

# Vulnerability Assessment of Passive Tamper-Indicating Seals

■  
Roger G. Johnston, Anthony R.E. Garcia, and W. Kevin Grace  
Chemical Science and Technology Division  
Los Alamos National Laboratory  
Los Alamos, New Mexico  
■

## Abstract

We examined 79 different passive, tamper-indicating seals. A number of these are in use, or under consideration, for nuclear security, nonproliferation and weapons control, radioactive waste management, and nuclear materials accountability. We learned how to spoof all the seals using rapid, low-tech methods. Cost was not a good predictor of seal security. It appears to us that many of the seals can be dramatically improved with minor, low-cost modifications. Seal users and manufacturers are encouraged to contact us to discuss specifics.

## Introduction

Tamper-indicating seals are widely used in industry and government for a variety of applications, including access control, records integrity, inventory, shipping integrity, hazardous material accountability, customs control, theft prevention/detection, counterterrorism, counterespionage, protecting instrument calibration, testing for illegal drug use, and consumer protection.<sup>1,2</sup> For nuclear applications, seals are widely used for nonproliferation and weapons control, access control, nuclear materials accountability and radioactive waste management.<sup>3-5</sup>

Seals do not stop unauthorized access, but are intended to leave unambiguous, nonerasable evidence of entry or tampering. Passive seals require no external power. They are popular for nuclear applications because of their safety, low cost, small size, portability, ruggedness, disposability, simplicity and ease of use. Passive seals often utilize pressure sensitive adhesive tapes, brittle materials, fiber optics, crimped cables or other (supposedly) irreversible mechanical assemblies.

As part of a comprehensive project on vulnerability assessment, we studied 79 different passive seals. They are categorized in Table 1 to the right. All but four are commercially available. These seals are widely used by both industry and government. To our knowledge (which is probably incomplete), at least 20 are currently in use for nuclear ap-

plications, and at least seven others are under consideration.

We devised and demonstrated successful attacks on all 79 seals. A total of 91 different successful attacks were developed for the 79 seals (one, two or three per seal). All attacks were low-tech and can be successfully performed by anyone who has access to a hardware store and a standard machine shop, has sufficient practice, and is reasonably skilled with his/her hands, at the level of an average artist or artisan. For some attacks, none of these attributes are required.

The results of our vulnerability analysis are presented here solely in statistical form. We do not wish to single out specific commercial products for criticism, nor to freely disseminate information on how to defeat widely used seals. Rather, we emphasize the lessons and conclusions that this vulnerability assessment has to offer.

## Definitions

A successful attack is defined as opening the seal, then re-sealing it or replacing it with a counterfeit such that the entry or tampering goes undetected. We classified successful attacks into three categories: type 1, 2 or 3. In all three types,

Table 1: Types of the 79 seals used in this study.

Type of seal	Number of seals
Plastic loop	15
Wire loop	4
Metal cable	13
Metal ribbon	10
Bolt type	7
Fiber optic	2
Adhesive tape	27
Other	1

the seal is broken, then repaired or counterfeited.

In a type-1 attack, the tampering will not be detected if the usual inspection process is followed. The usual process is either that recommended by the manufacturer of the seal, or the inspection process typically employed by end-users. The tampering will be detected, however, if unusual efforts are taken. For many seals, an example of an unusual inspection process would be to disassemble the seal and examine it in great detail to look for tampering.

In a type-2 attack, the tampering will not be detected even if unusual (but low-tech) inspection occurs, such as is disassembling the seal and examining it in detail by eye.

In a type-3 attack, the tampering cannot be detected even if the most advanced postmortem analysis is done. State-of-the-art techniques in forensics, material science or microscopy will not be able to tell that the seal was broken or counterfeited.

## Results

Only demonstrated attacks are considered here. For most of the seals, we have devised, but not yet fully demonstrated, one or more alternative (usually low-tech) attacks. Out of our 91 demonstrated attacks, 37 were classified as type 1, 42 as type 2 and 12 as type 3 (the most thorough defeat). Most of the attacks can be completed using tools and materials that will fit inside a briefcase or, in some cases, a pocket or the palm of a hand.

Only four of the 91 attacks developed in this study involve counterfeiting, that is, removing the original seal, then replacing it with a counterfeited duplicate. The majority of the attacks involved opening the seal, then resealing it and repairing the damage (if any).

Counterfeiting, nevertheless, appears to be relatively simple for most of the seals. Manufacturers frequently make counterfeiting easier by providing free samples of the seals to anyone who asks; using readily available materials or components; using easily replicated colors, logos or numbering; and using embossing or stamping for logos or numbers that is so shallow it can be easily buffed off and replaced with an alternative embossing or impression.

With practice, the time to successfully complete the attacks varied from three seconds for three of the seals to 125 minutes for the most difficult. The mean time was 5.7 minutes, with a standard deviation of 14.5 minutes. Figure 1 (page 26) shows the histogram of defeat times for the 91 attacks. (Two attacks are off-scale at 45 and 125 minutes.) The defeat time is the total time required to open the seal, reseal it or counterfeit it, and then cover up any evidence of entry at the appropriate level of attack (type 1, 2 or 3).

Figure 2 (page 27) shows little correlation between the defeat time and the unit cost of the seal. The linear correlation coefficient is only  $r = 0.25$ . Figure 3 (page 27) shows that there is also little correlation ( $r = 0.10$ ) between the defeat time and type of defeat (1, 2 or 3).

In Figure 4 (page 28), we see a histogram of the time required to initially develop the successful attacks. This is the time needed to devise the attack, gather up materials, make any specialized tools that were needed and demonstrate the attack for the first time. This first demonstration might not be fully successful; it sometimes took two to 20 times longer to become proficient at the attack. In all cases, however, the attacks could be developed relatively quickly. The mean time to develop an attack for the 91 attacks was 3.9 hours. There is little correlation between the time to develop an attack and other parameters, such as the cost of the seal (Figure 5 (page 28),  $r = 0.61$ ), the defeat time (Figure 6 (page 29),  $r = 0.37$ ), or the type of defeat (Figure 7 (page 29),  $r = 0.01$ ).

## Caveats

Ideally, vulnerability studies should evaluate seals in the specific, real-world context in which they are used. For about 11 percent of the seals we studied, we developed attacks in terms of an actual application. For most of the seals, however, we investigated the vulnerability in a generic sense, without one specific application in mind.

Another potential problem with this work is the classification of the attacks. Classifying an attack as successful and of what type (1, 2 or 3) is, for many of the seals, primarily our own estimation. Out of the 91 attacks we developed, only 43 were discussed with independent seal, security or nuclear experts, usually outside Los Alamos National Laboratory. An additional 13 were demonstrated to them. In each case, the experts agreed with our assessment that the attack was successful and our categorization (type 1, 2 or 3).

For only three of the 91 attacks did we do a rigorous double blind test. We had security personnel familiar with the seal try to determine which samples had been defeated. We did a blind test on three additional attacks. In these six cases, the security personnel were unable to detect which seals had been defeated, at the appropriate level of inspection (type 1, 2 or 3). (In a double blind test, the seals are independently coded so that neither the experimenter nor the test subjects are aware of which seals have been defeated until after the test is completed. In a blind test, only the experimenters are aware of the which seals have been defeated.)

The reasons for so few rigorous blind and double blind evaluations of our attacks include limitations on time and funds available for such tests, limited availability (and often a surprising lack of interest) of security personnel, and uncertainties about the context and real-world applications for the seals. Ideally, double blind tests of vulnerability should be conducted on security personnel unaware that a test is taking place. To ask security personnel which seal has been defeated is not a realistic way to evaluate real-world vulnerability. Adversaries do not usually announce to security personnel that they have defeated some of their seals. Tests on unaware security personnel, however, tend to be expensive,

time-consuming and difficult to arrange.

In analyzing this work, it is also appropriate to bear in mind that classifying an attack as type 3 is problematic. It is difficult to prove a negative — that no technology can detect the tampering. We are unable to envision any method of detecting our type-3 attacks, but that does not guarantee that such a method does not or will not exist.

### Concluding remarks

We believe this is the most comprehensive vulnerability assessment of passive seals ever undertaken. The major finding of this work is disturbing: All the tamper-indicating seals we examined can be defeated quickly, using low-tech methods available to almost anyone. Many of these seals are widely used for critical applications, including nuclear applications. The Department of Energy recognizes the vulnerability of seals and considers their safeguards effectiveness to be minimal unless they are combined with other containment/surveillance measures as part of an integrated system.<sup>6</sup>

For most of our attacks, minor modifications to the seal would substantially increase the difficulty of an attack. These modifications would usually add little to the cost. Most seals would also benefit significantly from changes in the manufacturer's suggested protocol for use and inspection. Most of the changes we would suggest are relatively minor. For many seals, we believe having security personnel aware of the most likely attack scenarios, and having them watch for these attacks, would dramatically improve tamper detection. Seal users and manufacturers with a legitimate interest in vulnerability issues are welcome to contact us to discuss specifics.

Finally, we were surprised to discover that neither the seal defeat time, nor time to develop an attack, are strong functions of unit cost (Figures 2 and 5) or the type of defeat (Figures 3 and 7). Prior to this study, we anticipated that the most costly seals would be the most effective, and that type-3 attacks would be most difficult.

### Acknowledgments

This work was performed under the auspices of the U.S. Department of Energy. We benefited from the input of Chuck Mansfield, Padmini Sockappa and James Jones.

### References

1. Rosett J.L. *Improving Tamper-Evident Packaging*. Lancaster, PA: Technomic Publishing, 1992, pp. 1-34.
2. Department of Transportation Physical Security Manual, GPO Report #982-2-5, Nov. 29, 1977.
3. Fuller J., et al, DOE's Tags and Seals Program, *Verification Technologies*, DOE Report DOE/DP/OAC/VT-92B, October 1992, pp. 4-41.
4. Tsipis K., et al., *Arms Control Verification: The technologies that make it possible*. New York: Pergamon-Brassy, 1986.
5. Gritton D. and Riley M., "Electronic Identification Devices," *Energy & Technology Review*, UCRL-52000-92, January 1992.
6. "Safeguards Seal Reference Manual," December 1986, DOE/DP-0035 (issued by the U.S. DOE, Assistant Secretary for Defense Program, Office of Safeguards and Security.)

Figure 1: Histogram for the demonstrated time to defeat a seal (with practice) for our 91 attacks. Two attacks are off the scale at 45 and 125 minutes.

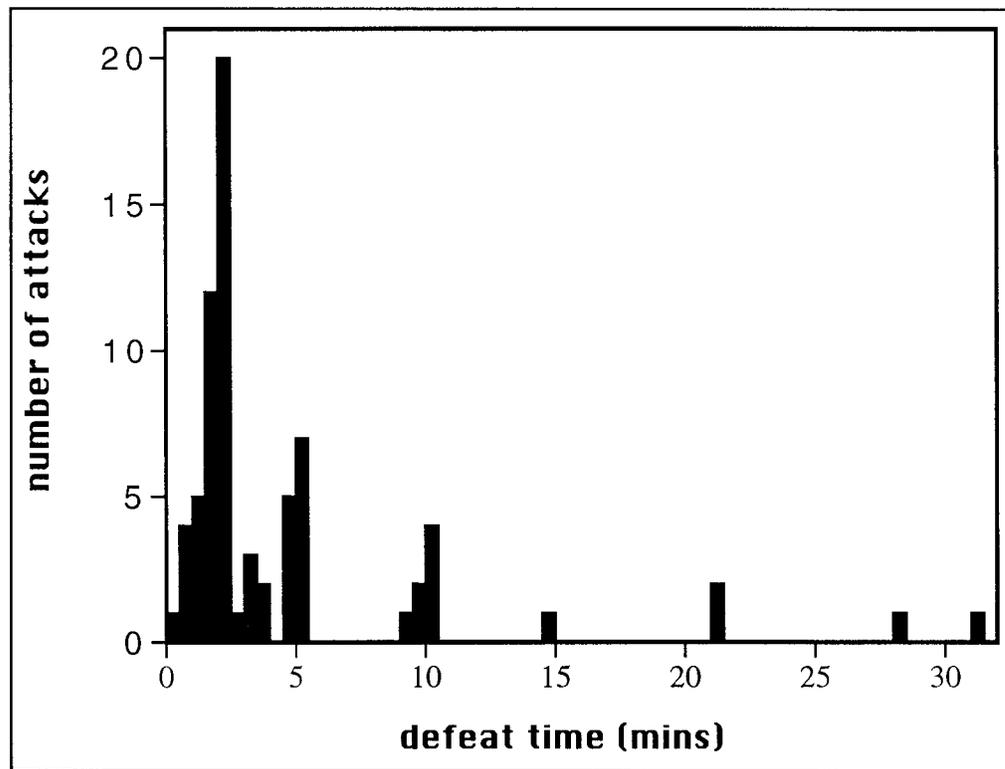


Figure 2: Seal defeat times vs. unit cost in quantities of 1,000. Each point corresponds to one attack. Data are plotted on a log-log graph. Note the lack of correlation, i.e., defeat times are not a strong function of seal cost.

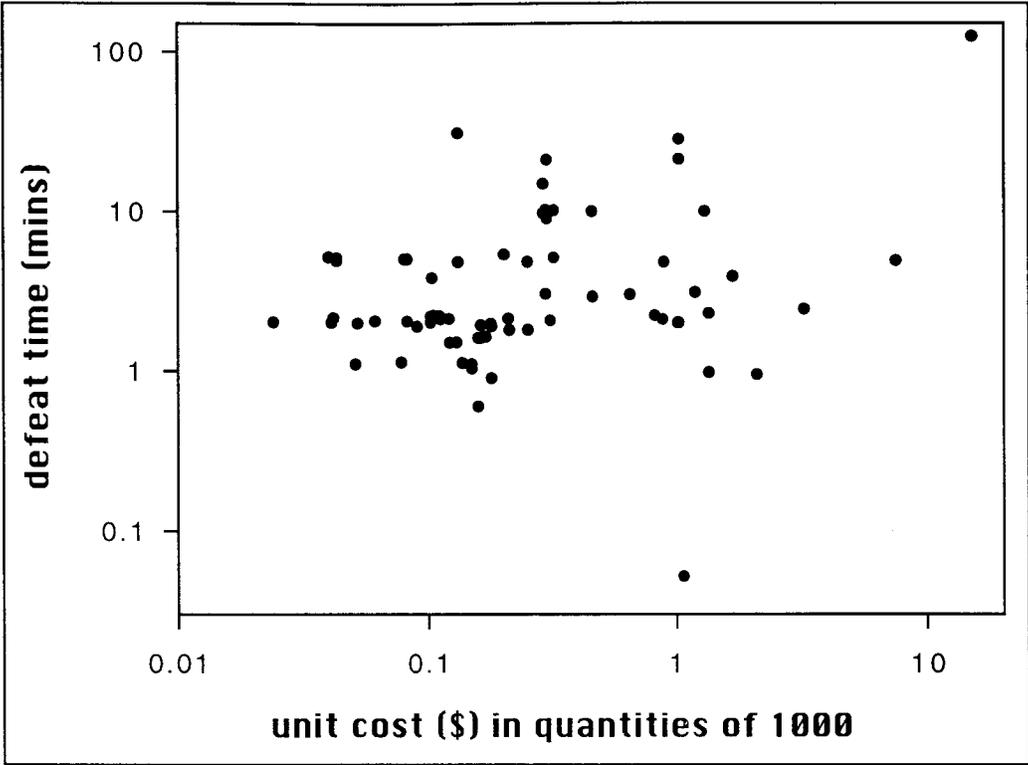


Figure 3: Defeat times vs. type of defeat (1, 2 or 3). One type-1 attack is off the scale at 125 minutes, and one type-2 attack is off the scale at 45 minutes. Note the lack of correlation.

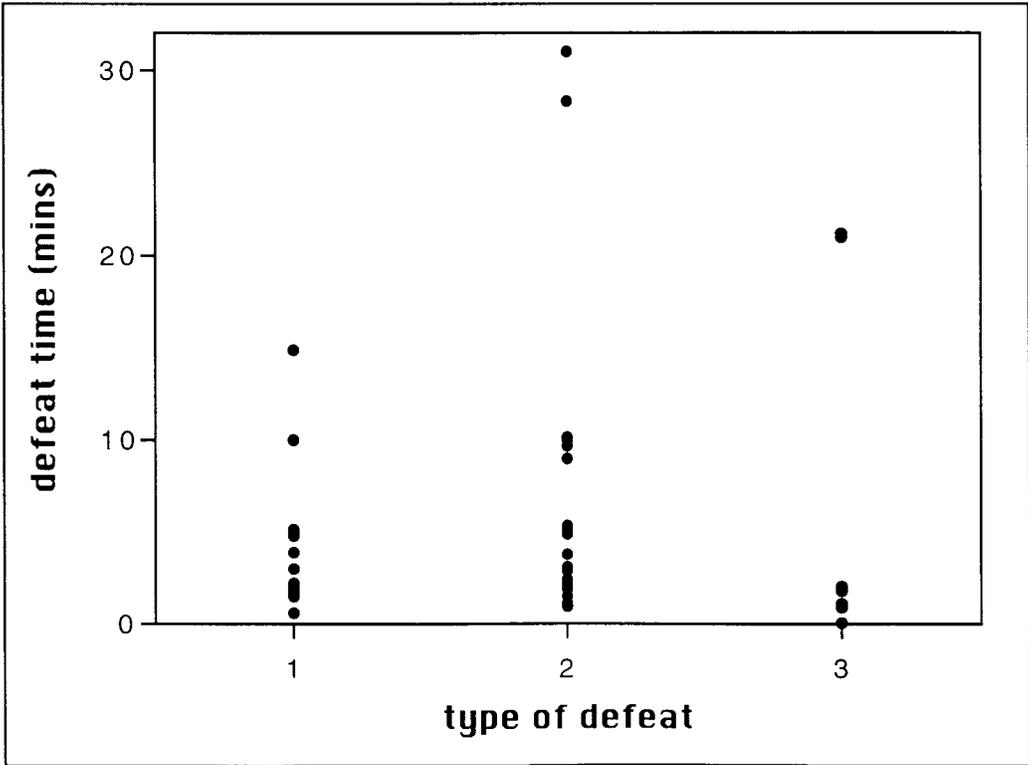




Figure 6: Time to initially develop at attack vs. the time to successfully complete the attack (with practice). This is a log-log plot. Note the relatively weak correlation.

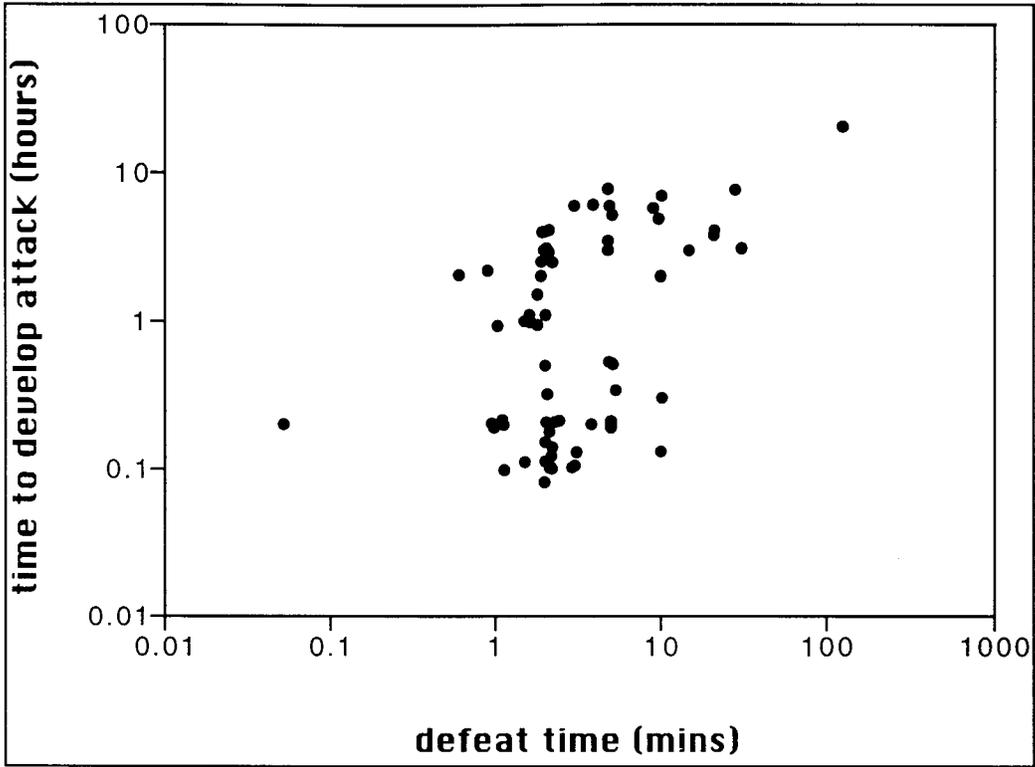


Figure 7: Time to initially develop an attack vs. the type of defeat (1, 2 or 3). Two attacks are off the scale at 20.5 hours and 240 hours. Note that there is little correlation between the development time and the type of defeat.

