ATLANTIC COUNCIL
IDEAS. INFLUENCE. IMPACT.

ISSUEBRIEF

Jason Healey
Leendert van Bochoven

Strategic Cyber Early Warning: A Phased Adaptive Approach for NATO

The perfect is the enemy of the good in both military alliances and cyber security.

The dream cyber warning network would detect most attacks before they occur and quickly detect and stop the rest, preferably automatically. Unfortunately, this is currently feasible only for extremely small and heterogeneous organizations willing to commit significant resources, such as financial institutions. The challenges for a military alliance of twenty-eight nations with widely varying budgets and needs are much harder to meet.

As NATO looks ahead to building its needed cyber capabilities, then, it should recognize that few if any of these criteria for success are likely within the Alliance—or most of its member militaries—soon.

However, all is not lost. NATO could have a credible cyber Distant Early Warning Line in the near future if its leadership were to approach the problem incrementally rather than waiting for the perfect solution. The Alliance is on the right track with its recent investments in monitoring. But this should be seen as the first step in a Phased Adaptive Approach for cyber defense. Later phases could consider a more sophisticated sensor grid integrated with militaries and national grids, once there is sufficient budget and trust in the Alliance.

In the meantime, NATO should supplement its technical monitoring with information from the private sector and also intelligence-based indications and warning, to give advance notice of geopolitical situations that might lead to serious cyber conflict. The most important warnings must be for

The Smarter Alliance Initiative

This issue brief is part of the Smarter Alliance Initiative, a partnership between the Atlantic Council and IBM, established in response to Secretary General Anders Fogh Rasmussen's call for NATO members to adopt a "smart defense" approach to leveraging scarce resources to develop and sustain capabilities necessary to meet current and future security challenges in an age of austerity.

Working with recognized experts and former senior officials from Europe and the United States, the Atlantic Council and IBM have produced a set of policy-oriented briefs focused on NATO reform and cyber security, with the aim to provide thought leadership and innovative policy-relevant solutions for NATO's continued organizational reform and role in cyber security.

The publications and their findings will be showcased at public and private events for the defense policy and NATO communities on both sides of the Atlantic.

For more information about the Smarter Alliance Initiative, please contact Barry Pavel, Director of the Atlantic Council's Brent Scowcroft Center on International Security, at **bpavel@acus.org** or Leendert van Bochoven, NATO and European Defense Leader, IBM, at **L_van_Bochoven@nl.ibm.com**.

Jason Healey is the director of the Cyber Statecraft Initiative at the Atlantic Council. You can follow his comments on cyber issues on Twitter at **@Jason_Healey**.

Leendert van Bochoven is NATO and European Defense Leader for IBM.

The real goal for early warning is to detect attacks in time to put sufficient countermeasures in place beforehand to stop the attack or minimize its effects.

critical attacks—major incidents that could potentially invoke Article 4, Article 5, or disrupt ongoing combat operations—and these are most likely during times of ongoing physical conflict or tensions.

Understanding Early Warning

One of the most well known examples of early warning is the Distant Early Warning (DEW) Line set up by the United States and Canada during the early days of the Cold War.[1]

But detecting an inbound attack is only a small part of early warning, a fact much overlooked in the cyber field, rooted in a mistaken belief that cyber incidents happen so quickly that "early" has to be measured in milliseconds. Accordingly, much of what passes for cyber early warning is actually tactical warning and attack assessment (TW/AA) which uses radars, satellites, and other sensors to detect attacks immediately and determine how serious they may be. However, any warnings provided through TW/AA—whether for nuclear forces or cyber attacks—are not very early, providing only hours to minutes for decision makers to react.

Since the real goal for early warning is to detect attacks in time to put sufficient countermeasures in place beforehand to stop the attack or minimize its effects, the short timelines of TW/AA generally are not enough on its own.

Accordingly, providing earlier, strategic warning of attacks —weeks or even months ahead—has been an intelligence task, a key element of which was to determine if the geopolitical situation becoming so tense that the Soviet Union (or indeed the United States) would be willing to launch a strategic first strike.

Warning is not just about informing decision makers that an attack is likely or inbound but also the converse: for the time being no attacks are likely and they can turn their attention to more pressing matters.

Cyber warning is no different: to have the maximum time to respond, defenders must not only be able to detect inbound attacks but also look for the intent of adversaries before they actually decide to turn the launch key (or press the enter key). Figure 1 illustrates this relationship between TW/AA and strategic and tactical warning for strategic and tactical attacks.

Figure 1: Relationship Between Kinds of Early Warning

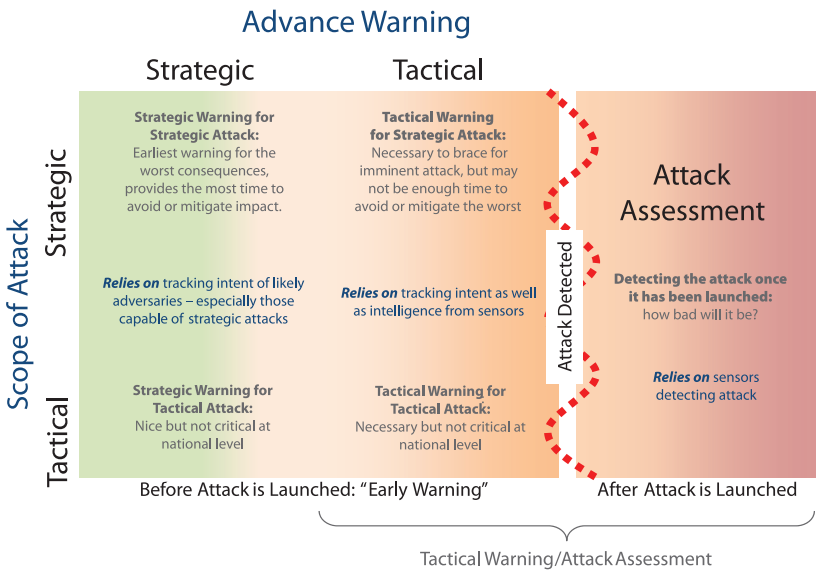

Moreover, warning is not just about telling generals, ambassadors, and ministers that an attack is likely or inbound but just as importantly the overlooked converse: having confidence to tell them that for the time being significant attacks are not likely and they should turn their attention to more pressing matters.

Early Warning for Cyber Attacks—And Its Difficulties

Despite the speed of individual attacks, warning in the cyber realm has the same goal as it did in the Cold War: detecting attacks in

1 Though DEW was only the northernmost of three such lines of radar sites to detect bombers, the term has somewhat become genericised to mean any generic line of sensors providing warning of an inbound attack.

time to put sufficient countermeasures in place beforehand to stop the attack or minimize its effects. Within the United States, the role of early warning is perhaps clearest with the government's EINSTEIN program. Early versions detect attacks as they unfolded (attack assessment), the newest EINSTEIN 3 variant will recognize an attack and stop it, hopefully before any negative impact (tactical warning). The speed of cyber attacks has shortened tactical warning times, blurring tactical warning and attack assessment (see Figure 2).

Unfortunately, because of the technical nature of the domain and its practitioners, "early warning for cyber attacks" often is equated to only gadgets like EINSTEIN monitoring for hostile zeroes and ones inbound on the wire. This kind of cyber warning is plagued by multiple problems, each difficult and together sometimes insurmountable, such as issues detecting malicious attacks in massive flows of Internet traffic and determining if different attacks are part of the same campaign, especially when multiple organizations have been targeted.

Figure 2: Relationship Between Kinds of Early Warning Applied to Cyber Conflict

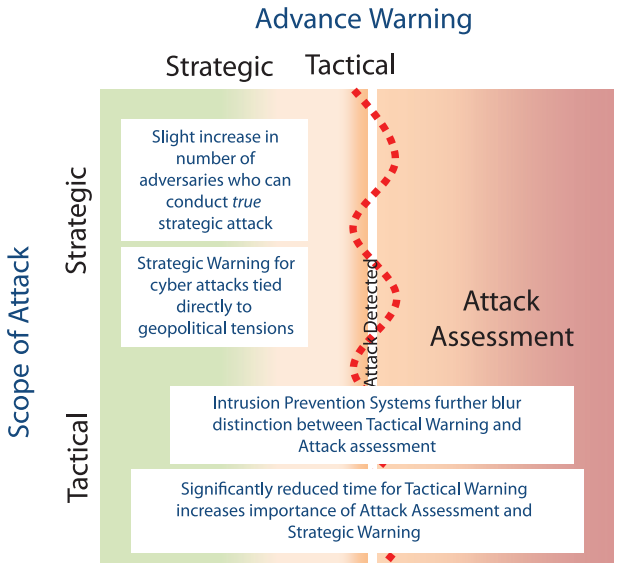

So, though detecting inbound attacks is critical, it is not nearly enough, especially not for military organizations working in a political (and politicized) context. To give defenders more time, monitoring must be complemented with intelligence-driven strategic warning based on known or presumed changes in intent of adversaries capable of strategic cyber attacks.

Warning does not have to be perfect, just "good enough" to enable leaders to start making hard choices.

NATO's Present and Future Cyber Early Warning

In early 2012, NATO began accepting bids for a $42 million expansion to their monitoring network which "will need to collect and sift through vast amounts of data across NATO networks stretching from the US to Afghanistan," according to press, and report the results back to the NATO Computer Incident Response Capability (NCIRC). In addition, the defense ministers gave the go-ahead for NCIRC to monitor and respond to incidents not just directed at the military organizations of the Alliance, but related civilian agencies (like the NATO Defense College).

NATO has not looked deeply beyond technical sources for warning information, though this is starting to change. According to one NATO official, "Traditional intelligence support for cyber aspects of crises is nascent, but developing," with a dedicated Cyber Threat Assessment Cell to help with attack assessment by integrating data from NATO networks with open-source intelligence.

Nations and national militaries also maintain significant early warning capabilities as do private sector companies. As a result of previous summits, Alliance members agreed to sign memoranda of understanding with NATO to share warning and other information. Progress, though, has been slow with only 20 member nations having signed memoranda (as of September 2012) and many practical hurdles (like trust, format, and content) to leap. A better alternative, described later in this issue brief, is to rely much more on the private sector for warning information. Much of the best information resides there already and can be had immediately for the price of a yearly subscription.

NATO and Strategic Cyber Attacks

As NATO and its members are routinely hounded with cyber attacks, this section describes ways to unpack the issues involved and understand which might be "strategic" and which "tactical."

The reasons for NATO's existence are collective defense and cooperative security. Accordingly, NATO's true cyber priorities should be incidents so serious that they will lead to potential collective defense or political response within the Alliance or put soldiers' lives directly at risk by disrupting combat operations. These should be strategic and all others cyber incidents would be tactical.

At a deeper level, a strategic attack would differ from a tactical attack in at least four key ways, **purpose, target, context, and scale**. That is, NATO would have to assess whether the adversaries disrupting systems or merely spying on them? Are the adversaries targeting critical infrastructure of allies? Or are the attacks against NATO itself? Is the incident during wartime or peacetime? And most importantly, is the incident extremely serious or a mere nuisance? Table 1 below illustrates how these four factors can help determine whether the attacks NATO has suffered so far may have been strategic.

Much of the best warning information is in the private sector and can be had for the price of a yearly subscription.

Applying these criteria more generally, there are five kinds of incidents that clearly fall short of the "strategic" threshold:

- Quotidian crime and nuisance attacks on companies, governments, individuals
- Significant disruptive attack on NATO member countries (but not crossing the Article 4 threshold because of lack of scope, duration, intensity or national responsibility for the attack)
- Criminal or nuisance attack on NATO
- Major espionage or disruptive attack on NATO headquarters or forces *during peacetime*
- Espionage or intrusion against NATO headquarters or forces *during wartime*

The first several of these incidents should be of no immediate concern to the Alliance while those at the bottom are important but are unlikely to be strategic attacks that seriously impugn defense and security at the national level.

Contrast these with incidents that are likely to cross the "strategic" threshold.

- Disruptive intrusion or major espionage against NATO combat operations *during combat operations*
- Potential or actual Article 4 cyber incident (leading to high-level political consultation)
- Potential or actual Article 5 cyber incident (leading to collective defense)
- Potential or actual Article 4 or 5 cyber incidents *during military operations* (for example, a devastating cyber attack that affects out-of-area operations in Afghanistan)

Table 1

	Incidents Against NATO	Purpose	Target	Context	Scale	Strategic?
ALLIED FORCE (1999)	Defacements and DoS	Disruptive	NATO, Allies	Wartime	Nuisance	No
Estonia (2007)	DoS	Disruptive	Ally	Crisis	Significant	**Yes**
Georgia (2008)	Defacements and DoS	Disruptive	Partner	Wartime	Significant	Possibly
UNIFIED PROTECTOR (2011)	DoS	Disruptive	NATO, Allies	Wartime	Nuisance	No

These strategic cyber incidents and must be the top priorities for cyber early warning. They require the longest warning times, whether to enable defenders to erect sufficient technical countermeasures or for political leaders to enlist allies or warn off the probably attackers.

Fortunately, these strategic level cyber incidents are actually the easiest kind for which to provide warning, through intelligence-based earliest assessments of adversaries' intent.

Making Strategic Cyber Warning Easy

The most significant cyber incident NATO has seen to date was the 2007 denial-of-service attacks against Estonia. According to participants at a recent Atlantic Council event, the technical defenders had "a couple of weeks" advanced warning as the mostly Russian hackers gathered their forces and organized online. Yet this warning did not reach the policy makers at NATO until the attacks were already underway, a needless strategic surprise.

For reasons already noted, strategic cyber incidents must be NATO's highest warning priority. Fortunately, and despite myths to the contrary, strategic warning for strategic cyber incidents is actually easier and cheaper than detecting the actual inbound attacks themselves and yields longer warning times.

The intelligence community has an existing indications and warning methodology (I&W) that fits strategic cyber warning extremely well. Capable of spotting trends in adversary intent, I&W tracks observable phenomena that can help explain and predict adversary activity—or the lack of it. The process is based on defined "warning problems" and indicators to track observable phenomena (or analytical assessments if the phenomena are not directly observable).

To guide analysts, over the past fifteen years, a small group of intelligence analysts, both in and out of government, have recognized many useful rules of thumb (see Text Box 1).[2] These rules, which often directly contradict popular myths of cyberspace, are directly relevant to cyber warning.

The most important warning problem for the Alliance is likely to be "Will major Country X conduct or support a large-scale disruptive cyber attacks against a NATO member?" which would include indicators such as:

Rules of Thumb for Strategic Cyber Warning

Strategic disruptive effects are incredibly difficult, far harder than are usually thought.

No known attacks have been *both* widespread and persistent enough to make have a significant prolonged international impact.

Cyber adversaries are people not ones and zeros traveling down a wire.

Adversaries with the capability to cause a strategic effect are not individuals but organizations, linked to nation states.

Adversaries with the intent to conduct a strategic attack usually lack the capability, while those with the capability lack the intent.

Nations have *never* launched a major disruptive or destructive attack against another nation unless during a period of significant tensions.

Physical conflict begets cyber conflict.

Therefore:
Strategic warning for cyber attacks works best to predict major disruptive attacks where there are a limited number of potential adversaries, those adversaries are usually associated with a foreign nation, and intent can be gleaned from and moves in response to world events.

This is a *perfect* match for NATO's traditional collective defense mission.

1. Are tensions with Country X very high or getting worse?
2. Are youth or other patriotic groups staging physical protests by against NATO, such as at member embassies?
3. Are senior Country X officials agitating and encouraging protests? Are the protests being egged on by media outlets tied to the government of Country X?
4. Is the Country X military taking an aggressive stance towards NATO, such as by mobilizing or conducting large-scale exercises near border areas?

2 These analysts include Matt Devost, Bob Gourley, Tom Parker, Ned Moran, Michael Tanji, and Jason Healey.

5. Has there been a marked increase in incidents against NATO by patriotic Country X hackers?
6. Are patriotic hackers conducting lower-level attacks or are they more coordinated, sophisticated, or resourced?

Note that these indicators are all readily observable and become gradually more severe as the list progresses. A large-scale disruptive cyber attack from Country X is extremely unlikely if all of all of the indicators are green. The more that are "tripped" and become red, the more likely an attack. If all are red, then a large-scale disruptive cyber attack from Country X might come at any time.

Besides transparency and ease of use, this methodology has another great advantage: it does not necessarily rely on secret intelligence sources. Since the indicators are tied to major, real-world security events, most or all of them can usually be tracked using open-source information gleaned from newspapers, television, and online sources like social media and bloggers.

The color table illustrates how such open-source indicators work in real life, used by the US finance sector to help predict whether there might be disruptive cyber attacks against banks because of the US invasion of Iraq in 2003. Developed by one of this paper's authors, Healey, these indicators made it clear what evidence was needed and made this process transparent to other analysts and to decision makers. The color coding of the indicators show that almost none of the likely precursors to a major disruptive campaign against the finance sector was likely. Decision makers could direct their attentions elsewhere. While this model is not very useful for crime or espionage that can take place at any time, this is not a significant drawback as these incidents are below the level that would demand a NATO response.

Real-World, Private-Sector Example

Cyber indicators used by the Financial Services ISAC before the Iraq war in 2003 to help track the likelihood of attacks intentionally targeting our sector

Indicators for Cyber Attacks Against Finance Sector During Iraq War, April 2003			
Indicator	**Current Status**	**Expected Trend**	**Past Trend**
War begins with Iraq	High	Steady	Increasing
Finance-related themes ("Big Banking supports the war") of protestors or hackers	Low	Increasing	Decreasing
War-related terrorist attacks or large scale protests against US or US Interests overseas	Low	Steady	Steady
War-related terrorist attacks or large scale protests over financial sector	Low	Steady	Steady
Protests or attacks against perceived "Jewish" targets in US or UK	Inactive	Steady	Steady
Protest hacking movement strengthens	Low	Steady	Increasing
Hacks against specific government war-related targets (OSD, CENTCOM, White House, DHS, etc.)	Inactive	Steady	Steady
Hacks against specific commercial war-related targets (weapons manufacturers, oil companies)	Inactive	Steady	Steady
Hacks against specific financial war-related targets	Inactive	Steady	Steady
Hacks against general financial targets	Inactive	Steady	Steady
"Cyber war" between US/UK hackers and anti-US/UK hackers	Low	Increasing	Steady
Incidents or reports of strictly defined "cyberterrorism"	Inactive	Steady	Steady
Reports of offensive computer network attack on Iraq	Unknown	Steady	Steady

Recommendations for NATO

As this Issue Brief has described, early warning for cyber incidents requires both technical monitoring and non-technical intelligence assessment, targeted at the most worrying strategic attacks. NATO has some capacity, but there is far more in national governments and militaries and the private sector. The Alliance should therefore undertake to implement the following recommendations.

NATO should **continue critical improvements already underway** to monitoring and protecting its own systems, but **begin planning to include these into a Phased Adaptive Approach** similar to that being used for NATO missile defense. This plan should ensure the Alliance masters the easiest and most important tasks in the short term while determining which more sophisticated defenses are needed or feasible. Earlier phases of the plan must include continuing current efforts and emphasize on the need to **create links with national and military computer emergency response teams, and companies working in cybersecurity and national network providers**. The private sector, in particular, detects many of these critical attacks on their own and often only need trusted partners in their governments to pass along sufficient warning. Even better, getting threat data from cybersecurity companies does not require international agreements or trust relationships, just a credit card number. For later phases of an adaptive plan, NATO can investigate the requirements and feasibility of other options, more expensive and expansive, such as a tighter regional monitoring network or defenses that automatically detect and stop attacks.

Summary of Recommendations

1. Continue critical improvements already underway
2. Begin planning for a Cyber Phased Adaptive Approach or equivalent
3. Create links with national and military CERTs
4. Create links with cybersecurity companies and network providers
5. Set a political goal of “no surprises” for disruptive critical attack
6. Expand the Cyber Threat Analysis Cell to a full cyber intelligence team
7. Define and track cyber warning problems and indicators for NATO
8. Create and practice the communication paths for warning to reach defenders and policy makers

NATO leadership should **set a goal to never be surprised by a disruptive, critical attack**—one that might have strategic effect by disrupting ongoing military operations or potentially creating an Article 4 consultation or invocation of Article 5 collective defense provisions. This should be the top warning priority for the NATO intelligence and cyber organizations and requires an emphasis on the traditional intelligence skill of indications and warning. Early warning is, first and foremost, an intelligence task and not a technical one since most critical incidents, as defined above, will come from adversary nation-states (or their proxies) during heightened political tensions.

Threat analysis cells often help bring intelligence sources and methods to help respond to ongoing attacks, not predict new ones. Accordingly, NATO should **expand its Cyber Threat Analysis Cell to a full cyber intelligence team** (as the Department of Defense did in 1998 when creating their first cyber joint task force). In the meantime, the Cell should work with companies working in cybersecurity and national military intelligence organizations, such as the US Defense Intelligence, to **define and track cyber warning problems and indicators for NATO**, particularly to warn about the most likely major disruptive cyber attack scenarios.

Lastly, NATO must have a **process to react to warnings**. NATO must build paths of communications to connect those providing the warning to the cyber defenders and the political leadership. These paths should be established in peacetime and practiced during exercises for example during the annual Crisis Management Exercise (CMX) at NATO HQ. It will take years for NATO to be able to provide early warning for all cyber attacks. Still, the Alliance could have a reasonable cyber DEW Line much earlier if it focuses on only strategic attacks, teams with nations and the private sector for information, and develops an indications and warning system—none of which are prohibitively expensive or difficult. This affordable Cyber DEW Line will improve NATO cybersecurity but perhaps it’s most valuable contribution will be to ensure the political leadership of the Alliance will not be surprised by large-scale incidents which can be entirely foreseeable using longstanding practices already in place in member militaries.

OCTOBER 2012

The Atlantic Council is a nonpartisan organization that promotes constructive US leadership and engagement in international affairs based on the central role of the Atlantic community in meeting today's global challenges.

1101 15th Street, NW, Washington, DC 20005 (202) 463-7226
www.acus.org