

LAUR-04-9004

Invited Article for *Port Technology International* 25, 155-156 (2005)

Assessing the Vulnerability of Tamper-Indicating Seals

Roger G. Johnston, Ph.D., CPP
Vulnerability Assessment Team
Los Alamos National Laboratory
MS J565, Los Alamos, NM 87545 USA
rogerj@lanl.gov, phone: 505-667-7414, fax: 505-665-4631
URL: <http://pearl1.lanl.gov/seals/default.htm>

Introduction

The Vulnerability Assessment Team (VAT) at Los Alamos National Laboratory (LANL) has studied tamper detection for 13 years. We have conducted vulnerability assessments (VAs) on hundreds of seals and cargo security programs, and undertaken research and consulting for over two dozen government agencies and private companies. This article discusses how we conduct VAs and what we have learned about seals.

Terminology & Concepts

Tamper-indicating seals are widely used for cargo security and customs applications. Unlike locks, seals are not meant to resist or delay unauthorized access. Instead, they record that it took place.

A **barrier seal** is a single device that is both lock and seal. It's a compromise, neither the best seal for a given application nor the best lock. Barrier seals are unfortunately sometimes called **security seals** but this confuses matters. Even flimsy seals made of paper or plastic are security devices.

A seal's effectiveness depends critically on its **use protocol**. This is the official and unofficial procedures for using the seal, including procurement, shipping, checkout, installation, inspection, removal, disposal, postmortem examinations (if any), record keeping, interpretation, response to evidence of tampering, and training.

The focus of a seal VA is on the use protocol, and on how to **defeat** the seal. Defeating a seal means removing it, then resealing using either the original seal or a counterfeit, without being detected. Simply cutting a seal off a container is not defeating it. The fact that the seal is damaged or missing will be noted at the time of inspection. To **attack** a seal means to undertake actions intended to defeat it.

People sometimes talk about "**tamper-proof**", "**tamper-resistant**", or **antipilferage seals** but these terms are misleading and unhelpful. "Tamper-proof" implies invincibility, which is unlikely and wholly unprovable. (Besides, a seal that can't be tampered with, can't leave evidence of tampering!) And resisting tampering or pilferage is what locks do, not seals.

Understanding Vulnerability Assessments

The purpose of a vulnerability assessment (VA) is often misunderstood. A VA is undertaken to improve the security of a seal or a cargo security program. The goal is not to "certify" the seal (or security program)--and especially not to bless a seal for all time, for any possible application, against all potential adversaries. A seal or cargo security program does not "pass" a VA anymore than a person "passes" an IQ test.

It is important not to confuse **seal testing** with VAs. Seals can be tested for mechanical strength, environmental durability, ergonomics, quality control, and other characteristics.

These tests are useful, and may even have important security implications. They are not, however, the same thing as a VA.

Another common misconception is that a VA should ideally find no vulnerabilities. In fact, such a VA is worthless and needs to be redone correctly. Vulnerabilities are always present in any security device, system, or program--and present in very large numbers. Finding a vulnerability is actually good news, not bad news, because we then have an opportunity to mitigate it.

In addition to finding vulnerabilities and devising/demonstrating attacks, an effective VA should recommend countermeasures and possibly demonstrate them. Countermeasures may involve modifying the seal and/or its use protocol.

The VAT does several different kinds of VAs. We may analyze a new seal at various stages during its design process. Unfortunately, we are more often asked to assess a seal that is ready for, or already in, production. By then, it is usually too late to make changes. It is most useful to assess a seal in the context of a specific user, application, and use protocol. Sometimes, however, we must study seals in isolation; we then recommend appropriate users, applications, and protocols. The VAT has also conducted VAs on entire cargo security programs, where seals are but one component.

How the VA Process Works

Real adversaries may or may not be creative (though it is dangerous to underestimate them). Vulnerability assessors must be. Whereas bad guys need to stumble upon only one successful attack, assessors must worry about many. The types of individuals who tend to be good at VAs include smart alecks, cynics, trouble makers, schemers, loophole finders, questioners of tradition and authority, artists, mechanics, tinkerers, inventors, and hackers.



Figure 1 - The LANL Vulnerability Assessment Team. Members include an artist, a former automobile body repair man, an industrial psychologist, Ph.D. physicists, mechanical technicians, and students with security clearances.

In the VAT, we follow a 12-step process in conducting VAs:

1. Study the device, system, or program to learn how it is really used. Be sure to talk to low-level personnel because security managers, designers, vendors, and manufacturers often have incomplete or unrealistic perspectives.
2. Play and experiment with the device, system, or program.
3. Brainstorm potential attacks. Like brainstorming in any field, it is absolutely essential that there be no premature criticism of ideas. Wild, impractical ideas are not just acceptable, they are essential. Thus, we encourage thinking about attacks that involve, for example, flying monkeys, space aliens, or Elvis impersonators. Crazy ideas can sometimes be morphed into

something practical. More importantly, they encourage us to further think “outside the box”. Brainstorming is a tricky balance between individual and group psychology. Creativity comes from individuals, not from groups, but the right group dynamics can create a productive, idea generating environment.

4. Play with the device, system, or program again based on ideas developed during brainstorming.
5. Edit and prioritize potential attacks. (Only at this stage do the crazy ideas get modified or rejected.) It is essential to set priorities given the large number of possible seal attacks, the finite time and funding available, and the fact that VAs do not have a clear endpoint.
6. Partially develop some of the most promising attacks.
7. Determine feasibility of the attacks.
8. Devise countermeasures.
9. Perfect attacks.
10. Demonstrate attacks.
11. Rigorously test attacks.
12. Rigorously test countermeasures.

Despite the linear process outlined here, it is important to stay flexible. The best attacks can pop into one’s head at any point in the process. Note also that sponsors of VAs are rarely willing to pursue steps 9-12 because they are expensive, time-consuming, and challenging to make realistic. Moreover, steps 9-12 vividly demonstrate that vulnerabilities really exist, typically making security managers uncomfortable.

VA Principles

We keep in mind the following principles when conducting VAs:

- Simple, low-tech attacks (even on high-tech seals) should be examined first because they are usually sufficient.
- If an attack is simple, quick, and cheap (as most are), extreme realism in testing the attack is probably unnecessary. This is not the case if an attack takes substantial skill, resources, and/or time.
- The bad guys, not the good guys, get to define the problem. For example, just because a seal has anti-counterfeiting features does not require an adversary to attack via counterfeiting. Similarly, adversaries are not obligated to attack barrier seals with brute force.
- A competent VA will produce more suggestions and countermeasures than are likely to be implemented. It is up to the appropriate security manager—not the vulnerability assessors—to decide which (if any) make sense to employ.
- VA findings and recommendations should be reported to the highest appropriate level without editing, interpretation, or censorship by middle level personnel.
- Vulnerability assessors need to praise the good things they find. We want the effective features or practices to be recognized and to continue. A VA report that contains nothing but criticisms won't constructively engage security managers, nor make them eager to arrange for future VAs.
- Don't forget Rohrbach's Maxim: No security device or system will ever be used properly (the way it was designed) all the time.
- Don't forget Shannon's Maxim: Secrecy is not a viable long-term security strategy. We must assume the adversary knows and understands the security strategies and hardware being used.
- High-tech, inventory tags such as bar codes, RFIDs, and contact memory buttons are rarely of much use for tamper detection. They are generally easy to remove/reapply, counterfeit, or spoof.

The Large Universe of Possible Attacks

The VAT has identified at least 105 distinct types of seal attacks, most falling into 1 of 10 categories:

Pick Attacks: pick the seal so that it opens without damage or evidence.

Unsealing Attacks: open the seal, then repair or hide any damage or evidence.

Backdoor Attacks: put a defect in the seal prior to use that can be exploited later. Do this during the design or manufacturing process, during shipping or storage, or just prior to use.

Tampering with the Seal Data: tamper with data (such as the seal serial number), or reports and interpretations about the seal inspection.

Seal Reader Attacks: physically tamper with, or otherwise spoof, the electronic or optical seal verifier (if any).

Electronic Attacks: for electronic seals, attack various components such as the sensors, microprocessor, signals, power, annunciator, encryption, or stored alarm condition.

Replicating: make a duplicate seal at the factory using procurement, breaking and entering, bribery, coercion, or social engineering.

Counterfeiting: make a duplicate seal outside of the factory. This can be relatively easy because new or used seal parts are available, and because only the seal's superficial appearance and apparent performance usually need to be mimicked.

Failure Mode Attacks: challenge the seal security program directly or with misdirection, or wait until an error is made and then exploit it.

Sabotaging the Sealing Process: use an insider or outsider to compromise the sealing process, such as applying the wrong seal or not closing the door prior to sealing it.

Better Seals

All the seals we have studied can be defeated. We believe this will always be the case. Better seals, however, are certainly possible. The fundamental problem is what to do with the information that tampering has occurred. In a convention seal, the “alarm condition” must be stored until the seal can be inspected. But typically, an adversary can too easily hide or erase the alarm condition, or make a fresh counterfeit seal.

“Anti-evidence” seals are a better approach. We store information at the very start (when the seal is first installed) that tampering has NOT occurred. This information gets quickly erased when the seal is opened. Adversaries thus have no alarm condition to hide, erase, or counterfeit.

Anti-evidence seals have a number of interesting attributes, including 100% reusability (even if mechanical), and no need for hardware outside the container. They also provide their own intrinsic check for “**gundecking**”, i.e., when the seal inspector claims to have checked the seal, but really didn’t.

Seal Psychology & Double Standards

Defeating a seal (unlike defeating locks, safes, or vaults) is primarily about fooling human beings, not about beating hardware. This is true even for high-tech electronic seals read with an automated reader.

Few people are under the illusion that locks, safes, or vaults provide absolute security. The idea that seals can be defeated, however, often produces hysteria, vehement denial, or (even worse) the belief that “seals are no good so we shouldn’t use them”. It is a mystery why tamper detection invokes such absolutist, emotional, unrealistic attitudes.

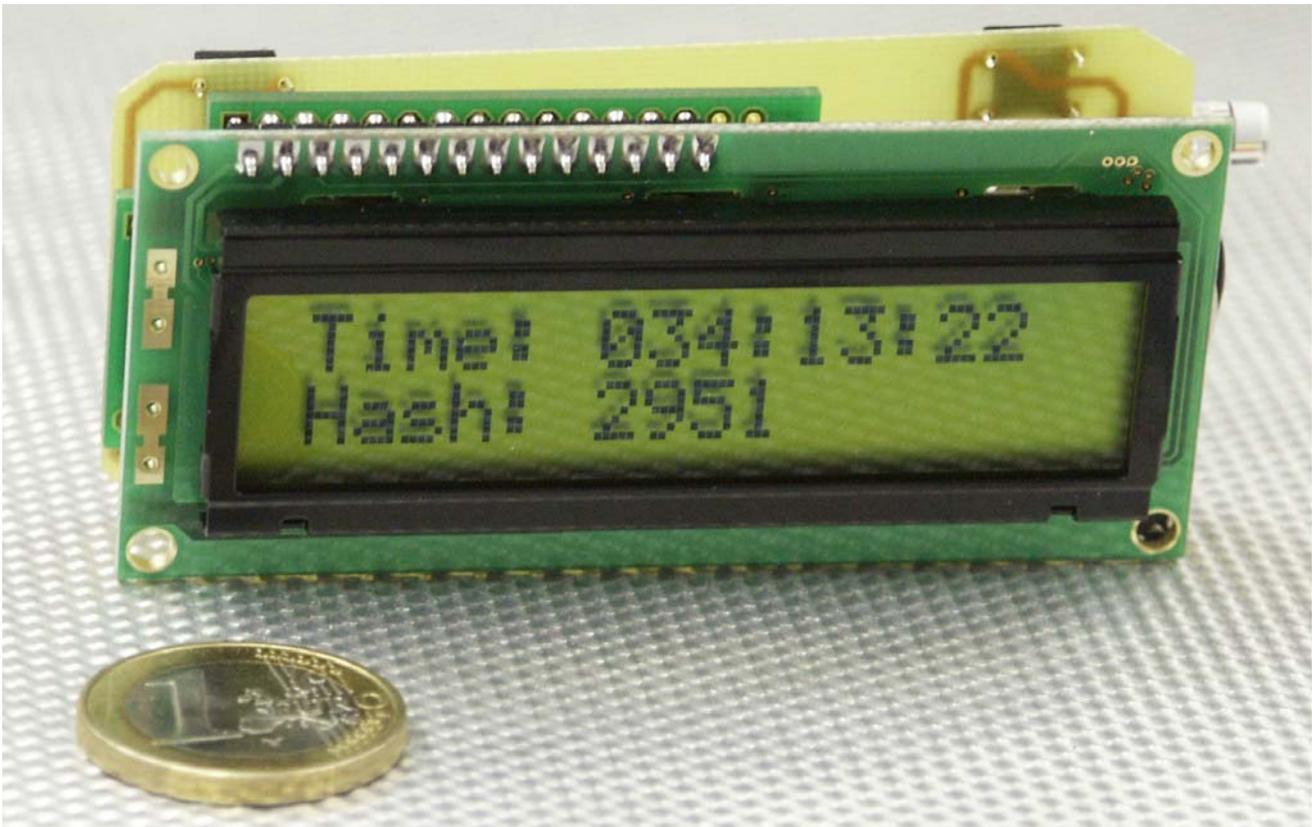


Figure 2 - The “Time Trap”: One of 20 Different Kinds of LANL Anti-Evidence Seals. This seal’s “serial number” (hash) changes over time in ways that the bad guys can’t predict or counterfeit. The hash is only displayed when the container is opened. Future hashes get instantly erased. In this photo, the container was opened on February 3 at 1:22 PM. Parts for this prototype cost under \$8 in retail quantities of 1, and the seal is fully reusable.

The fact that any security device or system can be defeated is a given. The more interesting question is, what practical measures can we introduce to improve security? In our experience, there almost always are simple, inexpensive countermeasures that dramatically improve a given seal’s ability to detect tampering. Absolute security is unobtainable, but a seal (and its container) can provide effective tamper detection if used intelligently, diligently, and with a good understanding of likely attack scenarios.

There are other psychological problems that commonly plague tamper detection. These include naive overconfidence in high-technology, fantasies about silver bullets and free lunches, choosing seals based on gossip or innuendo rather than rigorous analysis, obsession with unit cost while ignoring more important economic factors, a curious mix of fear and arrogance associated with many security programs, and an eagerness to “shoot the messenger”, i.e., retaliate against vulnerability assessors and others who raise security concerns.

Acknowledgments & Disclaimer

Jon Warner, Anthony Garcia, Ron Martinez, Leon Lopez, Adam Pacheco, and Sonia Trujillo contributed substantially to this work. The views expressed here are those of the author and should not necessarily be ascribed to LANL or the US Department of Energy.