LAUR-02-0492 *Journal of Homeland Security, April 2002*
http://www.homelandsecurity.org/journal/Articles/displayarticle.asp?article=50

Efficacy of Tamper-Indicating Devices

Roger G. Johnston, Ph.D., CPP
Anthony R.E. Garcia
Adam N. Pacheco

Vulnerability Assessment Team
Los Alamos National Laboratory
MS J565, Los Alamos, NM 87545
phone: 505-667-7414
fax: 505-665-4631
email: rogerj@lanl.gov

Abstract:

Tamper-indicating devices ("seals") have many important security applications including counter-terrorism, cargo security, law enforcement, nuclear safeguards, and protecting against product tampering. We studied 198 different seals and demonstrated how all can be defeated quickly using low-tech methods available to almost anyone. The seals ranged from inexpensive low-tech seals through expensive high-tech seals. Many are currently in use for critical applications. A total of 289 different attacks, at least one per seal, were devised and demonstrated. In many cases, simple changes to a seal and/or to how it is used can dramatically improve its effectiveness. Unfortunately, these changes are almost never implemented.

Introduction

Tamper-indicating devices (TIDs) are meant to detect unauthorized access, entry, or tampering [1-4]. TIDs are widely used for a variety of government and private sector applications. These include access control, cargo security, inventory control, banking, courier services, document and records integrity, customs, law and drug enforcement, hazardous materials accountability, nuclear safeguards & nonproliferation, treaty monitoring, counterespionage, counterterrorism, computer physical security, and preventing utility theft [2-6]. TIDs are also commonly used to protect food, drink, and drugs from tampering [7, 8]. They can help guarantee instrument calibration and the sterility of medical supplies, plus assist in maintaining a chain of custody for forensics and law enforcement evidence.

TIDs are frequently called "tamper-indicating seals", "security seals", or just "seals". They take a variety of forms. Seals can be frangible foils or films, plastic wraps, pressure-sensitive adhesive tapes, crimped cables or other (theoretically) irreversible mechanical assemblies; security containers or enclosures that give evidence of being opened; devices or materials that display irreversible damage or changes when manipulated; and electronic or electrooptic devices and systems that continuously monitor for changes, such as a break in an electrical cable or fiber-optic bundle.

Seals differ from locks in that they are intended to leave unambiguous, nonerasable evidence of unauthorized access, rather than impeding or delaying access. Also unlike locks, seals must be inspected, either manually or electronically, to do their job. Seals differ from intrusion ("burglar") alarms in that unauthorized access or entry is not reported immediately. This has both advantages and disadvantages [9].

The Vulnerability Assessment Team at Los Alamos National Laboratory (LANL) has previously reported briefly on a broad vulnerability assessment of various seals [10]. This paper presents the results of a more thorough vulnerability assessment and analysis of over 100 additional seals, plus new results for the seals discussed previously.

We define terminology in the next section, then characterize the seals that were studied. Results of this study are presented next, along with discussion and comments. This is followed by a review of the significant limitations of this study. We conclude with a brief discussion of the some of the serious problems with tamper detection in general.

Terminology

We define a "passive" seal to be a seal that is never electrified with either AC or DC electricity. An "active" seal, in contrast, is electrified at least at some point in its use cycle. (A seal can be read or inspected by an electronic "reader" or "verifier" and still be considered "passive" if the seal itself is never electrified.)

Like all security devices (and security programs), seals have vulnerabilities [11, 12]. We define a seal vulnerability to be a weakness or problem with the seal that can be exploited by an adversary to achieve surreptitious (undetected) access, entry, tampering, or theft. One of the critical factors for seals is the so-called "use protocols". The seal use protocols are the procedures for seal procurement, shipping, storage, checkout, installation, inspection, removal, record keeping, interpretation, disposal, and personnel training. Use protocols are important because a seal is no better than the protocols for using it.

To "defeat" a seal means to open the seal, then reseal using either the original seal or a counterfeit, without being detected. Not being detected is critical. Simply removing a seal from a container or object is not a "defeat". Indeed, many seals are made of paper or plastic and can be easily torn off by hand. This does not necessarily make them ineffective as tamper-indicating devices. The fact that they are damaged or missing provides the evidence of unauthorized access. We also often talk about "attacking" a seal. This means undertaking a sequence of actions intended to try to defeat the seal. A successful attack is also called a "defeat".

The emphasis in this work is on low-tech attacks--primarily because high-tech attacks don't seem to be necessary, even for high-tech seals. We define a "low-tech attack" as one that uses only relatively low cost methods, tools, and supplies readily available to the general public. An attack can still be considered low-tech even if, to be successful, it requires considerable practice and/or manual dexterity at the level of an average artist or craftsman.

Because seal defeats are a matter of degree, we have found it useful to categorize the thoroughness of a defeat. Towards this end, we have developed the Los Alamos Seal Defeat Categorization Scheme [13]. We classify defeats at type 1, 2a, 2b, or 3, based on the thoroughness of the seal inspection procedure that still gets fooled by the seal attack.

In a type 1 defeat, tampering is not detected if the usual or nominal seal inspection process is followed. The usual process is that routinely or typically employed by the end-user. For many seals, this is the use protocol recommended by the developer or manufacturer of the seal. For the 23% of the seals in this study that were analyzed in terms of a specific application, we defeated the actual inspection process used for that application. A type 1 defeat, however, will be detected if unusual efforts are taken. For many seals, an example of an unusual inspection protocol would be to disassemble the seal and examined it in great detail to look for tampering or counterfeiting.

In a type 2a defeat, tampering is not detected if the usual inspection protocol is followed AND if the user visually studies the exterior of the seal (plus any internal parts that can be seen without opening the seal) in great detail to look for evidence of an attack. The visual inspection can be done with either the naked eye or a hand-held magnifier.

In a type 2b defeat, tampering is not detected if the usual inspection protocol is followed AND if the user disassembles the seal and meticulously examines the interior and the exterior of the seal visually (with the naked eye or a hand-held magnifier) to look for evidence of an attack.

In a type 3 defeat (the most thorough), tampering cannot be detected, even if the most advanced postmortem analysis is undertaken. State-of-the-art techniques in forensics, material science, or microscopy will not be able to tell that the seal has been defeated. This designation is problematic because it is not possible to prove there are no techniques (now or in the future) that can detect the attack. Nevertheless, this study resulted in seal defeats that we have rated as type 3 because we do not know how to detect the attacks, even in principle.

The Seals Studied

The 198 different seals used in this study include both commercial and government seals, passive (192 seals) as well as active (6 seals). They range from very inexpensive low-tech seals, through expensive high-tech devices. Most of these seals are in widespread use. At least 56% of the seals are currently in use for applications that can reasonably be considered “critical” or “high-security”. To our knowledge, at least 16% of these seals are currently in use for nuclear safeguards applications somewhere in the world.

We--somewhat subjectively--judged 4% of the seals to be “high-tech”, i.e., utilizing advanced technologies or materials. 19% were considered “medium-tech”, with the remainder (77%) being judged “low-tech”.

We also briefly studied the design of 4 additional active seals. We identified simple low-tech attacks for these seals that look very promising. Because we have not yet attempted to demonstrate the potential defeats, however, they are not included in our results.

Results and Discussion

We demonstrated 289 different defeats for the 198 seals--at least one defeat per seal. (We devised and demonstrated as many as 6 different defeats for some seals.) Table 1 summarizes the average results. All 289 attacks were demonstrated by a single, well-practiced individual, executing the attack alone and using only low-tech methods, tools, and supplies. The fact that the average defeat time is under 4 minutes is significant given that many seal applications involve leaving the seal unattended for weeks to months, or even years [14].

Table 1 - Results for the 289 seal attacks (on 198 different seals) demonstrated in this study.

	mean (average)	median (midpoint)	range
defeat time:	3.9 mins	1.4 mins	3 secs - 2 hrs
cost of attack tools & supplies:	$126	$8	2¢ - $3000
marginal cost of tools & supplies:	40¢	10¢	1¢ - $40
time to devise attack:	6.1 hrs	36 mins	2 secs - 10 days

In Table 1, the “cost of attack tools & supplies” does not include the cost of labor to devise, practice, or execute the attack. The “marginal cost of tools & supplies” is what it costs in tools and supplies to attack a second seal of the same design. Because attack tools and supplies can often be re-used, the marginal cost is quite low. The “time to devise the attack” is how long it took us to think up an attack that ultimately proved successful. The time to become highly proficient at an attack, however, was typically longer by a factor of 1 to 20 than the time to devise the attack, depending on the seal and the attack.

Table 2 shows the results for the fastest attack on each of the seals. Again, the average defeat time and attack costs are low. The average time to devise a successful attack is also low. Figure 1 similarly demonstrates that the defeat times are quite modest. It shows the percent of the 198 seals that can be defeated in less than a given amount of time by a lone individual. For some attacks, an assistant or assistants could speed up the attack. For other attacks, an assistant would just get in the way.

Table 2 - Results for the **fastest** attack on each of 198 different seals.

	mean (average)	median (midpoint)	range
defeat time:	2.9 mins	1.1 mins	3 secs - 45 mins
cost of tools & supplies:	$128	$5	2¢ - $3000
marginal cost of tools & supplies:	43¢	9¢	1¢ - $40
time to devise attack:	5.1 hrs	12 mins	2 secs - 10 days

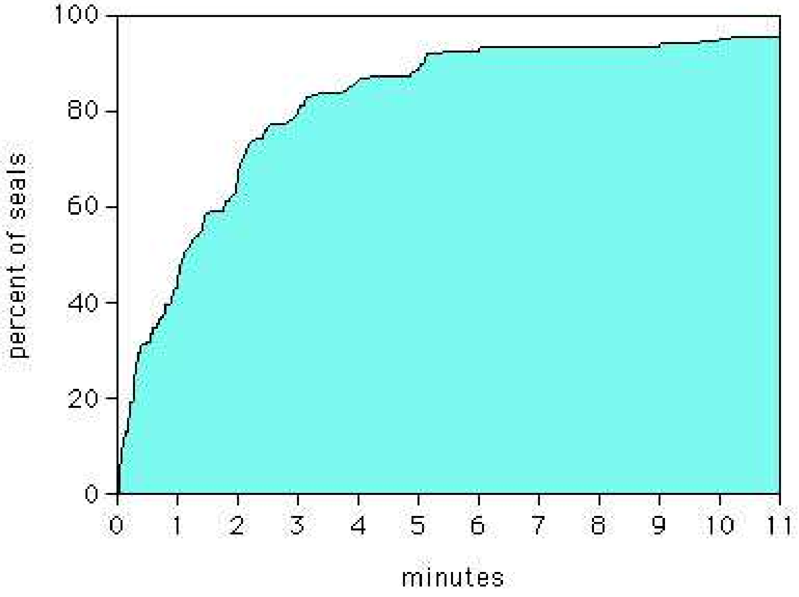

Figure 1 - Percent of the 198 seals that can be defeated in less than a given amount of time by one person executing a low-tech attack.

Throughout this study, we found that high-tech seals are not automatically superior to low-tech seals. Indeed, some of the high-tech seals were considerably easier to defeat than many of the low-tech seals--at least the way that the high-tech seals are currently used. “Easier” in this context can mean a faster defeat time, lower attack cost, lower marginal attack cost, and/or shorter time to devise the attack. Table 3 shows that the ease of defeat (defined a number of different ways) is very weakly correlated with either our subjective judgment of a seal’s level of high-technology, or with its cost. Cost should presumably have some correlation with the sophistication of the technology employed by a seal.

Table 3 - There is a very weak correlation between how easy it is to defeat a seal and its sophistication.

	r	average slope
defeat time vs. cost of seal:	0.12	1.5 secs/dollar
defeat time vs. seal tech level:	0.12	2 mins/tech level
cost of attack tools/supplies vs. seal cost:	0.03	27¢/dollar
cost of attack tools/supplies vs. seal tech level:	0.08	$50/tech level
marginal attack cost vs. seal cost:	0.08	0.4¢/dollar
marginal attack cost vs. tech level:	0.13	54¢/tech level
time to devise attack vs. seal cost:	0.22	6.5 mins/dollar
time to devise attack vs. tech level:	0.07	2.5 hours/tech level

The value r in Table 3 is Pearson's linear correlation coefficient, and measures the degree of correlation between the two parameters listed on each line [15]. The correlation coefficient takes on values between r=-1.00 and r=+1.00. A value of r=+1.00 means perfect correlation, while r=0.00 means no correlation. (A value of r=-1.00 would mean perfect anti-correlation.) As can be seen, the r values are very low, indicating remarkably little correlation between the ease of defeat and the degree of sophistication of the seals.

The costs of the seals used for Table 3 were only a rough estimate for some of the government seals that have not been commercially produced in large quantities. Specifying the cost of many of the cheaper commercial seals is also challenging because they are sold in large quantities as commodities, and the price can fluctuate significantly. The (subjective) tech level of each seal was scored as by us 1, 2, or 3, where 1=low-tech, 2=medium-tech, and 3=high-tech.

While the low correlation coefficients shown in Table 3 are the main story, some of the slopes are also quite interesting. These indicate how much one quantity changes (on average) with the other. For example, spending an extra dollar per seal--which is a lot in the commercial world--only adds 1.5 seconds to the defeat time according to Table 3. Similarly, spending an extra dollar on a seal only increases the attacker's tools and supplies cost by 27¢, and the marginal cost by less than a penny. This suggests you cannot outspend an adversary.

Similarly, going up one level in degree of technology from 1 (low-tech) to 2 (medium tech), or from 2 (medium-tech) to 3 (high-tech) adds, on average, only 2 minutes to the defeat time, $50 to the cost of attack, 54¢ to the marginal cost of the attack, and only 2.5 hours to the time to devise the attack. High-technology does not, therefore, seem to be a silver bullet for tamper detection.

Figure 2 shows in more detail the weak correlation between defeat time and seal unit cost for the 289 attacks developed in this study.

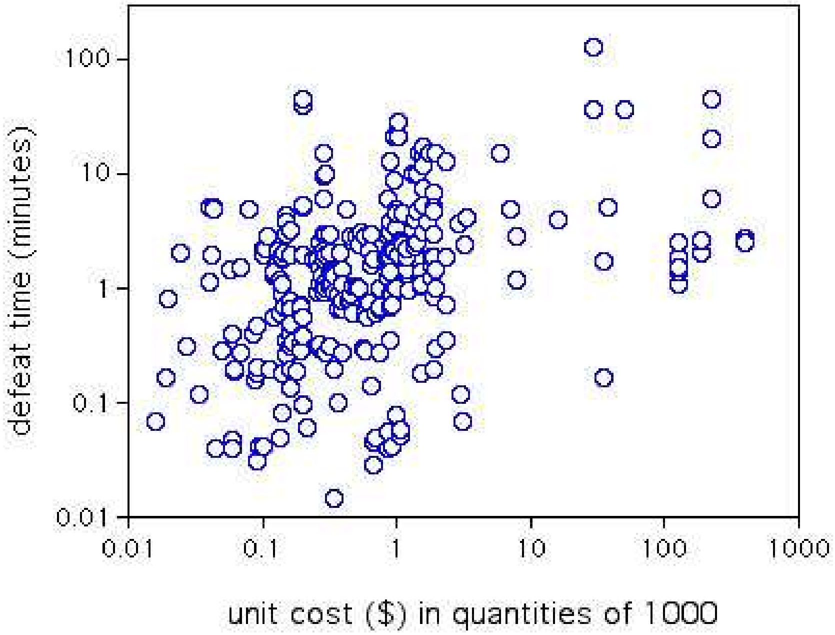

Figure 2 - Defeat time vs. seal unit cost for the 289 attacks developed in this study.

The fact that the seal cost does not correlate well with seal effectiveness is perhaps surprising. Many seal users, vendors, and manufacturers have long maintained, at least informally, that the seal user gets what he or she pays for. It may also be unexpected that high-tech seals are not harder to defeat than many low-tech seals. (This might not necessarily be the case if the high-tech seals were used differently and/or re-designed.) But why should high-tech seals be so vulnerable? We speculate that the reasons might include:

- High-tech seals must still be coupled physically to the real world, which makes the seals susceptible to simple physical attacks.

- Developers of high-tech seals tend to focus more on high-tech electronic attacks than on low-tech physical attacks.

- Developers of high-tech seals may be experts at electronics, electrooptics, microprocessors, encryption, etc. but may lack expertise in real-world tamper detection issues and physical attacks.

- There tends to be more legs for an adversary to attack in a high-tech device because of the greater complexity.

- The “Titanic Effect” may also be a factor. This is over-confidence in, or arrogance about, high-technology that may inhibit the implementation of effective design features and use protocols.

- Low-tech seals typically require hands-on inspection and handling. This forces the seal inspector to pay close attention to seal details. This is not necessarily the case with high-tech seals, or seals inspected with high-tech readers (verifiers). In our experience, if the reader is happy, the inspector is happy, even when the seal has been crudely attacked.

- Inspectors for high-tech seals may not fully understand the devices and will tend to mindlessly follow use instructions, rather than paying attention to the overall scene. An adversary can exploit this fact.

In addition to being low-tech, the tools and supplies needed to implement our 289 attacks occupy a small volume: only 1.2 liters on average per attack. For many of the attacks, the necessary tools and supplies fit in the palm of one hand. This small size should make it easier for an adversary to surreptitiously execute an attack. It would, in contrast, be harder for an adversary to avoid notice if his attack required truckloads of tools and supplies. We probably could decrease the average volume by at least a factor of 2 if we designed more custom attack tools, rather than using mostly commercial, off-the-shelf items as was done in this work. This would, however, increase the cost of the attacks.

Figure 3 charts the thoroughness of the 289 attacks, i.e., how many were type 1, 2a, 2b, or 3 defeats. A significant number (15%) were type 3, the most comprehensive and difficult to detect. Interestingly, the type 3 attacks take an average of 1.5 minutes *less* time to complete than the type 1 attacks, which are much less comprehensive.

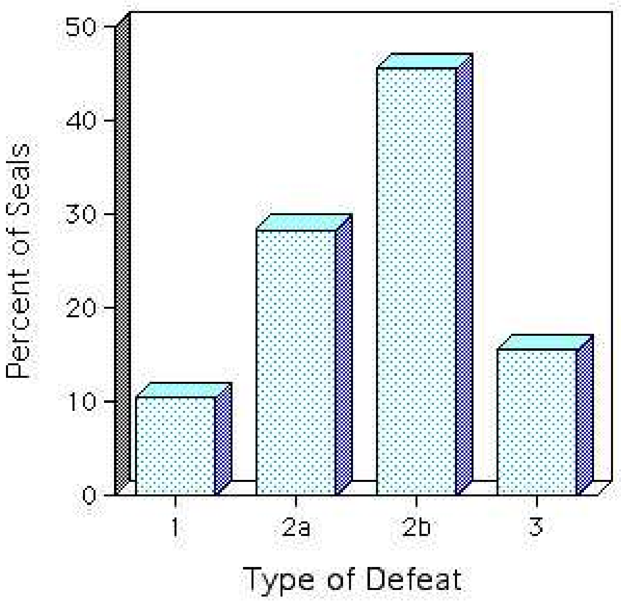

Figure 3 - Percent of the 289 defeats that are type 1 (least thorough) through type 3 (most thorough).

It is encouraging to note that many of the attacks developed in this study appear to have effective countermeasures. Unfortunately, these countermeasures are rarely implemented by seal users. About 58% of our 289 attacks have simple and inexpensive countermeasures, while 30% have countermeasures that we judge to be more involved and/or expensive, but probably still quite practical for many applications. Only 12% of the 289 total attacks have no apparent practical countermeasures. (Even some of the type 3 defeats have reasonable countermeasures involving a modification to the seal.)

Caveats

There are a number of serious problems and limitations with this work. We have not, for example, considered attacks employing "social engineering", i.e., recruiting, bribing, coercing, or compromising security personnel involved with seals. Such attacks, however, are often quite attractive [16]. Indeed, no existing seals appear to have been designed with the idea that the manufacturer, vendor, installer, inspector, or custodian of the seal data may have a hidden agenda, though such seals are possible [17, 18].

We also did not consider attacks that involve using the original seal manufacturer (knowingly or unknowingly, willingly or unwilling) to make a replicate seal. This kind of attack is probably relatively easy to implement, especially for low-cost seals, and may have a high probability of success [19]. Other types of counterfeiting attacks are probably also quite straightforward, but were not, for the most part, considered in this study. This is because re-using the original seal was typically easier, faster, and cheaper. Of the 289 attacks developed in this study, only 9% involved counterfeiting the entire seal or parts of the seal. Seal vendors and manufacturers sometimes emphasize the difficulty of counterfeiting their seals. While we suspect this is often overstated, counterfeiting is usually not the most attractive attack strategy for an adversary.

There were also no cold attacks considered in this work. A “cold attack” is one where the adversary attempts to defeat a seal on-site which he or she has never seen before. We consider such scenarios to be unrealistic because we are unaware of ANY seal currently in use, including for very high-security applications, that an adversary would need to attack cold.

A critical issue in this work is the determination of whether an attack is successful. This is hard for the reader of this paper to judge for herself given that we have not discussed specifics. Indeed, we have not explained here how the attacks are done. We believe to do so would be irresponsible, especially given the critical applications for some of these seals. We also did not identify the specific seals studied. There are two reasons for this. Firstly, since all the seals we studied appear to have readily exploitable vulnerabilities, we know from experience that identifying a given seal as a subject of this work will be taken by some as a *de facto* black mark on that seal. This is simple-minded in our view because we believe ALL seals have vulnerabilities. Secondly, we believe briefly discussing specific seals and their vulnerabilities in a short paper outside of a detailed consideration of concrete applications, facilities, use protocols, security personnel, adversaries, and an overall context would be a serious disservice. The field of tamper detection is already plagued--at least in our view--by hearsay, over simplification, misconceptions, wishful thinking, sloppy terminology, the absence of useful standards, a “one-size-fits all” attitude, and a general lack of thoughtfulness and rigor [20]. We do not wish to contribute to these problems by inadvertently appearing to offer simplistic answers. While we think our rationalizations for omitting specifics are valid, this does not solve the problem that readers will have difficulty evaluating the viability of our attacks when our results are presented solely in statistical form. As a partial solution to this problem, we invite security managers and planners with legitimate security responsibilities to contact us to discuss specific tamper detection issues within a more holistic framework.

Ideally, the success of an attack should be established using rigorous blind or double tests involving actual seal inspectors. (In a blind test, the test subject does not know which seals have been attacked and which have not. In a double blind test, the experimenter supervising the test also does not know.) In reality, however, rigorous blind and double blind tests are difficult and expensive to arrange, and nearly impossible to make realistic. The reasons are too complex to delve into here, but are partially discussed elsewhere [21-23]. In any event, most of our sponsors are unwilling to commit the time, labor, and money needed for rigorous testing of our devised attacks. Most are quite content--indeed more comfortable--being told about or shown the attacks, then being allowed to judge for themselves if the seal vulnerabilities require mitigation, and if our suggested countermeasures are practical. As Table 4 shows, we did do blind and double blind testing to establish the success of some our attacks, but most determinations were made in other, less rigorous ways, as indicated in the table.

Table 4 - Ways in which the 289 attacks were judged to be defeats.

	percent of attacks

attack discussed with seal user:	39%
attack demonstrated to seal user:	10%
samples of attacked seals shown to seal user:	6%
rigorous blind test with seal user:	5%
rigorous double blind test with seal user:	1%
rigorous double blind test with non-experts:	1%
not yet presented to outsiders:	38%

Thus, for 38% of the attacks, the idea that the attack is successful is, at least at this point, solely our own opinion, with no independent, outside confirmation. Most of these attacks are for seals that had no specific user to discuss/demonstrate/test the attack with (or show samples to), no specific application to consider, and/or no well-defined use protocol to defeat. Given our previous experience, however, we are confident that these attacks would be judged successful by independent outsiders.

For all the other attacks (62%), the outsider(s) with whom the attacks were discussed, demonstrated, tested, or samples shown to always agreed with us that the attacks were viable. For most of the attacks, the “outsiders” were security managers and/or their seal inspectors who could reasonably be considered to be knowledgeable in use of the seals in question. For 1% of the attacks, we instead blind-tested the attacks using LANL personnel who were unfamiliar with this work. This was done because the actual seal inspectors were not available to participate in the testing.

Towards Better Tamper Detection

Tamper detection is a field that is over 7000 years old [24]. Despite its antiquity, however, there is little in the way of a formal theory for tamper detection, few meaningful standards for seals or seals testing [25], and a great deal of confusion over concepts, strategies, and terminology [26]. Moreover, seals are often used very poorly, even for critical applications. We have studied in detail the seals or the complete tamper detection programs used by over one dozen government agencies and 11 private companies. We have also reviewed in lesser detail seals and use protocols employed by a number of other seals users. In our view, few seal users have chosen the most appropriate seal for their application. Fewer still understand the seals they are using, have any substantial awareness of their vulnerabilities, or provide seal installers and inspectors with the information and practice they need for effective tamper detection. Few seal users employ seal use protocols even close to optimal. This is very unfortunate because we believe dramatic improvements in tamper detection are often possible with relatively minor changes to a seal, its use protocols, and/or to the training provided to seal installers and inspectors.

Having seen the problems with existing seals, we firmly believe that better seals are possible. We've developed 15 new seals ourselves, but this work is largely unsupported. In fact, the U.S. Government has undertaken very little development of new seals since 1994. Work on new seals continues in the private sector, but much of this is focused on lowering unit cost and improving ease of use, not on security. Most seal users in the private sector seem astonishingly unconcerned about seal effectiveness.

In the case of critical government applications such as nuclear safeguards, a two-person rule is often used in conjunction with tamper-indicating seals. A "two-person rule" stipulates that no individual can be alone with the critical assets. In theory, use of a two-person rule should make it harder to attack a seal.

We certainly agree that a two-person rule is a sound security strategy. We are not, however, convinced that a two-person rule should be used as an excuse to avoid optimizing seal effectiveness. Moreover, we suspect that the supreme confidence typically placed in the two-person rule is unwarranted. There are many potential problems [27, 28]. These include the fact that there is no real standard within DoD, DOE, or the U.S. Government; the two-person rule is implemented quite differently inside various government facilities, and training varies considerably. Moreover, with the exception of certain narrow applications, the two-person rule is poorly studied, researched, and tested. Effective training and practice exercises are often lacking. Another problem for tamper detection is that personnel authorized to serve on two-person rule teams rarely know what a seal attack looks like for the seals used in their facility, nor do they know what attack tools might be needed. This, plus the fact that many of the attacks demonstrated in this work can be implemented in phases over an extended period of time, a few seconds per phase, makes it less than certain that personnel involved in a two-person rule environment will be able to automatically spot a sophisticated (or even unsophisticated) seal attack.

In summary, tamper detection can and should be more effective given its importance to security and societal well-being. Better seals, more optimal use protocols, additional research, greater awareness of seal vulnerabilities, more critical thinking, increased sophistication on the part of seal users, and better tamper detection training are sorely needed.

Acknowledgements and Disclaimer

The views presented here are those of the authors and should not necessarily be ascribed to Los Alamos National Laboratory or the United States Department of Energy. Debbie Martinez, Ron Martinez, Eric Baca, W. Kevin Grace, and Anna Nogar contributed to this work.

References

1. David L. Poli, "Security Seal Handbook," Report SAND78-0400 (Albuquerque, NM: Sandia National Laboratories, December, 1978), http://www.sandia.gov/doe-oss/DOC-Reports.html.

2. Naval Facilities Engineering Services Center (NFESC), "Antipilferage Seal User's Guide" (Port Hueneme, CA: October, 1997), http://locks.nfesc.navy.mil/Security_seals/guides/seal_ug/rev_sealguide.pdf.

3. Naval Facilities Engineering Services Center (NFESC), "DoD Training Course on Effective Seal Use" (Port Hueneme, CA: Spring, 2000), http://locks.nfesc.navy.mil/Security_seals/security_seals/sp2086.pdf.

4. Roger G. Johnston, "The Real Deal on Seals," Security Management, Vol. 41, September 1997, pp. 93-100, http://lib-www.lanl.gov/la-pubs/00418795.pdf.

5. Lou Tyska (editor), Guidelines for Cargo Security & Loss Control (Annapolis, MD: National Cargo Security Council, 1999), pp. 29-38.

6. U.S. Nuclear Regulatory Commission (NRC), "Tamper-Indicating Seals for the Protection and Control of Special Nuclear Material," Regulatory Guide 5.15, Revision 1 (Washington, D.C.: March 1997), http://www.sandia.gov/doe-oss/DOC-Standards.html#NRC.

7. Title 21 of the Code of Federal Regulations, Sec 211.132.

8. U.S. Food and Drug Administration (FDA), "Tamper-Resistant Packaging Requirements for Certain Over-the-Counter (OTC) Human Drug Products," FDA Compliance Guides, Chapter 32A, Drug Alteration Guide 7132A.17 (Washington, D.C.: May 21, 1992).

9. Roger G. Johnston, "The 'Town Crier' Approach to Monitoring," Report LAUR-01-3726 (Los Alamos, NM: Los Alamos National Laboratory, July 2001).

10. Roger G. Johnston and Anthony R.E. Garcia, "Vulnerability Assessment of Security Seals," Journal of Security Administration, Vol. 20, No. 1, June 1997, pp. 15-27, http://lib-www.lanl.gov/la-pubs/00418796.pdf.

11. James L. Jones, "Improving Tag/Seal Technologies: the vulnerability assessment component," Report 95/00599, (Idaho Falls, ID: Idaho National Engineering and Environmental Laboratory, December, 1996).

12. R.G. Johnston and A.R.E. Garcia, "An Annotated Taxonomy of Tag and Seal Vulnerabilities," Journal of Nuclear Materials Management, Vol. 28, No. 3, Spring 2000, pp. 23-30.

13. Roger G. Johnston, "Effective Vulnerability Assessment of Tamper-Indicating Seals," Journal of Testing and Evaluation, Vol. 25, July 1997, pp. 451-455, http://lib-www.lanl.gov/la-pubs/00418792.pdf.

14. Roger G. Johnston, “Tamper Detection for Safeguards and Treaty Monitoring: Fantasies, Realities, and Potentials,” The Nonproliferation Review, Vol. 8, No. 1, Spring 2001, pp. 102-115, http://lib-www.lanl.gov/la-pubs/00367047.pdf

15. William Mendenhall and Terry Sincich, A Second Course in Statistics: Regression Analysis (New Jersey: Prentice-Hall, 1996), pp. 101-171.

16. R.G. Johnston and A.R.E. Garcia, "An Annotated Taxonomy of Tag and Seal Vulnerabilities," pp. 27-29.

17. Roger G. Johnston, “Tamper Detection for Safeguards and Treaty Monitoring: Fantasies, Realities, and Potentials,” pp. 107-108.

18. Roger G. Johnston, “Tamper-Indicating Seals for Nuclear Disarmament and Hazardous Waste Management,” Science & Global Security, Vol. 9, No. 3, 2001, pp. 105-107, http://lib-www.lanl.gov/la-pubs/00818333.pdf.

19. R.G. Johnston and A.R.E. Garcia, "An Annotated Taxonomy of Tag and Seal Vulnerabilities," pp. 27.

20. Roger G. Johnston, “Tamper Detection for Safeguards and Treaty Monitoring: Fantasies, Realities, and Potentials,” pp. 102-109.

21. Roger G. Johnston, Debbie D. Martinez, and Anthony R.E. Garcia, “Were Ancient Seals Secure?,” Antiquity, Volume 75, No. 288, June 2001, pp. 302-303.

22. Roger G. Johnston, “Effective Vulnerability Assessment of Tamper-Indicating Seals,” p. 454.

23. Roger G. Johnston and Anthony R.E. Garcia, "Vulnerability Assessment of Security Seals," pp. 19-20.

24. Roger G. Johnston, Debbie D. Martinez, and Anthony R.E. Garcia, “Were Ancient Seals Secure?,” p. 299.

25. Roger G. Johnston, “Effective Vulnerability Assessment of Tamper-Indicating Seals,” pp. 451-452.

26. Roger G. Johnston, “Tamper Detection for Safeguards and Treaty Monitoring: Fantasies, Realities, and Potentials,” pp. 102-115.

27. Oleg Bukharin and Helen M. Hunt, “The Russian-U.S. HEU Agreement: Internal Safeguards to Prevent the Diversion of HEU,” Science & Global Security, Vol. 4, No. 2, 1994, pp. 199-201.

28. Roger G. Johnston, “Tamper Detection for Safeguards and Treaty Monitoring: Fantasies, Realities, and Potentials,” p. 106.

Roger G. Johnston, Ph.D., CPP heads the Vulnerability Assessment Team at Los Alamos National Laboratory (LANL). He received a B.A. degree from Carleton College in 1977, and M.S. and Ph.D. degrees in physics from the University of Colorado in 1983. Johnston holds 6 U.S. patents and has published over 60 technical papers in various fields. He served as a Science Fellow at the Center for International Security and Cooperation (CISAC) at Stanford University for 2000-2001.

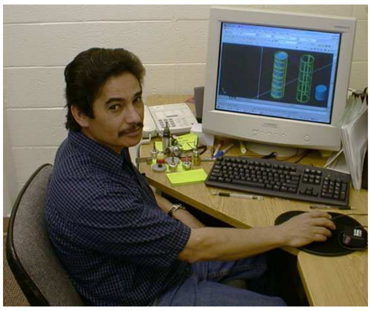

Anthony R.E. Garcia is Senior Technician with the Vulnerability Assessment Team at Los Alamos National Laboratory (LANL). Garcia has contributed to research and development projects involving national security, tamper detection, nuclear nonproliferation, superconductivity, thin film physics, and laser applications. In his free time, he is an award-winning wood carver specializing in traditional Northern New Mexican art.

Adam N. Pacheco is an Undergraduate Research Assistant with the Vulnerability Assessment Team at Los Alamos National Laboratory (LANL). He is completing an Associates Degree Program in Electro-Mechanical Technology at the University of New Mexico-Los Alamos.