tags (RFIDs) and contact memory buttons[3], Global Positioning System (GPS) receivers[3,5,6], nuclear safeguards[7,8,9], and techniques for vulnerability assessments[10]. It has often been unclear who should be warned of these vulnerabilities and in what detail, even given existing government rules, regulations, classification guidelines, and policies for dealing with sensitive information.

In the world of computer software, security vulnerabilities can typically be dealt with in a more straightforward manner. When a new cyber vulnerability is discovered, it is widely considered best

[*] Editor's Note: This paper was not peer reviewed.

practice to keep the vulnerability quiet until the software developer or computer manufacturer can be (quickly) contacted, and allowed time to fix the problem.[11,12,13,14] The software upgrade that results can then be rapidly and easily disseminated via the Internet to customers. Indeed, computer and network users know they should frequently (or even automatically) check for software patches and upgrades.

With physical security hardware or procedures in contrast, there is usually no equivalent simple, inexpensive way to provide updates and security fixes, nor even to contact customers. Many physical security devices and systems are sold through a complex network of dealers, vendors, and integrators. The purchaser may even be several layers removed from the end-user. And unlike software fixes, security upgrades to physical security devices, systems, measures, and programs often take a long time to develop and install, and can be quite expensive. Meanwhile, physical security may be at great risk.

Another complicating factor for physical security is that vague, generalized warnings about security vulnerabilities rarely result in countermeasures being implemented. Security managers and security programs tend to be inherently cautious and traditionalist, and are often severely restricted in terms of budget. Typically, attacks must be thoroughly described or demonstrated in detail, along with possible countermeasures, before either the vulnerability will be acknowledged, or any security improvements will be seriously considered. Unfortunately, implementing a countermeasure is often viewed by bureaucratic organizations as an admission of past negligence on the part of security managers, so security managers are often—understandably—less than eager to make changes[11,15,16,17]

With any detailed disclosure of vulnerabilities, we must worry about helping the “bad guys” (nefarious adversaries) more than the “good guys” (security providers). This is especially a concern if—as often happens—security managers or programs ultimately fail to implement recommended security countermeasures. Common reasons for this include a lack of funding, commitment, follow-through, or support from superiors, or an unwillingness to be proactive about security or to admit that security vulnerabilities exist. Sometimes the only way that necessary security countermeasure will be implemented (particularly within government organizations) is if there is public pressure to improve security. But detailed, public discussion of security problems is often a prerequisite for this kind of public awareness and pressure.

The purpose of this paper is to provide a tool to help decide if and how security vulnerabilities should be disclosed. This tool, called the Vulnerability Disclosure Index (VDI), is not presented here as the ultimate, authoritative method for dealing with this complex issue. It is offered instead as a first step, and as a vehicle for thinking about and discussing some of the factors that need to be pondered when vulnerability disclosures are being considered.

The VDI tool is a semi-quantitative method. A high VDI score suggests that public or semi-public disclosure of the vulnerability in at least some detail may well be warranted. A medium score supports the idea that it would be appropriate to discuss the vulnerability, but perhaps in lesser detail and/or to a more limited audience of security professionals and end-users. A low VDI score indicates the vulnerability should probably be kept in confidence, or shared discretely only with those having an explicit and immediate need to know.

THE VDI TOOL

The Vulnerability Disclosure Tool (VDI) works by considering 18 different factors (A-R), and subjectively scoring each for the vulnerability in question. The higher the score for each factor, the greater that factor supports full public, detailed disclosure.

The tables of points appearing below for each factor A-R are meant to serve as a guide to help the user decide on a score. Users should feel free to choose any integer number of points for each factor between the minimum and maximum given in each table. (Thus, users are not restricted to just the values shown in the table.) Scores are meant to be roughly linear, i.e., if a factor doubles in quantity or intensiveness, the number of points assigned to it should approximately double.

One of the most important factors involved in decisions about vulnerability disclosures has to do with the characteristics of the good guys and the bad guys. Factors C-M, P, & Q attempt to deal with this.

Exactly who constitute the "good guys" and who are the "bad guys" should usually be clear from the context. Note, however, that the good guys will often not be 100% good (few government agencies are, for example), nor do the bad guys necessarily have completely malicious goals. For example, while the tactics and extremism of eco-terrorist may well be nefarious, their fundamental concern—protecting natural resources—is not necessarily evil. We should also be careful not to automatically assign "good guy" status to government or authoritarian organizations. A totalitarian regime that uses security measures to suppress its citizens and their civil liberties, for example, does not deserve the title of "good guys".

It is often the case that knowledge of security vulnerabilities is of more help to the good guys than to their adversaries. This is because the good guys usually outnumber the bad guys. (There are, for example, far more bank employees than there are people who are currently active as bank robbers.) Moreover, bad guys usually need to stumble upon only one vulnerability for one target, and can often attack at the time of their own choosing. Security managers, on the other hand, must deal with many vulnerabilities and many possible targets, often extended in time and space. They must even try to manage unknown vulnerabilities. Furthermore, while the bad guys usually fully understand the good guys, the identity of the bad guys is unknown for many security applications. Given this asymmetry between good and bad guys, vulnerability information frequently has more marginal value to the good guys than to the bad guys.

FACTOR A: RISK (0-300 POINTS)

Generally speaking, vulnerabilities that represent minimal risk can be publicly discussed in detail without much concern. Worries about helping the bad guys more than the good guys grow as the risk increases. High-risk vulnerabilities are often best discussed with security managers via private channels, if possible.

With the VDI tool, risk is thought of as the product of the probability of an attack succeeding times the seriousness of the consequences. The term "attack" means an attempt by the bad guys to defeat a security device, system, measure, or program by exploiting the vulnerability in question.

Typically, attacks on the government or public welfare will need to be considered more consequential than attacks on private companies or property.

Table A below provides a lookup table for points to assign to factor A based on the probability of an attacking succeeding, as well as the seriousness of its consequences.

Table A - Factor A, Risk.

Consequences
↓ Probability of attack succeeding ------>

	negligible	low	medium	high	very high
negligible	300	250	200	150	100
low	250	206	162	119	75
medium	200	162	125	88	50
high	150	119	88	56	25
very high	100	75	50	25	0

FACTOR B: OBVIOUSNESS OF THE VULNERABILITY (0-200 POINTS)

If the vulnerability is blatantly obvious to almost any reasonably resourceful person, or if similar attacks have already been suggested publicly thereby making them obvious, there is little point in keeping quiet. Motivated bad guys can figure out obvious vulnerabilities on their own, anyway. If, on the other hand, there has been no previous speculation on this or related vulnerabilities, and only extraordinarily creative, knowledgeable, and clever individuals can figure it out after extensive thought and experimentation, it may well be smart to limit public or detailed discussion of the vulnerability and how to exploit it. (The vulnerability assessors themselves will know if discovering the vulnerability required extensive time and effort, or whether it was spotted almost immediately.)

Security managers often fail to recognize even obvious vulnerabilities—presumably because they are not mentally predisposed to doing so.[2,10]

Table B - Factor B, Vulnerability Obviousness.

obviousness of the vulnerability	points
none	0
a little	50
some	100
a lot	150
very substantial	200

FACTOR C: ATTACK TIME, COST, AND MANPOWER (0-100 POINTS)

If the attack is trivial to prepare, rehearse, and execute—though not necessarily to think up (Factor B)—then a detailed public discussion may be unwise. On the other hand, if few adversaries can marshal the necessary resources, the risk associated with a public disclosure may be minimal.

For this factor, if some of the sub-factors (time, cost, and manpower) are needed in large quantities but others are not, score each separately from 0-100 points, then average them together to get the net score.

If the conditions for preparing and practicing the attack are considerably different from that for executing the attack, consider which is the more important constraint for the given vulnerability, and choose the score for factor C accordingly. (Some attacks, for example, must be executed quickly to be effective, but may take months for preparation and practice.)

Table C - Factor C, Attack Time/Cost/Manpower.

time, cost, & manpower for practice & execution	points
very minimal	0
minimal	25
some	50
a lot	75
very extensive	100

FACTOR D: LEVEL OF SKILL, SOPHISTICATION, AND HIGH TECHNOLOGY (0-100 POINTS)

If the average person on the street can easily exploit the vulnerability, a public airing of details may be unwise. On the other hand, if only highly trained, sophisticated adversaries can pull off the attack, and only after extensive practice with expensive high-tech or social engineering tools, there is probably minimal harm in discussing the attack in some detail. This will allow security managers to better appreciate the problem—and be motivated to fix it.

Attacks on some security devices require great skill, but little technological expertise. (Picking a lock is an example.) If some of the sub-factors (skill, sophistication, and level of technology) are high, but others are low, score each separately from 0-100 points, then average them together to get the net score for this factor.

Table D - Factor D, Attack Skill/Sophistication/High-Technology.

required skill, sophistication, & high-technology	points
very minimal	0
minimal	25
some	50
a lot	75
very extensive	100

FACTOR E: COST, TIME, AND COMPLEXITY OF THE COUNTER-MEASURES OR ALTERNATIVE SECURITY (0-200 POINTS)

If the suggested countermeasures are cheap and easy, a full public disclosure of both the vulnerability and the countermeasures may be warranted. If, however, there are no known countermeasures or alternatives, or they are impractical, expensive, and/or time consuming to put in place, there is typically little chance they will be widely implemented. Being discreet about the vulnerability is therefore indicated. (There is the chance, of course, that somebody else might be able to devise more practical countermeasures if she were made aware of the vulnerability.)

Table E - Factor E, Countermeasures.

cost & complexity of countermeasures	points
very high (or there are no countermeasures)	0
fairly high	50
moderate	100
fairly low	150
very low	200

FACTOR F: RATIO OF CURRENT TO FUTURE USE (0-100 POINTS)

This factor considers the ratio of current use of security to the extent of use likely in 3 years. If the security device, system, measure, or program hasn't been fielded to any great extent, there should be ample time and at least some willingness to fix problems, so a public discussion of vulnerabilities may be warranted. If, on the other hand, the fixes would mostly have to be retrofitted in the field, the odds that this will actually happen is less, and a detailed public disclosure of vulnerabilities may be risky.

Table F - Factor F, Ratio of Current Use of the Device, System, or Program to Use 3 Years in the Future.

ratio of current to future use	points
>5	0
2-5	25
0.5-2	50
0.2-0.5	75
<0.2	100

FACTOR G: NUMBER OF ORGANIZATIONS FOR WHICH THE VULNERABILITY IS RELEVANT (0-200 POINTS)

If the vulnerability is highly localized, e.g., the local ice cream shop has a vulnerability because the manager frequently forgets to lock the back door at night, it clearly makes little sense to widely publicize the vulnerability and alert the bad guys. The vulnerability should quietly be pointed out to the manager or shop owner. If, on the other hand, the vulnerability is shared by a large number of diverse organizations, a public disclosure may be prudent.

The reasons this factor is not the sole, overriding consideration in vulnerability disclosures include the following:

1. We cannot always be 100% certain exactly how many organizations may actually be subject to a given vulnerability.
2. Going public can potentially contribute to better security for organizations and security applications we have not considered. For example, publicly discussing the ice cream shop's vulnerability may remind other unrelated businesses to lock their doors at night.
3. Going public may also help ensure good security practice at future ice cream shops and unrelated businesses that don't currently exist. (Factor G.)
4. Even if we try to carefully channel the vulnerability information by disclosing it to just one or a small number of organizations, there is still a risk that the information will leak out anyway, especially if the organization(s) are large and/or have a poor security culture. (Factors H, I, L, & M.)
5. A public disclosure may pressure the ice cream shop into implementing better security than if the issue is just discussed privately.
6. Even if only one or a small number of organizations are relevant, a public disclosure is relatively safe if the security of those organizations is poor in other ways than just the vulnerability in question. (Factors L & M.)

Note: When there are no relevant organizations, the physical security device, system, measure, or program in question is not in use. Thus, full public disclosure (200 points in the first row) in warranted for factor G because there is no immediate risk.

Table G - Factor G, Number of Vulnerable Organizations

number of organizations	points
0	200
1	0
2 or 3	20
4-9	50
10-20	90
20-50	140
50-100	180
100-200	190
>201	200

FACTOR H: NUMBER OF SECURITY PERSONNEL (0-100 POINTS)

This factor concerns how many people inside the good guys' organizations will ultimately learn about the vulnerability if management is informed. (For many organizations, this nearly equals the number of total security employees, because few organizations are good at compartmentalizing information for any length of time.) The larger the number of people involved, the more likely the vulnerability will be deliberately or inadvertently leaked anyway, so the lower the risk of going public with the vulnerability in the first place.

Table H - Factor H, Number of Security Personnel

typical size of good guys' security force	points
very small	0
small	25
medium	50
large	75
very large	100

FACTOR I: RATIO OF GOOD GUYS TO BAD GUYS (0-200 POINTS)

When good guys greatly outnumber bad guys, openly sharing vulnerability information tends to do more good than harm. For example, there are probably more child care providers than there are pedophiles at risk for molesting children. Thus, publicly providing information on how to protect children is probably prudent. On the other hand, in the case of underage drinkers, there are likely to be more minors interested in illegally obtaining alcohol than there are store clerks and bar bouncers to check IDs, so it may make more sense to disclose vulnerabilities directly to alcohol vendors than to the general public.

Note that for Factor I, only personnel directly involved in relevant security operations should be considered—not the total number of general employees.

Table I - Factor I, Ratio of Good to Bad Guys

ratio of good guys to bad guys	points
<< 1	0
< 1	50
~ 1	100
> 1	150
>> 1	200

FACTOR J: THE ADVERSARY IS KNOWN (0-100 POINTS)

If the bad guys are well known, it may be prudent to carefully direct the flow of vulnerability information away from them. On the other hand, when the identity of the bad guys is largely unknown, e.g., they might even be unknown insiders within the security organization, we have less of an opportunity to effectively direct the flow of vulnerability information. A public disclosure is then more warranted.

Table J - Factor J, Bad Guys Identity.

how well the bad guys are known	points
fully identified	0
fairly well known	25
somewhat known	50
slight idea	75
total mystery	100

FACTOR K: THE DEGREE TO WHICH THE SECURITY DEPENDS ON SECRECY (0-100 POINTS)

Secrecy is not usually a good long-term security strategy. [18] That's because people and organizations are typically not very good at keeping secrets. Thus, if security is largely based on a misplaced faith in secrecy, taking actions to end over-reliance on secrecy could actually be healthy.

A public discussion of vulnerabilities may force good guys who rely mostly on secrecy to implement better security measures. It is, for example, believed that publicly discussing software vulnerabilities forces manufacturers to fix security problems faster and better.[11,19] In any event, holding private discussions with security managers who rely mostly on secrecy is unlikely to result in improved security because they will (at least in the author's experience) tend to foolishly count on the vulnerability remaining a secret.

Table K - Factor K, Secrecy.

security is primarily based on secrecy	points
not at all	0
just a little	25
some	50
a lot	75
completely	100

FACTOR L: THE EFFICACY OF THE OTHER SECURITY MEASURES (0-120 POINTS)

If an organization has extremely poor general security, there are already multiple vulnerabilities to exploit. Thus, the risk from a public disclosure of a single vulnerability is greatly lessened. Moreover, a public disclosure might pressure the good guys into improving overall security, not just deal with the immediate vulnerability in question. If, on the other hand, the security is generally outstanding except for the sole problem(s) that have been identified, a public disclosure might help the bad guys succeed where they would otherwise have failed.

Table L - Factor L, Overall Security Effectiveness.

overall effectiveness of security	points
excellent	0
good	30
fair	60
poor	90
very poor	120

FACTOR M: THE SOPHISTICATION OF THE GOOD GUYS (0-300 POINTS)

When security managers and other security personnel don't fully understand the security devices, systems, or programs they are using, and lack awareness of the important vulnerabilities, we are probably better off being very public and detailed in discussing the vulnerability in question. If the good guys think no vulnerabilities are even possible—a distressingly common situation in the field of physical security—this factor should be assigned a large number of points.

Table M - Factor M, Security Sophistication

sophistication of the good guys	points
excellent	0
good	75
some	150
just a little	225
none	300

FACTOR N: "SILVER BULLET" ATTITUDES (0-200 POINTS)

This factor considers the degree to which the security device, system, measure, or program is generally viewed by government, business, end-users, potential end-users, and the public as a security panacea. If the security is thought to magically provide invincible security, a detailed public discussion of the vulnerability is probably healthy. Even though the bad guys might also temporarily believe in the myth of invincibility, the good guys cannot count on this indefinitely because the bad guys will tend to think more critically about security vulnerabilities than the good guys.

Examples of security technologies that have clearly been viewed—quite incorrectly—as "silver bullets" (panaceas) include RFIDs, GPS, biometrics, encryption, and tamper-indicating seals.[3]

Table N - Factor N, Panacea & Overconfidence Illusions.

security is viewed as as largely invincible	points
not at all	0
a little	50
some	100
a lot	150
completely	200

FACTOR O: THE EXTENT OF OVER-HYPING (0-120 POINTS)

If the security device, system, measure, or program is being over-hyped by manufacturers, vendors, or other proponents, a detailed public discussion of the vulnerabilities is probably healthy and will ultimately result in better security. Over-hyping is a serious problem for physical security because of the relative lack of rigorous standards, metrics, principles, and testing guidelines, as well as effective research & development.[2,9,10]

Symptoms of over-hyping include sloppy terminology, or exaggerated and absolutist phrases such as "tamper-proof", "completely secure", "*impossible* to defeat", "passed all vulnerability assessments". Other indications of over-hyping are the misuse or misrepresentation of statistics and tests, deliberate obfuscation, or comparing apples and oranges.[2]

Table O - Factor O, Over-Hyping.

amount of over-hyping	points
none	0
a little	30
some	60
a lot	90
completely	120

FACTOR P: HOW MUCH ARE THE BAD GUYS LIKELY TO BENEFIT? (0-120 POINTS)

If the bad guys have (or believe they have) little to gain from exploiting a vulnerability, then there is probably little risk to a full public discussion. Of course, what the bad guys hope to gain depends on the context. Crooks would be interested in economic gain, disgruntled individuals in retaliation, terrorists in disruption and death, radicals in making political statements, hackers in demonstrating prowess, and vandals in entropy.

This factor deals with how the bad guys can benefit, whereas the factor A (risk) dealt with how much the good guys have to lose (and the probability).

Table P - Factor P, Bad Guys Benefit.

bad guys stand to gain	points
a tremendous amount	0
a lot	30
some	60
just a little	90
nothing	120

FACTOR Q: HOW SUBSTANTIAL ARE THE PENALTIES TO BAD GUYS IF THEY ARE CAUGHT? (0-80 POINTS)

Some illegal activities, such as product counterfeiting or copyright violations, carry relatively light legal penalties, or else the laws are rarely enforced. If the bad guys face little risk from exploiting a vulnerability, they may be more likely to proceed. A public disclosure of the vulnerability is therefore more risky.

Table Q - Factor Q, Penalties.

extent of likely penalties	points
negligible	0
a little	20
some	40
a lot	60
very substantial	80

FACTOR R: MOTIVATION OF THE INDIVIDUALS CONTEMPLATING A VULNERABILITY DISCLOSURE (0-160 POINTS)

While good things can be done for bad reasons, and vice versa, it is worth considering the motivation of the would-be discloser. If he or she wants to disclose the existence of vulnerabilities primarily for selfish reasons, it might be prudent to exert at least a partial restraint on full disclosure. Obvious conflicts of interest need to be considered as well, e.g., the vulnerability assessors are evaluating a product made by a competitor of their employer.

This factor requires the VDI tool user to attempt to gauge motivation. If the vulnerability assessor himself is using the tool, he will need to undertake a certain amount of honest introspection that may be healthy when considering disclosure issues.

Table R - Factor R, Assessor Motivation.

motivation	points
entirely self-promotion or self-interest; major conflict of interest	0
partially self-promotion or self-interest	40
a mix of self-interest and altruism	80
mostly altruistic	120
entirely altruistic; zero conflict of interest	160

INTERPRETATION

The overall VDI score is computed as follows. The sum of the points from all the factors (A-R) is computed, then normalized to (divided by) the maximum possible number of points (2800), and finally multiplied by 100 to produce a VDI value in percent. The higher the VDI percent, the more appropriate it is to widely disseminate detailed information about the vulnerability in question. Thus, VDI in percent = [Σ(scores for factors A through R) / 2800] x 100%

The recommendations that the model makes for various VDI scores are shown in table S. The term “fully enabling” means enough detail about the vulnerability is presented to allow anyone sufficiently qualified to reproduce a viable attack on the relevant security device, system, measure, or program with minimal effort. “Partially enabling” means only incomplete information is provided, while “not enabling” means the disclosure provides little practical guidance to an adversary about exactly how to exploit the discovered vulnerability.

Table S - Recommended Course of Action Based on VDI Scores.

VDI score	Recommended level of vulnerability disclosure
>75%	public release, fully enabling
68%-75%	public release, partially enabling
60%-67%	public release, non-enabling
50%-59%	restricted release (security trade journals & meetings), fully enabling
40%-49%	restricted release (security trade journals & meetings), partially enabling
34%-39%	restricted release (security trade journals & meetings), non-enabling
12%-33%	highly restricted, private release: contact the relevant good guys directly
<12%	no disclosure at all

Note that for VDI scores in the range 34%-59%, the recommendation in table S is for disclosure, but only to an audience of security professionals. This can be done by using security trade journals and security conferences. While such forums cannot be guaranteed to be free of bad guys, they probably have a higher ratio of good guys to bad guys than would be the case for the general public.

It is also important to bear in mind that the recommended choice of action from table S does not automatically preclude those actions listed below it in the table. For example, if the VDI score calls for a non-enabling public disclosure of the vulnerability, this does not preclude more detailed, enabling discussions in private with good guys at a later time. The publicity surrounding the disclosure of a vulnerability (even if non-enabling) may elicit inquiries from good guys who have a legitimate need to know more details. The typical problems with vague public disclosures, however, are that (1) they may not reach the most important audience, and (2) they may not be taken seriously if details or demonstrations are not provided.

EXAMPLES

Five examples are presented in this section, with 1-3 being hypothetical. These 5 examples are used to check whether the guidance offered by the VDI index is reasonable. At least in the author's view, the recommended courses of action that come from the VDI tool seem sensible for all 5 examples. This, however, is far from a rigorous validation of the model.

Table T shows the points assigned to each factor for the 5 examples, as well as the total points and the resulting VDI scores.

Example 1: The mascot for Dunderhead State University is a billy goat. Loss or harm to the mascot could cause serious damage to the University's pride, and undermine the morale of the Fighting Scapegoats football team and their supporters. A subtle vulnerability has been discovered in the security provided for the mascot, making it very easy for students and fans from competing schools to kidnap or otherwise harm the mascot. Fixing the problem is possible, but complicated. The vulnerability is unique to Dunderhead State and the one location where the mascot is kept. The overall VDI percentage computed from Table T is 29%, indicating (from table S) that we should discuss the matter only with University students and staff responsible for the mascot's security and welfare. A public disclosure would be imprudent.

Example 2: A simple but non-obvious method is found for stealing candy bars from vending machines. The attack can be eliminated by quickly snapping a cheap piece of plastic into the interior of the machine the next time it is refilled. From table T, the overall VDI score is 44%, indicating (from table S) that we should do a partially enabling disclosure to security professionals and vending companies, including possibly some discussion of the countermeasure.

Example 3: A (widely respected) company hired by many organizations to perform background checks on security personnel is discovered to have done poor quality work, and may even have faked much of the data. The company's competitors do not seem to have this problem, though switching vendors is somewhat expensive. The overall VDI percentage in table T is 51%, indicating that we should do a fully enabling disclosure to general security professionals about the problem, probably going so far as to even identify the company.

Example 4: Lawrence M Wein raised a controversy about whether a paper discussing terrorist poisoning of milk with botulinum toxin should be openly published.[20,21] Here, we will assume that this theoretical attack would have major consequences, but a relatively low probability of success[22]. In addition, we shall assume—as Leitenberg and Smith maintain[22]—that a terrorist would need considerable sophistication, skill, time, and effort to obtain significant quantities of the botulinum toxin. Under these assumptions and the author's view of the situation (which may or may not be correct), table T shows an overall VDI percentage of 62%, indicating that the vulnerability should be discussed openly in a non-detailed manner. Given that the paper itself is not very enabling[22], this is essentially what the National Academy of Sciences actually decided to do when it chose to publish the paper despite government objections.[23]

Example 5: The VAT has demonstrated how easy it is for relatively unsophisticated adversaries to spoof—not just jam—civilian GPS receivers using widely available commercial GPS satellite simulators.[5,6] Unlike the military signal, the civilian GPS signal is not encrypted or authenticated. Even though it was never designed for security applications, it is frequently used that way. Most GPS users are unaware of the vulnerability. Prior to developing the VDI tool, the VAT made the decision to publicly disclose the vulnerability. This disclosure was partially enabling in that the use of a GPS satellite simulator was discussed. After developing the VDI tool, the VAT scored the GPS vulnerability as shown in Table T. The VDI score of 69% supports our prior intuitive decision to do a partially enabling public release.

Table T - Scores for Each VDI Factor for the 5 Examples.

	Example 1 (mascot)	Example 2 (candy bars)	Example 3 (bkg checks)	Example 4 (toxic milk)	Example 5 (GPS)
Factor A	130	119	56	119	60
Factor B	25	20	25	150	100
Factor C	25	10	10	75	60
Factor D	25	10	5	75	60
Factor E	40	190	110	120	150
Factor F	50	50	65	45	100
Factor G	0	200	200	200	200
Factor H	10	50	80	20	80
Factor I	50	0	60	195	180
Factor J	25	95	50	90	90
Factor K	70	5	90	20	40
Factor L	10	40	60	70	60
Factor M	70	110	150	150	290
Factor N	100	50	150	110	195
Factor O	10	10	90	75	115
Factor P	70	90	50	60	35
Factor Q	20	30	50	70	40
Factor R	80	140	120	80	80
Sum of Points	810	1219	1421	1724	1935
VDI	29%	44%	51%	62%	69%

DISCUSSION

The VDI score computed in this model is meant to provide guidance on the maximum amount of vulnerability information (if any) that should be disclosed. Generally, it is prudent to release no more information about a vulnerability to no more people than is necessary to accomplish what needs to be done, i.e., alert security managers to a problem, create more realistic views about security, and/or get countermeasures implemented. Minimizing the amount of information and the people who receive it reduces the odds that it will benefit the bad guys—but, as discussed above, it also reduces the odds that the good guys will take necessary actions.

At best, the VDI tool should be considered only a preliminary attempt to encourage thinking and discussion of vulnerability disclosure issues. The tool cannot be the final arbitrator for whether to disclose security vulnerabilities, in what degree of detail, when, or to whom. Every case is different, and there are other, sometimes overriding factors that must also be considered but are missing from the VDI model. These include government classification regulations, state and federal laws, organizational & employer rules, proprietary and intellectual property issues, legal liabilities[24], contractual obligations such as who sponsored the vulnerability assessment and who owns its results, and personal views on morality, fairness, and social responsibility. The author of this paper and the VDI tool can make no claim to any unique insight or wisdom on any of these matters.

There are other limitations to this tool as well. While the various factors (A-R), their scoring, and relative weights seem plausible, it is very difficult to rigorously defend specific details of the VDI tool. Questions very much open for debate include:

- What factors are missing?
- What factors A-R are correlated or "non-orthogonal" and should be combined into some other, more general factor?
- Are the relative weights of the factors (i.e., the maximum possible number of points for each factor) appropriate?
- Does the roughly linear assignment of points in the table for each factor make sense?
- Should the recommended course of action for the various ranges of VDI scores in table S be different? (Admittedly the break points in column 1 of table S are somewhat arbitrary.)

In terms of weighting, the factor weights are as follows:
A=M > B=E=G=I=N > R > L=O=P > C=D=F=H=J=K > Q.
This weighting, while very much open for debate, is not arbitrary. In the view of the author, the factors with the highest possible scores (or weights) probably are indeed the most critical.

It also is very important to avoid the "fallacy of precision". This is thinking that because one has assigned numeric values to complex parameters, then he or she automatically has a rigorous understanding of them. The fact is that quantified ambiguity is still ambiguity.

Despite the myriad potential problems with the VDI tool, it does nevertheless serve as a means for raising many of the critical issues associated with the disclosure of vulnerabilities. Anyone conscientiously using the tool automatically demonstrates that he or she has at least made a rudimentary attempt towards sincerely considering the risks and implications of disclosing vulnerabilities. The VDI score can help to justify the decision to disclose or not to disclose. As such, the tool may be of some value for protecting vulnerability assessors and others from the retaliation and recrimination that all too commonly arises when vulnerability issues or questions about security are raised in good faith.[10,11,15,1625] The VDI tool might also help the user choose a more appropriate channel, medium, or forum for vulnerability disclosures than he or she might be otherwise inclined to pursue, e.g., the popular press or the Internet vs. security conferences and journals vs. private discussions with manufacturers or end-users.

REFERENCES

[1] Vulnerability Assessment Team Home Page, http://www.ne.anl.gov/capabilities/vat.

[2] Roger G. Johnston, "Assessing the Vulnerability of Tamper-Indicting Seals", Port Technology International 25(2005): 155-157.

[3] Roger G. Johnston and Jon S. Warner, "The Dr. Who Conundrum", Security Management 49(2005): 112-121.

[4] Roger G. Johnston, Anthony R.E. Garcia, and Adam N. Pacheco, "Efficacy of Tamper-Indicating Devices", Journal of Homeland Security, April 16, 2002, http://www.homelandsecurity.org/journal/Articles/displayarticle.asp?article=50.

[5] Jon S. Warner and Roger G. Johnston, "A Simple Demonstration that the Global Positioning System (GPS) is Vulnerable to Spoofing", Journal of Security Administration 25(2002): 19-27.

[6] Jon S. Warner and Roger G. Johnston, "GPS Spoofing Countermeasures", Homeland Security Journal, December 12, 2003, http://www.homelandsecurity.org/bulletin/Dual%20Benefit/warner_gps_spoofing.html.

[7] Morten Bremer Maerli and Roger G. Johnston, "Safeguarding This and Verifying That: Fuzzy Concepts, Confusing Terminology, and Their Detrimental Effects on Nuclear Husbandry", Nonproliferation Review 9(2002): 54-82, cns.miis.edu/pubs/npr/vol09/91/91maerli.pdf.

[8] Roger G. Johnston and Morten Bremer Maerli, "International vs. Domestic Nuclear Safeguards: The Need for Clarity in the Debate Over Effectiveness", Disarmament Diplomacy, issue 69, February-March 2003, http://www.acronym.org.uk/dd/dd69/69op01.htm.

[9] Roger G. Johnston and Morten Bremer Maerli, "The Negative Consequences of Ambiguous 'Safeguards' Terminology", INMM Proceedings, July 13-17, 2003, Phoenix, AZ.

[10] Roger G. Johnston, "Effective Vulnerability Assessments", Proceedings of the Contingency Planning & Management Conference, Las Vegas, NV, May 25-27, 2004.

[11] Bruce Schneier, "Is Disclosing Vulnerabilities a Security Risk in Itself?", InternetWeek, November 19, 2001, http://www.internetweek.com/graymatter/secure111901.htm.

[12] M. Rasch, "'Responsible Disclosure' Draft Could Have Legal Muscle", SecurtyFocus, November 11, 2002, http://online.securityfocus.com/columnists/66.

[13] A. Arora, R. Krishnan, A. Nandkumar, R. Telang, and Y. Yang, "Impact of Vulnerability Disclosure and Patch Availability—An Empirical Analysis", April 2004, http://www.dtc.umn.edu/weis2004/telang.pdf.

[14] A. Arora and R. Telang, "Economics of Software Vulnerability Disclosure", Security & Privacy 3(2005): 20-25.

[15] E. Hall, "Risk Management Map", Software Tech News 2(2004), http://www.softwaretechnews.com/technews2-2/stn2-2.pdf.

[16] M.A. Caloyannides, "Enhancing Security: Not for the Conformist", Security & Privacy 2(2004): 86-88.

[17] Roger G. Johnston, "Tamper Detection for Safeguards and Treaty Monitoring: Fantasies, Realities, and Potentials", Nonproliferation Review 8(2001): 102-115, http://www.princeton.edu/~globsec/publications/pdf/9_2johnston.pdf.

[18] Roger G. Johnston, "Cryptography as a Model for Physical Security", Journal of Security Administration 24(2001): 33-43.

[19] A. Arora, R. Telang, and H. Xu, "Timing Disclosure of Software Vulnerability for Optimal Social Welfare", November 2003, http://www.andrew.cmu.edu/user/xhao/disclosure.pdf.

[20] Lawrence M. Wein, "Got Toxic Milk", New York Times, May 30, 2005, http://www.nytimes.com/2005/05/30/opinion/30wein.html?ex=1275105600&en=e56b2b8b96d56f1e&ei=5088.

[21] Rebecca Carr, “Publication Heeds U.S., Pulls Terror Article”, Atlanta Journal and Constitution, June 26, 2005, http://www.ajc.com/hp/content/auto/epaper/editions/sunday/news_24ebc541731e70fe0050.html.

[22] M. Leitenberg and G. Smith, “’Got Toxic Milk?’: A Rejoinder”, (2005), http://www.fas.org/sgp/eprint/milk.html.

[23] Scott Shane, “Paper Describes Potential Poisoning of Milk”, New York Times, June 29, 2005, http://www.nytimes.com/2005/06/29/politics/29milk.html?ex=1277697600&en=06b46176c5d1a2cf&ei=5088&partner=rssnyt&emc=rss.

[24] J. Stisa Granick, “Legal Risks of Vulnerability Disclosure”, (2005), http://blackhat.com/presentations/win-usa-04/bh-win-04-granick.pdf.

[25] Anonymous, “Don’t Shoot the Messenger”, CSO 5(2006): 52-53, http://www.csoonline.com/read/080106/col_undercover.html.