Progress and Evolution of Critical Infrastructure Protection over the Last Ten Years?

by Fred Cohen, Fred Cohen & Associates

The question was posed as to whether and to what extent critical infrastructure protection has changed in the last ten years. In my opinion, this includes the question of whether such changes have improved the situation. Let us start with the underlying nature of the question.

It seems that while many changes can be readily identified, there is no basis today for determining whether or to what extent any such changes have made protection more or less effective. Indeed, the most fundamental change needed to address this issue is the movement toward a science of protection, which has not been seriously undertaken, at least in some areas. As a result of the lack of scientific study and a widely accepted scientific basis, there is no meaningful way to make sensible measurements of protection, and thus the notion of improvement cannot be realistically assessed. But perhaps more detail would be helpful in bringing clarity to this issue.

What Changes have been Made?

As a starting point, it should be noted that there have been many changes in protection associated with critical infrastructures of late. Starting in the late 1990s, when the President's Commission on Critical Infrastructure Protection (PCCIP) identified the need at a national basis to make changes in the United States, a vision of a future state was notionally present. But, at the time, that vision was not being realized at a rapid rate. Then the September 11, 2001 attacks on U.S. soil took place, leading ultimately to a dramatic change in the posture of the United States and much of the rest of the world with regard to protection.

Unfortunately, to many knowledgeable and skeptical observers, many of these changes were of little technical value other than from a perception perspective. Issues such as increased levels of invasiveness and inconvenience in airport searches were not met by published or identified measurements of effectiveness. Testing repeatedly showed that the countermeasures in place could be readily bypassed. A system of continuous "improvement" based on reducing the potential for the "last attack" were applied with little apparent effect other than to make travel less convenient, more costly, and more invasive to the traveling public with perhaps a marginal but unmeasured benefit in reducing copy-cat attacks. The feeling of safety and level of suspicion engendered by such measures has the — again unmeasured — likely effect of making individual and poorly trained threat actors more nervous and more likely to be detected and interdicted while making the public at large feel safer, with the benefits of these results quantitatively unknown and somewhat speculative.

Other less obvious changes were perhaps somewhat more technically effective, although again, measurement has not been widely published if it has been undertaken in a meaningful way at all. For example, water systems, pipelines, power systems, telecommunications, and other similar systems have been examined at some level of depth to better understand the potentials for harm and the limitations of existing protective mechanisms. While the increased scrutiny is, presumably, a benefit in terms of general and situational awareness, and the increased focus of attention and resources on these issues has the potential to increase knowledge, evidence of the utility of these presumed benefits is lacking and no system of measurement has been widely identified and applied to provide meaningful metrics.

Lacking a system of measurement and a standard against which to measure, it is doubtful we can determine, based on fact, whether protection is better or worse today than it was ten years ago.

(Continued on Page 5)

Evolution *(Cont. from 4)*

Some Things We May be able to Measure in Some Areas

By focusing on select areas of protection, additional insight may be gained. The particular areas chosen here include the information protection arena with focus on general purpose computing at enterprises, by individuals, and in government, which is vital to several critical infrastructure sectors (e.g., finance, government, telecommunications, logistics), and on industrial control systems, which are vital to more traditional critical infrastructure sectors (e.g., power, water, fuel, manufactured goods, and transportation). Within these areas, there are various issues that are commonly considered, some of which are identified here and discussed in terms of changes over the last ten years. This represents only my qualitative assessment and is without adequate basis in measurement to be treated as definitive in any meaningful way. As discussed, measurement in these areas is currently problematic.

Risk Issues:

Threats: It appears that more and more skilled, organized, and resourced threat actors have been identified and are active in 2011 than in 2001. However, to my knowledge, there is no systematic reporting of or method for identifying and characterizing such individuals and groups available in any public forum. While there are groups that produce reports and provide databases on these issues, none of them appear to be adequately supported by base rates or detailed methodologies to support sound conclusions about situations at any given time or as it changes with time.

Vulnerabilities: It appears, based on publicly available databases like the Open Source Vulnerability Database (OSVDB) and national vulnerabilities databases such as the MITRE database on vulnerabilities, that the rate of widely published vulnerabilities has increased. However, this tends to ignore many of the widely known vulnerabilities that are not directly technical in nature, such as exploitations of people, systems, architectures, and similar issues. For example, malicious programs that are downloaded and run by users represent vulnerabilities that are ignored by such collections and often dominate the events of interest. Interconnection and architectural vulnerabilities have increased for ICS systems so that it is clear that there are more paths to attach such systems. Many such systems do not have the same sort of mechanisms or protections feasible today as more general purpose Internet connected computers. This can reasonably be seen as an increase in vulnerability if measured in terms of the size and number of links from external sources to targets in the attack graph. But as a metric, it is unclear whether this is useful for understanding the overall situation. For example, insiders remain responsible for much of the reported damage attributed to sources, and thus external connectivity changes may be of relatively limited impact in terms of overall security.

Consequences: The consequences of attacks are rarely available for review as they tend to be closely guarded secrets. While the widely publicized consequences are seemingly bigger and more frequent today than ten years ago, as seen by the DataLoss.Org website and other published reports, reporting requirements have increased as has publication of events, while speculation about consequences has been lacking in the published reports. When consequences are available, the lack of a common method of valuation for information-related losses or content leads to a set of numbers that are not comparable and thus fail to support any meaningful conclusions against a standard or over time.

Interdependencies and Risk Aggregation: The analysis necessary to support definitive answers in this regard is, again, not available. But it seems almost certain that analysis would show that aggregation of risks and interdependencies have dramatically increased over the last ten years. There is little available data to support substantial conclusions from losses in this regard; however, the effects of risk aggregation and common mode failures have produced dramatic results, including the nuclear reactor problems in Japan, the WikiLeaks classified leak incident, and any number of large-scale events involving information technology (IT). What is unclear is whether the situation is in fact getting any worse since, in

(Continued on Page 6)

Evolution *(Cont. from 5)*

addition to the lack of measurement, mandatory reporting has increased only for leaks of privacy-related information and select critical infrastructure incidents. Even for reported cases, valuation methods are not available and quantification is not based on a standard. The rationale for asserting that the problem has worsened is the increased complexity and integration of systems into larger information mechanisms. For example, identity management has substantially moved ahead. Along with it, federation of identities and the use of centralized servers for control of these processes obviously places increased risk on the small number of servers leveraged to larger effect and adds interdependencies that were not previously present. However, the economies of scale and quality of such systems may compensate for the aggregations by providing higher surety for more systems. Again, without a system of measurement, no basis for meaningful metrics can be supported.

Risk Management Approaches: The four basic approaches (transfer, acceptance, avoidance, and mitigation), have not changed over the last ten years, and this field appears to be relatively static. It appears that the tolerance for risk has increased in the information arena, but this may be the result of a lack of understanding of the issues rather than a conscious or knowing decision by decision-makers.

Protection Objectives:

• Integrity protection capabilities have improved to some extent in large numbers of computers as the Trusted Computing Group's trusted platform module integrity mechanisms have been increasingly deployed over the last decade; however, these mechanisms go largely unused for anything other than limited digital rights management. Meanwhile, systems like mobile devices, which are increasingly dominating the consumer market and pushing into enterprise markets, largely lack longstanding integrity mechanisms such as hardware process separation and file systems security. Again, we have no sound method for measuring any actual change in integrity protection.

• Availability of hardware has improved with technological advancements, and software checking has improved availability along with operating system improvements for widely deployed operating environments like Windows. At the same time, server software has become more complex and interdependent, and no widely available measurements of availability have been identified that could address the overall question of availability. Furthermore, the movement to mobile devices has dramatically changed the nature of availability since access and utility has increased from a wide range of locations where access was previously very expensive and not widely available. Again, depending on what is chosen to measure, different results will appear. It is, however, obvious to most observers that accessibility to services has made a dramatic shift in the positive direction with the increased availability of WiFi and cellular services supporting smart-phones, pad computers, and netbook platforms.

• Confidentiality has been problematic in the rapid growth of Internet technologies and the push toward gaining efficiencies of scale. Many major disclosures of leaks have driven the perception of a worsening situation, but again, no metrics are really available that are comprehensive in nature. Large-scale leakage of classified information to the media associated with the WikiLeaks case appears to be unprecedented, but it is unclear how this affects critical infrastructures.

• Use control is apparently being loosened as mobility increases and information work is increasingly outsourced and physically distributed. More and longer chains of interdependencies appear to be present today than ten years ago. This leads to use control problems, increasingly seen in published attacks, that gain access to user interfaces which are then exploited to attack systems with which the users have access. But again, metrics in this area are lacking and the lack of detection does not translate into a lack of actual attack. More detections may mean fewer attacks persist.

• Accountability is increasing in many systems and

(Continued on Page 28)

Transportation *(Cont. from 16)*

constrained within the borders of the United States and thus must deal with the complexity of global regulations and policy. The TSSSP seeks to manage risk in collaboration with industry. Both also seek to grow transportation systems. ❖

Evolution *(Cont. from 6)*

infrastructures as logging requirements are being added and reporting requirements increase. Laws like Sarbanes-Oxley Act and the Gramm-Leech-Bliley Act are reflected globally, and the financial crisis has increased the push toward greater accountability. But in truth, there was little accountability ten years ago and little accountability today. More or less, the lowest level actors involved in any act can be counted on to be punished for top-level acts of malice or incompetence.

Protective Technologies and Approaches:

In essence, there are no widespread new technologies for information protection that have been deployed in the last ten years. Identity management has increased market penetration; data loss prevention methods similar to previous methods for intrusion detection have become more widespread; and trusted platform modules are now widespread. But the overall effect of these changes is apparently negligible and no metrics exist to accurately measure any such effects. Furthermore, the technologies supporting ICS systems have been essentially unchanged over this period, despite dramatic increases in connectivity. This does not bode well for protection of these elements of critical infrastructures. We see more spending on protection in computers, which may reflect increased risk, but again, without metrics...

Conclusions

The one thing we can say with a fairly definitive conclusion with regard to the questions at hand, is that the science and measurement associated with critical infrastructure protection, at least in the information protection arena, has not yet reached the point where even the most basic questions about whether protection is better or worse can be meaningfully evaluated. Without developing a science and system of measurement, we will not be able to answer these sorts of questions ten years from now either. Unfortunately, the main area where we need to make progress in order to make progress in all other areas is an area where we have made no progress.

And time is not on our side. ❖