7. Geosynchronous Orbital Ion Cannon source code. Accessed Jan 2011. <http://pastebin.com/FH6njMew>.
8. Moyer, E. 'Report: FBI seizes server in probe of Wikileaks attacks'. CNET, 1 Jan 2011. Accessed Jan 2011. <http://news.cnet.com/8301-13578_3-20026908-38.html>.
9. 'Jester's Court'. <http://th3j35t3r.wordpress.com/>.
10. Leyden, J. 'Anonymous hacktivists fire ion cannons at Zimbabwe'. The Register, 31 Dec 2010. Accessed Jan 2011. <http://www.theregister.co.uk/2010/12/31/anon_hits_zimbabwe_sites/>.
11. Cluley, G. 'Pro-Wikileaks hackers bring down Tunisian government websites'. Naked Security blog, Sophos, 3 Jan 2011. Accessed Jan 2011. <http://nakedsecurity.sophos.com/2011/01/03/pro-wikileaks-hackers-tunisian-government-web-sites>.

Cyber attacks: awareness

Edward G Amoroso, AT&T

In this excerpt from his book, Cyber Attacks: Protecting National Infrastructure, cyber-security expert Edward G Amoroso looks at how you detect infrastructure attacks, manage vulnerability information and manage risk.

Real-time understanding

'Situational awareness' refers to the collective real-time understanding within an organisation of its security risk posture. Security risk measures the likelihood that an attack might produce significant consequences to some set of locally valued assets. A major challenge is that the factors affecting security risk are often not locally controlled and are often deliberately obscured by an adversary. To optimise situation awareness, considerable time, effort and even creativity must be expended.

Sadly, most existing companies and agencies with responsibility for national infrastructure have little or no discipline in this area. This is surprising, as a common question asked by senior leadership is whether the organisation is experiencing a security risk or is 'under attack' at a given time.

Awareness of security posture requires consideration of several technical, operational, business and external or global factors. These include the following:

- **Known vulnerabilities**: detailed knowledge of relevant vulnerabilities from vendors, service providers, government, academia and the hacking community is essential to effective situational awareness. Specific events such as prominent hacking conferences are often a rich source of new vulnerability data.
- **Security infrastructure:** understanding the state of all active security components in the local environment is required for proper situational awareness. This includes knowledge of security software versions for integrity management and anti-malware processing, signature deployments for security devices such as intrusion detection systems, and monitoring status for any types of security collection and processing systems.
- **Network and computing architecture**: knowledge of network and computing architecture is also important to understanding an organisation's situational security posture. An accurate catalogue of all inbound and outbound services through external gateways is particularly important during an incident that might be exploiting specific ports or protocols.
- **Business environment**: security posture is directly related to business activities such as new product launches, new project initiation, public relations press releases, executive action involving anything even mildly controversial, and especially any business failures. Any types of contract negotiations between management and employee bases have a direct impact on the local situational security status.
- **Global threats**: any political or global threats that might be present at a given time will certainly have an impact on an organisation's situational security posture. This must be monitored carefully in regions where an organisation might have created a partnership or outsourcing arrangement. Because outsourcing tends to occur in regions that are remote to the organisation, a global threat posture has become more significant.
- **Hardware and software profiles**: an accurate view of all hardware and software currently in place in the organisation is also essential to situational awareness. A common problem involves running some product version that is too old to properly secure through a programme of patching or security enhancement. A corresponding problem involves systems that are too new to properly characterise their robustness against attack. In practice, an optimal period of product operation emerges between the earliest installation period, when a product or system is brand new, and the latter stages of deployment, when formal support from a vendor might have lapsed (see Figure 1).

Each of these factors presents a set of unique challenges for security teams. An emerging global conflict, for example, will probably have nothing to do with the vulnerability profile of software running locally in an enterprise. There are, however, clear dependencies that arise between factors in practice

and will improve situational awareness. For example, when vulnerabilities are reported by a hacking group, the organisation's security posture will depend on its local hardware, software and security infrastructure profile. As a result, it is generally reasonable for an organisation to combine the value of all situational status factors into one generic measure of its security posture. This measure should be able to provide a rough estimate of the broad organisational security risk at a given time. It should then weigh the likelihood and potential consequences of serious attack against the normal, everyday level of risk that an organisation lives with every day. Presumably, risk on a day-to-day basis should be lower than during a serious incident, so it stands to reason that a rough metric could capture this status, perhaps as a high, medium and low risk characterisation (see Figure 2).

Unfortunately, the public perception of categorising high, medium and low security risks is that it does not provide useful information. This is certainly true for such measures as the public threat metric, which was used by the US Department of Homeland Security to characterise risk. The problem with this metric was that it dictated no concrete actions to be taken by citizens. If risk was characterised as low, citizens were warned to remain vigilant and on guard; if risk was characterised as medium or even high, the advice was essentially the same. Citizens were told to go on with their normal lives, but to be *somehow* more careful. Obviously, this type of advice causes confusion and is to be avoided in national infrastructure protection. The only way a posture metric can be useful is if it is driven by real-time events and is connected directly to an explicit incident response programme. When this is done, an ongoing rhythm develops where the situational status helps direct security management activity. This could involve some serious flaw being detected in an organisation (which would drive the threat level upward), followed by detection of a real exploit in the wild (which would drive the threat level further upward), followed by a patch activity that fixes the problem (which would drive the threat level back down – see Figure 3).

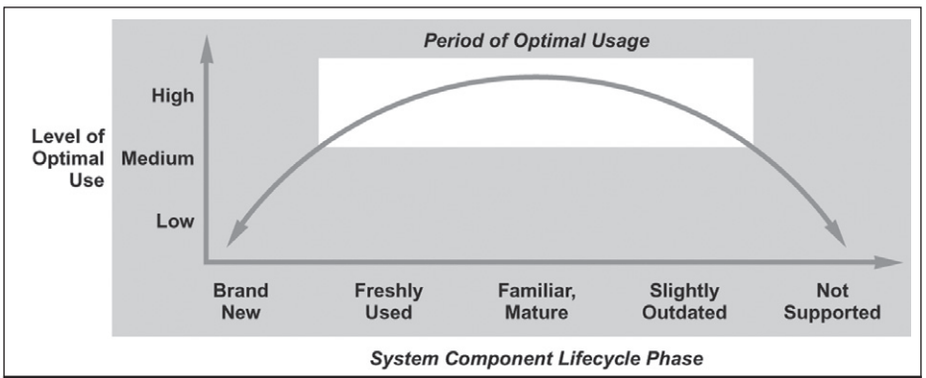

Figure 1: Optimal period of system usage for cyber-security.

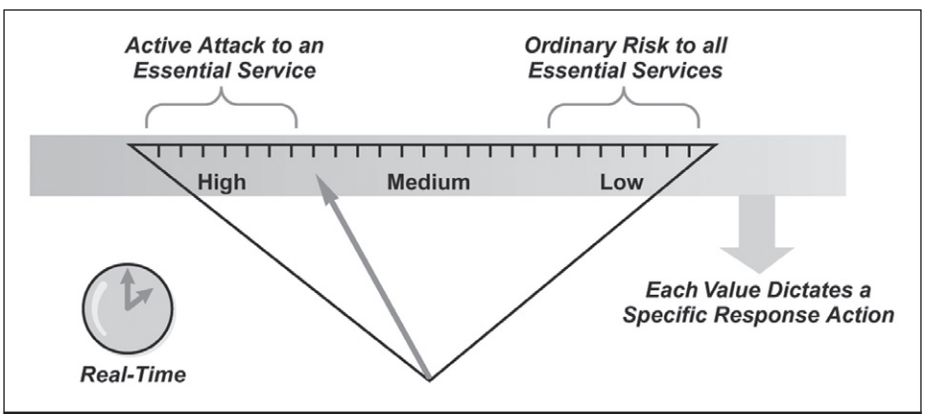

Figure 2: Rough dashboard estimate of cyber-security posture.

Regardless of public perception with respect to previous government threat metrics, any programme of situational awareness for cyber-security must include a broad characterisation of real-time risk. The attributes of this broad characterisation will be based on a much more detailed understanding of the real-time posture. Collectively, this posture is referred to as situational awareness and is based on an understanding of whether or not the infrastructure is under attack, which vulnerabilities are relevant to the local infrastructure, what sort of intelligence is available, the output of a risk management process, and information being generated by a real-time security operations centre. These elements are described in the sections that follow.

Detecting infrastructure attacks

The process of determining whether an attack on national infrastructure is under way is much more difficult than it

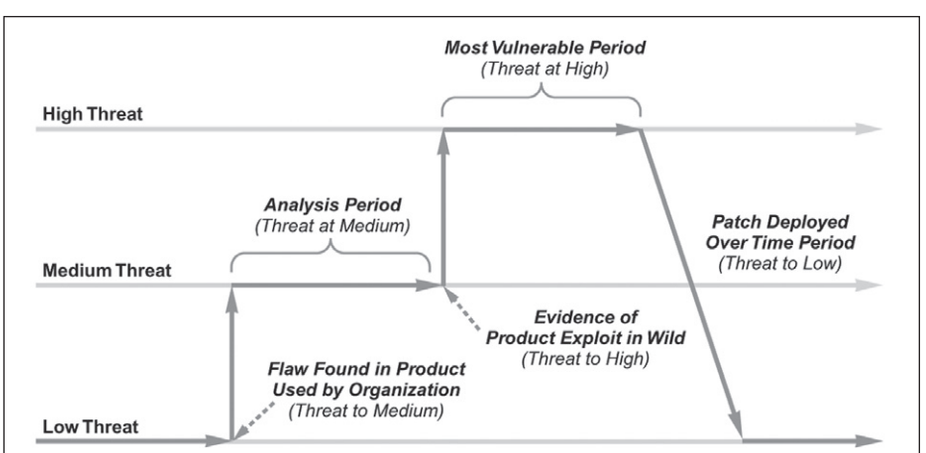

Figure 3: Security posture changes based on activity and response.

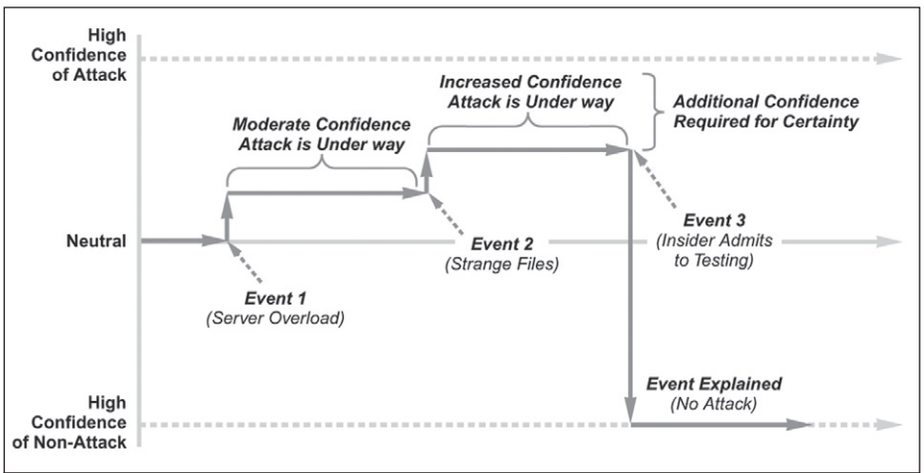

Figure 4: Attack confidence changes based on events.

sounds. On the surface, one would expect that, by observing key indicators, making the determination that an attack has begun or is ongoing would seem straightforward. Correlating observed activity with profiles, signatures and the like can provide a strong algorithmic basis, and products such as intrusion detection systems offer a means for implementation. These factors are misleading, however, and the truth is that no security task is more difficult and complex than the detection of an ongoing attack, especially if the adversary is skilled.

To illustrate this challenge, suppose you notice that an important server is running in a somewhat sluggish manner, but you cannot diagnose the problem or explain why it is occurring. Obviously, this is suspicious and could be an indicator that your server has been attacked, but you cannot state this with any certainty. There could be a million reasons why a server is running slowly, and the vast majority of them have nothing to do with security. Suppose, however, that you discover a recently installed directory on the server that is filled with unfamiliar, strange-looking files. This will clearly raise your suspicion higher, but there are still numerous explanations that do not signal an attack. Perhaps, finally, someone in the enterprise steps forward and admits to running some sort of benign test on the server, thus explaining all of the errant conditions. The point is that confidence that a target is under attack will rise and fall, depending on the specifics of what is being observed. Obviously, there is a threshold at which the confidence level is sufficiently high in either direction to make a sound determination. In many practical cases, analysis never leads to such a confidence threshold, especially in complex national infrastructure environments (see Figure 4).

In our example, you eventually became confident that no attack was under way, but many scenarios are not terminated so cleanly. Instead, events expose a continuing stream of ongoing information that can have a positive, negative or neutral effect on determining what is actually going on. In many cases, information that is incorrect or improperly interpreted has the effect of confusing the process. Relatively new technologies, such as mobile wireless services, tend to exhibit this property, especially in cases where a particular incident has never been seen before. The primary disadvantage of never determining the root cause of an attack is that the security posture cannot be accurately measured. This is especially troublesome when the attack is severe and targets essential national infrastructure services.

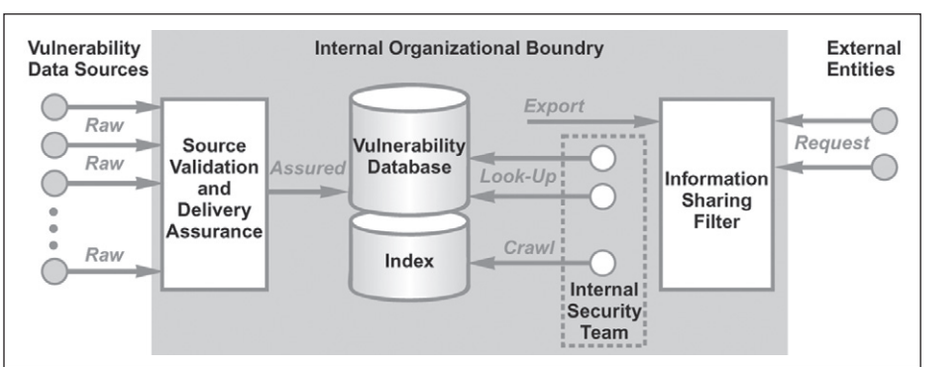

Figure 5: Vulnerability management structure.

Managing vulnerability information

A common cynical view of computer security is that its experts are engaged in nothing more than a game of Trivial Pursuit around attack and vulnerability information. Support for this view is evident in the security books published to date, most of which contain page after page of esoteric attack specifics that are often long-since irrelevant. It is also evident in social circles at security and hacking conferences, where the discussion rarely addresses foundational topics of software engineering or system design but instead focuses on such trivia as which systems have which bugs in which versions on which hardware. Some security experts and hackers have become walking encyclopaedias of such knowledge, even viewing information as the driver of power and skill. Anyone not possessing sufficiently detailed knowledge is thus tagged a newbie, lamer or perhaps worse – a *manager*.

In spite of this odd phenomenon, situational awareness for national infrastructure protection does require a degree of attentiveness to daily trivia around vulnerability information. We refer to the information as trivia simply because, once addressed and fixed, the value of the information drops very close to zero. Nevertheless, it is important information to collect, and most national infrastructure teams use the default approach of *active opportunism* , where a set amount of effort is expended to gather as much data as possible and anything else that comes in is welcomed.

The problem with active opportunism is that it will never be complete and cannot be depended upon for accurate management decisions. For example, the question of whether a given

vulnerability has been coded into an exploit and made available on the Internet can be researched by one, two or 50 people. If no evidence of such an exploit is found, then the weak conclusion can be drawn that it does not exist. Obviously, information about the vulnerability could be tucked away in some IRC discussion or on an obscure hacking site, but unless it is found or volunteered, the security team will never know for sure.

The best one can hope for is to create as active and complete a vulnerability information-gathering process as possible. Managers are strongly advised to follow three basic rules:

1. Always assume that the adversary knows as much or *more* about your infrastructure than you do.
2. Assume that the adversary is always keeping vulnerability-related secrets from you.
3. Never assume that you know everything relevant to the security of your infrastructure. Such complete knowledge is unattainable in large, complex national infrastructure settings.

Cyber-security intelligence reports

A technique commonly used in government intelligence community environments, but almost never in most enterprise settings, involves the creation and use of a regularly published (usually daily) intelligence report. For cyber-security, such a report generally includes security-related metrics, indicators, attack-related information, root-cause analysis, and so on for a designated period. It is typically provided to senior management, as well as all decision-makers on the security and infrastructure teams. The report should also be indexed for searches on current and previous information, although this is not a common practice.

Although the frequency and content of intelligence reports should be tailored to the needs of the local environment, some types of information that one would expect in any daily intelligence report include the following:

- **Current security posture**: the situational status of the current security risk would be required in any intelligence report, especially one issued over a daily or weekly interval (monthly intervals create too long a gap for information to be considered 'intelligence').
- **Top and new security risks**: characterisation of the top risks, as well as any new risks, is also important to include in an intelligence report. Visualisation and other techniques are often helpful to highlight changes in risk posture.
- **Automated metrics**: security systems that generate metrics should provide input to the intelligence report, but care must be taken to avoid the creation of a voluminous document that no one will read. Also, raw output from some devices is indiscernible and should be either summarised or avoided in the report.
- **Human interpretation**: ultimately, the most useful cyber-security intelligence includes analysis by experienced and expert human beings who can interpret available security data and recommend suitable action plans. It is unlikely that this interpretation function will be automated in the near future.

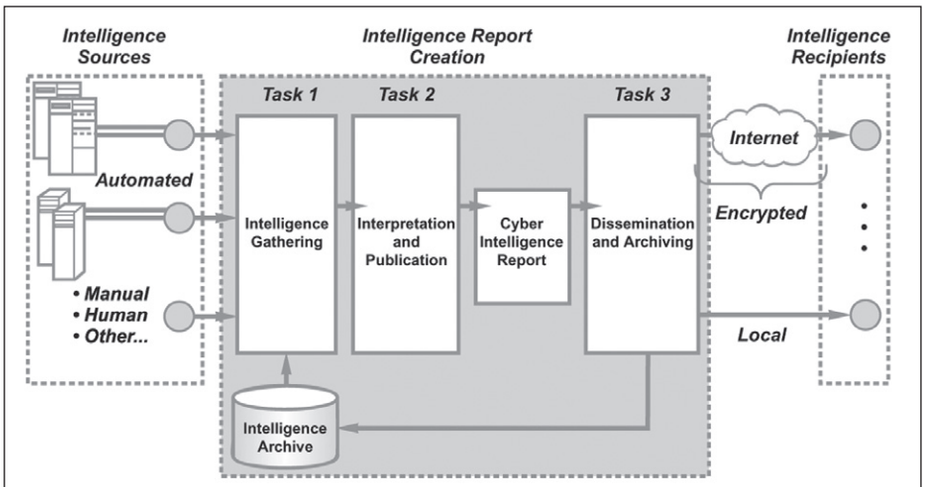

Figure 6: Cyber-security intelligence report creation and dissemination.

The activity associated with the realisation of a cyber-security intelligence report can be viewed as an ongoing and iterative process.

One by-product of creating an intelligence report is that it helps guide the local culture toward greater attentiveness to real-time security considerations. Everyone knows that, during an incident, response activity summaries will find their way to senior managers, which tends to heighten concentration on the accuracy and completeness of the report. In addition, when an incident occurs that does not find its way into the report, managers

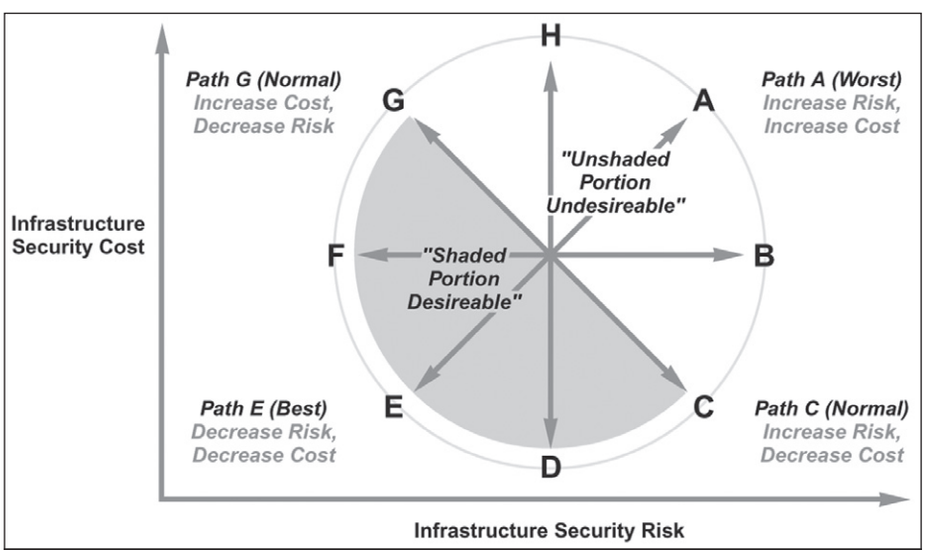

Figure 7: Risk versus cost decision path structure.

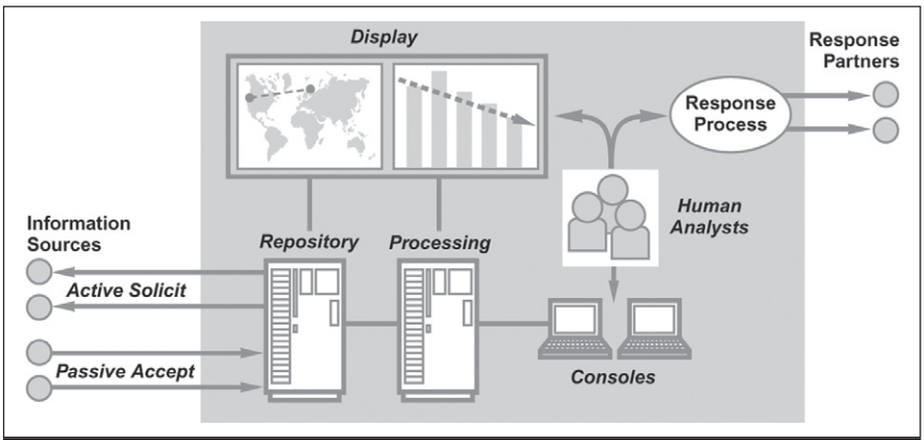

Figure 8: Security operations centre (SOC) high-level design.

can justifiably question the completeness of reporting around the incident.

Risk management process

Managers of essential national services must understand the security risks associated with their underlying infrastructure. Although this can be done using all sorts of fancy risk taxonomies, tools and methodologies, the recommended approach is to simply maintain a prioritised list. Depending on the severity of the risks in the list, managers can decide to focus on a subset of the top ones, perhaps the top 10 or 20. Funding and resource allocation decisions for cyber-security can then be driven by the security risk profile of the organisation, keeping in mind that the list of risks will change with any adjustments in threat environment, technology deployment, or reported vulnerabilities.

The generally agreed-upon approach to measuring the security risk associated with a specific component begins with two estimations:

- **Likelihood**: this is an estimate of the chances that an attack might be successfully carried out against the specific component of interest.
- **Consequences**: this is an estimate of how serious the result might be if an attack were carried out successfully.

These two estimates must be performed in the context of an agreed-upon numeric range. The actual values in the range matter less than the relative values as the estimates increase and decrease. The simplest and most common values used are 1, 2 and 3, corresponding to low, medium and high for both estimates. Once the likelihood and consequences have been estimated, risk is obtained by multiplying the values. Thus, if some component has a high likelihood of attack (value 3) and medium consequences resulting from an attack (value 2), then the associated risk is 3 times 2, or 6. If security measures are put in place to reduce the likelihood of an attack to medium (value 2), then the risk is now 2 times 2, or 4. Again, the absolute value of risk is less important than the relative value based on security decisions that might be made. A useful construct for analysing security decisions in infrastructures compares relative security risk against the costs associated with the recommended action. The construct allows managers to consider decision paths that might increase, decrease or leave unaffected the security risk, with the balancing consideration of increased, decreased or unaffected associated costs (see Figure 7).

To interpret the choices in the decision path structure, start at the middle of the diagram and consider the effects of each path labelled A through H. The path labelled G shows a security decision that increases costs in

Practical heuristics for managing vulnerability information

- **Structured collection**: the root of all vulnerability management processes must be some sort of structured collection approach with means for assuring proper delivery of information, validating the source, cataloguing the information in a suitable taxonomy and maintaining a useful database for real-time reference with provision for indexing and crawling vulnerability data in real-time. This structured approach should be integrated into all day-to-day cyber-security activities so that accurate vulnerability information is available across the entire security infrastructure and team. Filters should exist to assure incoming data, as well as to ensure that external entities only obtain appropriate information.
- **Worst-case assumptions**: many situations arise where a security team cannot determine whether some important piece of vulnerability-related information has actually been disclosed or has become known to an adversary group. The most mature and healthy approach in such scenarios is to assume the worst possible case. Most experts would agree that if the possibility arises that some vulnerability *might* be known externally, then it probably *is* known.
- **Non-definitive conclusions**: making definitive statements about national infrastructure security is not recommended. Too many cases exist where a security team draws the confident conclusion that a system is secure only to later obtain vulnerability-related information to the contrary. Experienced managers understand, for example, that they should always include caveats in security posture reports given to senior leaders in government or industry.
- **Connection to all sources**: managing vulnerability information should include connections to all possible sources such as industry groups, vulnerability-reporting services, hacking conferences, internal employee reports, and customer data. Sometimes the most critical piece of vulnerability information comes from the most unlikely source.

order to reduce risk. This is a normal management decision that is generally considered defensible as long as sufficient budget is available. Similarly, the path labelled C is also normal, as it accepts increased risk in order to reduce costs, which is unfortunately a common enough decision. Interestingly, any decision path in the area shaded on the figure will be generally acceptable in most cases because the relationship between cost and risk is reasonable. The decision paths in the unshaded portion of the graph, however, are generally considered unacceptable because of the odd balance between the two factors. Decision path H, for example, increases costs with no impact on security risk. This case corresponds to the situation encountered all too often where a security safeguard is put in place that actually has zero impact on the risk profile.

To summarise, all decisions about national infrastructure protection should be made in the context of two explicit management considerations: maintaining a prioritised list of security risks to the system of interest; and justifying all decisions as corresponding to paths in the shaded portion of the decision path structure shown in Figure 7. If these two simple considerations were mandatory, considerable time, effort and money would be immediately saved for many infrastructure management teams.

Security operations centres

The most tangible and visible realisation of real-time security situational awareness is the Security Operations Centre (SOC), also referred to as a fusion centre. The most basic model of SOC operations involves multiple data, information and intelligence inputs being fed into a repository used by human analysts for the purpose of operations such as interpretation, correlation, display, storage, archival and decision making. The SOC repository is constructed by active solicitation or passive acceptance of input information, and information processing combines human analysis with automated processing and visual display (see Figure 8).

Most SOC designs begin with a traditional centralised model where the facility is tied closely to the operations of the centre. That is, methods and procedures are created that presume that SOC resources, including all personnel, are located in one place with no need for remote co-ordination. All data is stored in a local repository that can be physically protected in one location. This approach has its advantages, because it removes so many co-ordination-related variables from the management equation. That said, an SOC can be created from distributed resources in geographically dispersed locations. Repositories can be distributed, and analysis can be performed using remote co-ordination tools. Generally speaking, this approach requires more work, but the main benefit is that more expert analysts can be recruited to such an approach, especially if the requirement is that 24/7 operations be supported. Experts can be hired across the globe in a 'follow the sun' support arrangement.

Typical operational functions supported in an SOC include all human interpretation of data by experts, management of specific incidents as they arise, support for 24/7 contact services in case individuals have security-relevant information to share, and processing of any alarms or tickets connected to a threat management or intrusion detection system.

The 24/7 aspect of SOC operation is particularly useful to national-level situational awareness, because key infrastructure protection managers will know that they can obtain a security posture status at any time from a human being on call in the SOC. Government procurement efforts for national services should include requirements for this type of coverage in the SOC.

National awareness programme

The goal of supporting a national-level view of security posture should not be controversial to most security and infrastructure managers. Everyone will agree that such a view is necessary and useful for supporting national infrastructure protection-related management decisions. The challenge, however, lies with the following important practical considerations:

Tasks for creating a cyber-security intelligence report

1. The first task involves *intelligence gathering* of available vulnerability and security posture data. This can be automated but should allow for manual submission from people who might have useful information to share. Many organisations do this gathering in the early morning hours, before the bulk of the business activity begins (a luxury that does not exist for global companies).
2. The second task involves *interpretation* and *publication* of the gathered data, not unlike similar processes in daily news publications. The interpretation should focus on the audience, never assuming too much or too little knowledge on the part of the reader. It is during this task that the human interpretive summary of the collected data is written.
3. The third task involves protected *dissemination* and *archiving* of the report for use by end users with a need to know. Report transmission is generally protected by encryption, and report archives and storage are protected by access controls.

- **Commercial versus government information**: to achieve full situational awareness at the national level will require considerable support from both commercial and government entities. Groups supplying security status information must be provided with incentives and motivations for such action. Patriotic justification helps, but global companies must be more deliberate in their sharing of information with any government.
- **Information classification**: when information becomes classified, obviously the associated handling requirements will increase. This can cause problems for data fusion. In fact, the essence of data compartmentalisation for classified information is to prevent and avoid any type of fusion, especially with unclassified data. The result is that situational awareness at

the national level will probably include two views: one unclassified and public, the other based on more sensitive views of classified information.

- **Agency politics**: government agencies are famous for using information as a basis for political agendas, including support for project funding, hiring plans and facility expansion. This tendency is counter to the goal of information sharing for situation awareness and must therefore be managed carefully.
- **SOC responsibility**: if a national SOC is to be realised, then some organisation must be designated to run it. The decision as to whether this should be a defence- or civilian-related initiative is beyond our scope, but most security experts agree that current defence-related awareness initiatives provide many of the elements required in a fully functioning SOC.

If these challenges are not addressed properly, the risk is that inaccurate views of situational awareness could arise. If an agency, for example, finds out about a vulnerability but decides to *not* share this information, then a hole emerges in any national level risk estimation. Similarly, if a commercial organisation is unable to receive and process classified information, then its view of current security risk posture will not be accurate. Attentiveness to managing these issues on a case-by-case basis, perhaps as part of a national SOC, would seem the best approach.

About the author

Edward Amoroso is currently senior vice-president and chief security officer of AT&T, where he has worked in cyber-security for the past 25 years. He has also held the adjunct professor position in the computer science department at the Stevens Institute of Technology for the past 20 years. Amoroso has written four previous books on computer security, and his writings and commentary have appeared in major national newspapers, television shows and books. He is a popular commentator on cyber-security. He holds a BS degree in physics from Dickinson College, and MS/PhD degrees in computer science from Stevens Institute of Technology. He is also a graduate of the Columbia Business School.

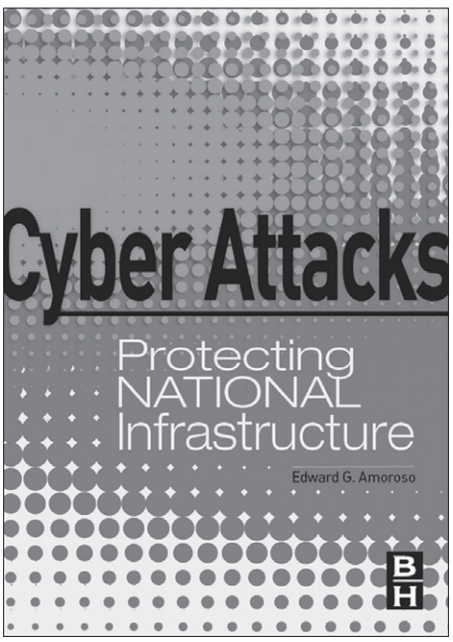

Cyber Attacks: Protecting National Infrastructure *by Edward G. Amoroso was published 1 Dec 2010 by Butterworth-Heinemann, print ISBN: 9780123849175, e-ISBN: 9780123849182. For additional information, including sample content, go to: <http://www.elsevierdirect.com/cyberattacks>.*

Advanced evasion techniques


Steve Gold

Steve Gold, freelance journalist

Much has been made in recent months of so-called Advanced Evasion Techniques (AETs) that could be used by hackers and cyber-criminals to bypass IT security defences. Do the claims stand up to technical scrutiny?

As any student of the Internet will tell you, even though its structure dates all the way back to the 1960s, its underlying packet construction is modular in nature. Put simply, this means that, like the C programming language in its many forms, the Internet Protocol Suite allows quite complex data constructs to be created and then streamed to/from various points on an IP-based network architecture.

Today, the Internet Protocol Suite is better known as Transmission Control Protocol/Internet Protocol (TCP/IP), which were the first two networking protocols defined in the Suite. In common with most network protocols, the Internet Protocol Suite consists of set of layers, each of which solves a given set of problems involving the transmission of network data. Because the TCP/IP model consists of four layers – the link layer, the Internet layer, the transport layer and the application layer – each layer must work on the basis that the others are intact – that is, that they are uncompromised and meet certain transmission formats.

In our four-layer TCP/IP model, a component of a given layer provides a well-defined service to the upper layer protocols, while at the same time using services from the lower layers. Upper layers are logically closer to the user and deal with more abstract data, relying on the lower layer protocols to translate data into forms that can eventually be physically transmitted. And it's this interrelationship between the layers that some crackers have observed as a dependence,