A National Model for Cyber Protection
Through Disrupting Attacker Command and Control Channels

Jeff Brown, CISO, Raytheon Company

In today's cyber security environment there is one inescapable truth. There is no way to prevent a determined intruder from getting into a network so long as one allows e-mail and web surfing –and no business today can long survive without these two bedrocks of the information age.

The reasons for this are simple. The vast majority of our Information Assurance architectures rely on patching and configuration control for protection, the consistent application of which has thus far proven elusive over large enterprises. It also relies on signatures for both protection and detection which, by definition, will not stop the first wave of the increasing volume of zero day attacks we are seeing today. Therefore, when you must let the attack vector (an e-mail or a web address) past your perimeter to the desktop, you are virtually guaranteed to have successful penetrations.

Raytheon believes the best way to address this new reality is to recognize that attackers will get into your network and expand our defensive actions to detect, disrupt, and deny attacker's command and control (C2) communications back out to the network. It is an acknowledgement of the fact that there are fewer, or perhaps relatively noisier, ways to get out of a network than to get into it. Such a strategy focuses on identifying the web sites and IP addresses that attackers use to communicate with malicious code already infiltrated onto our computers. While some of these sites are legitimate sites which have been compromised, the majority are usually new domains registered by attackers solely for the purposes of command and control. There is little danger of unintended consequences from blocking these web sites and their associated IP addresses for outbound traffic. Where they are legitimate sites, the benefit of protecting the enterprise far outweighs any inconvenience there might be if an employee needs to legitimately go to that site. Raytheon has had success with this strategy, but it requires a significant investment, unaffordable to most small and medium size entities and many larger ones.

One of the corollaries of recognizing that networks can always be penetrated is a shift in how we measure ourselves. Measuring ourselves against how many intrusions occur becomes a far less interesting. What counts, instead is the intruder's dwell time in our network, or how long an intruder has had access. It's more important to recognize how successful the penetrations were versus how many penetrations occurred. The ideal goal would be to have

advance notice of a new malicious C2 channel so that even if someone opened a malicious e-mail the outbound C2 channel would already be blocked—making the effective dwell time zero.

There are two ways to reduce the dwell time of an intruder, both of which we are pursuing in Raytheon. The first is to make a considerable investment in traffic analysis and analytical methods to detect the malicious outbound traffic in a network. We have had considerable success in this arena but it has required a large investment that a majority of organizations are not likely to match.

However, the other way to reduce dwell time is a method every organization, large and small, can match--collaboration with other operational entities. If we can take advantage of the good work of other organizations, we are eager to do so. We recognize that many other organizations regularly find and report C2 channels. Anti-virus vendors, CERT CC, managed security service providers, defense contractors, research institutions, intelligence agencies, other large government agencies, and law enforcement all see relatively narrow aspects of the C2 environment. But put them all together and they collectively see a very wide swath of the C2 threat environment. Many already aggregate and share the information formally or informally through ISACs, the Defense Industrial Base Cyber Task Force, Infraguard, or any number of other forums. But there is no central clearing house for this information or an operationally focused framework for rapid dissemination of this threat information to a broad national audience.

It is in the collaboration realm that Raytheon believes there is an opportunity for a national scale effort that can turn collective effort to our advantage in the cyber battle. The gaping hole in cyber collaboration (often called information sharing) is that the vast majority of small and medium–sized organizations, both commercial and government, do not participate in these groups or do not have the resources to take advantage of this information when they get it. Unfortunately, for many in critical infrastructure sectors, these small and medium-sized organizations represent a significant portion of our supply chain. We have a vested interest in their success.

While there is no national-scale framework in place, there is a model that has already proven effective fighting other cyber security problems. The model involves a set of trusted entities developing threat information and reporting voluntarily (with non-attribution) to a central source, which consolidates the information and rapidly disseminates it to a very large user community. The user communities, in return, implicitly trust the centralized service and expend little or no resources to validate the information. They simply let the automated processes protect them as a passive service rather than investing in active collaboration—and with much better results.

If this sounds familiar, it's because it is the model used for the highly successful anti-virus and spam filtering industries. We propose that this same model be used to disseminate information on attacker C2 URLs and IP addresses and automatically block outbound traffic to them. If attackers get into your network but cannot get back out the attack is effectively thwarted.

Such a model will have a tremendous impact against botnets and the advanced persistent threat both of whom make heavy use of web-based command and control. While the first wave of their attacks might initially succeed they would be short-lived after the first discovery because of the rapid and automated dissemination of the C2 channels. Subsequent waves would fail completely by virtue of rapid dissemination and automatic blocking of the C2 mechanisms. Of course, one could argue that an attacker could always rapidly change their command and control channels and make them unique to each attack. While this is true, the more we force intruders into greater costs and complexity, the more likely we are to change his cost-benefit calculations. It seems axiomatic that anything that is both simple and inexpensive while forcing this behavior is worth doing on our part.

This document, then, proposes a model for standing up a National Cyber Threat Protection Service to implement a C2 disruption strategy. It will describe the process, key relationships, and responsibilities of the participants and the incentives for each community of interest. This is a voluntary model. Within all the communities described below, not everyone has to participate for the model to be effective. The more the better, but once the process includes a critical mass, the benefits will quickly accrue to a wide swath of both the public and private sector.

An Industry-Government Cooperative Model for Disrupting Malicious Cyber Command and Control.

There are three types of entities involved in this process:

1. **Threat reporters** discover and report malicious C2 channels.
2. **A National Cyber Threat Response Center (NCTRC)** which acts as a central threat clearing house, collecting the threat reports, vetting them as necessary, and providing them to vendors in a standard format.
3. **Vendors for firewall devices** (the term here being used in its most generic sense) would accept the new threat information and push it out to their devices in the field the same way anti-virus and spam filtering vendors push new definitions today.

Certified Threat Reporters.

Threat Reporters are organizations with the detection and analytical capability to discover command and control sites via malware reverse engineering or traffic analysis. Organizations, be they commercial, private, or governmental, would apply to be certified as Threat Reporters and have their reports of C2 channels accepted as valid.

Some third party, presumably a government entity, an industry consortium or some hybrid of the two, would be responsible for certifying potential Threat Reporters against a moderate standard of in-house capabilities. The standard would measure both quality and quantity. Quality would be evaluated by a review of in-house detection and analytical capabilities designed to give *a priori* confidence in their reports' reliability. This would ensure the information the reporters provide is credible and allow for a more rapid automated dissemination process with minimum manual review. Quantity would be measured after certification to ensure the reporter was contributing enough unique threat information to the community to continue to merit the marketing advantage of being a Certified Threat Reporter.

It is important to note that submission of reports by Threat Reporters would not be the same as disclosing breaches required under other laws or agreements. A significant percentage of reports would come from intelligence or other detection activities not associated with any activity within the reporting organization's network. For this model to be viable the reporters have to be free to provide threat information without any implication that they experienced a breach or might get requests for involuntary disclosure of additional information.

Threat reporters would normally submit only malware command and control information, either web sites or IP addresses and the class of threat (e.g. botnet, advanced persistent threat, etc). That information, alone, is enough to make this model work if all parties trust the credibility of the assessment. Other detailed information on the malware involved could be voluntarily submitted, but not at the expense of rapid submission of the C2 channels.

The advantage to the Threat Reporters, especially managed security service providers, is in their ability to use the certification for branding purposes. Organizations that develop threat data internally but which do not wish to participate due to low risk tolerance or because they feel reporting might conflict with their business model would simply not apply to become Threat Reporters.

National Cyber Threat Response Center (NCTRC)

The role of the NCTRC is to serve as a clearing house for processing reports of C2 URLs and IP addresses from Threat Reporters and rapidly distributing them to the community of firewall device vendors. By having a central point disseminating the information to all vendors equally we avoid the problem we face with anti-virus today where not all vendors detect all threats. The NCTRC would also deconflict erroneous reporting that resulted in disruption to

legitimate activities. The NCTRC would maintain a "reputation index" (e.g. credibility rating) for each reporter much like seller ratings on eBay. By this feedback loop a Threat Reporter could be decertified (i.e. no longer have their reports accepted or be able to claim Threat Reporter status in their marketing).

The NCTRC must be a single organization focused on rapid dissemination of actionable information. Unlike the current anti-virus business model where organizations submit malware to their vendor of choice, there would be only one clearing house. The question of who operates the clearing house is largely irrelevant so long as everyone in the model trusts them. It could be a government entity or, more likely, a non-profit organization overseen jointly by the government and an industry consortium. Regardless of who operates the NCTRC, the government must be as secure reporting information to it as industry is. With the large amount of IP threat information the government sees simply because of the size of its network, the absence of threats detected in their networks would significantly reduce the value of the model.

Firewall Device Vendors

Producers of devices that are capable of blocking outbound web traffic would accept the data from the Clearing House, reformat it as appropriate for their device, and push it out to their customers as quickly as possible. Traditional desktop or network firewalls, web proxies, and routers would all be capable of performing this function, thus giving network owners a wide variety of products from which to select based on their architecture and investment tolerance. The vendors would differentiate themselves from each other not only on price, but also on their speed of updates and value-add services such as the ability of their customers to manually override the lists or their ability to provide reports to network owners.

Industry, Critical Infrastructure Providers, and Government

The real benefit from this model lies with the vast majority of network owners in business, industry, and government who cannot afford the deep detection and analytical capability needed to protect themselves. Today, these organizations are totally at the mercy of a determined intruder who is virtually guaranteed to be able to compromise systems with socially-engineered zero-day attacks. Most simply do not have the investment dollars to build a detection infrastructure dependent on traffic analysis or the expertise to make use of the various information sharing groups. With this model, though, these businesses could easily, and voluntarily, afford a single device that most already have anyway.

It would, however, now provide an order of magnitude increase in the level of protection by stopping in near-real time many of paths an attacker would use to get back out of the network. For those who had not been compromised yet when updates come out, they

would completely nullify any subsequent attack with that command and control channel. For those who had already been compromised in the first wave of a zero day attack, it would minimize the length of time when an attacker could access the compromised box and it would identify compromised computers that might otherwise have gone undetected. Best of all, assuming they implicitly trust the system, the organizations employing the model do not have to invest any additional resources to take full advantage of the model.

A secondary benefit would accrue to organizations whose websites have been hijacked and used as C2 sites (as opposed to dummy domains registered specifically for C2). These organizations would become aware of the infection more quickly as hits on their web sites dwindled or simply monitoring the NCTRC lists. They would be then able to exhibit good internet citizenship by quickly cleaning their systems and working with the NCTRC to be removed from the block list.

A third benefit, although perhaps more appropriate to a follow-on effort, would be the ability to tie the reported C2 channels to a library of instructions for finding and cleaning the specific malware where is was detected. This would be a much more complex and less automated process, but it would give smaller organizations a quick way to not only know they have a problem, but also allow them to short circuit the remediation process.

The Prospect of a Common Operational Picture

Perhaps one of the most tantalizing side benefits of this model is that it could be the basis of a true Common Operational Picture. If every firewall device supporting this model not only blocked the outbound traffic, but also—again, voluntarily—reported back to the Clearing House that there was a blocked C2 attempt from their IP address it would, given the potentially hundreds of thousands of devices reporting in, represent a very accurate picture of the scope of any given attack or campaign. Unlike today when organizations are loathe to report incidents because of the risk of bad publicity, data reported to this COP would not reveal any information beyond the fact that someone on their network tried to communicate with a bad URL or IP. Plus, by definition, if the firewall device blocked the outbound traffic, the attack failed or has been neutralized. But knowing the nationwide scope of attacks from the same source would yield invaluable information unavailable today.

If the IP addresses reporting in could be grouped by their critical infrastructure or agency, the COP could be filtered to that organization. For example, if the NCC knew the IP space of all nuclear power plants, a COP could show attempts to access the same C2 sites from multiple power plants. This might indicate a concerted effort to compromise the plants.

Similarly, the defense industry or financial community would see the scope of attacks across their community. Or the Department of Defense would see which attacks were unique to them since there might be no detections of specific C2 sites outside of DoD IP space. And all this in near-real time.

Incentives

This model for denying and disrupting attacker command and control on a national scale includes positive incentives for every participant.

1. Organizations, especially commercial entities, will have an incentive to be certified threat reporters for branding purposes. It shows that they have a robust, capable process and investments to become credible reporters of threat data. There could even be tiered levels for branding purposes based on the volume and accuracy of inputs, i.e. an anti-virus vendor who might report a lot of C2 URLs based on all the malware they get would be Platinum Reporters. A large company with robust internal capabilities might be a Gold level. Managed Security Service providers would be especially eager to participate since the number of C2 channels first reported by them would be a tremendous marketing tool.
2. The Government will greatly benefit by being provided a very large body of C2 URLs and IPs with very little investment on their part. They will also benefit, of course, by the overall increased security of the industrial base which is a major goal of US policy. Most important, however, is the promise of a near-real time common operating picture that truly reflects the current threat environment. The main burden on the government's part would be the up front effort to champion implementation and develop interface standards for receiving reports and disseminating them to vendors.
3. Firewall device vendors will have a great incentive to participate. They will be noticeable by their absence if they don't participate and it will most likely open up a whole new class of customers who see in a single device a high payoff defensive measure.
4. Best of all, small and medium sized organizations of all types will now have a way to take collective advantage of the investigative work of the best IA organizations in the country. By investing only in the firewall device that best fits their architecture, their security will increase by an order of magnitude or more simply because, like AV, a known bad domain will get blocked within hours of discovery.
5. This would also help to restore trust in the internet by identifying and isolating ISPs that do not maintain standards of good behavior on their networks. Their IP space and registered domains would frequently be blocked, presumably reducing their profitability and providing an incentive to good behavior.

6. Once this model is up and running it could easily be extended internationally. In fact many foreign producers would have a great incentive to have their devices capable of participating in this model. From there it is a short jump to an international model.

Risks

The main risk associated with this model is the risk of blocking a legitimate web site that has been taken over by an attacker for use as a Command and Control site or downloader site. While we believe this risk will be small compared to the gain, the model envisions a reclama or deconfliction process whereby a domain owner could get his domain removed from the list either as an error or after demonstrating his site was no longer hijacked. A secondary mitigation would be for the vendors to allow manual overrides on blocked domains at the local level, exactly as is done today with exceptions to web proxy vendors' predefined categories.

There is a secondary risk involved in building the trust relationships required to make this model work. Industry and government alike must be assured that there is no negative connotation to submitting threat data. The simple imperative of getting malware command and control data out to the broadest possible audience must take precedence.

Summary

This model, if implemented on a national scale, has the potential to be a game changer. For every attack, if a single organization discovered the attack, the entire nation would soon be protected. It would force an attacker to make the command and control channel unique for every attacked IP address. An attacker would have to either reduce the scope of attacks or greatly expand his domain registrations. In the later case, someone registering enough domains to operate on the level our attackers operate today would soon gain such a high profile they would be susceptible to other mitigations.

In the end, this model takes the best aspects of today's anti-virus, spam filtering, and proxy URL categorization to build a fourth service that is akin to anti-virus on outbound traffic. This National Model for Disrupting Attacker Command and Control proposed in this paper could set a new standard for effective public-private partnership in the Internet Age.