



Defence Research and
Development Canada

Recherche et développement
pour la défense Canada



Dynamic Defensive Posture for Computer Network Defence

Concepts and Research Directions

Craig Burrell, Julie Lefebvre and Joanne Treurniet

Defence R&D Canada – Ottawa

TECHNICAL MEMORANDUM

DRDC Ottawa TM 2006-250

December 2006

Canada

Dynamic Defensive Posture for Computer Network Defence

Concepts and Research Directions

Craig Burrell
DRDC Ottawa

Julie Lefebvre
DRDC Ottawa

Joanne Treurniet
DRDC Ottawa

Defence R&D Canada – Ottawa

Technical Memorandum

DRDC Ottawa TM 2006-250

December 2006

Principal Author

Original signed by Dr. Craig Burrell

Dr. Craig Burrell

Approved by

Original signed by Dr. Julie Lefebvre

Dr. Julie Lefebvre
Head/Network Information Operations

Approved for release by

Original signed by Dr. Cam Boulet

Dr. Cam Boulet
Chair/Document Review Panel

© Her Majesty the Queen in Right of Canada as represented by the Minister of National Defence, 2006

© Sa Majesté la Reine (en droit du Canada), telle que représentée par le ministre de la Défense nationale, 2006

Abstract

This paper examines the concept of dynamic defensive posture in computer networks. We propose that defensive posture is provided by knowledge of whether and how critical network resources are vulnerable to attack. After introducing the basic concept, we discuss the constituent elements of defensive posture. We then review relevant technologies, finding that current technology addresses only aspects of the problem. Finally, we propose a variety of research problems for which the solutions would contribute significantly to our ability to identify a network's defensive posture.

Résumé

Le présent article examine le concept de la position défensive dynamique dans les réseaux informatiques. Nous proposons que la position défensive soit déterminée par la connaissance de la vulnérabilité des ressources essentielles du réseau aux attaques et de la manière dont elles le sont. Après avoir présenté le concept de base, nous discutons des éléments qui composent la position défensive. Nous examinons ensuite les technologies pertinentes, et constatons que la technologie actuelle ne permet de régler que certains aspects du problème. Enfin, nous proposons une variété de problèmes de recherche pour lesquels les solutions contribueraient nettement à notre capacité de déterminer la position défensive d'un réseau.

This page intentionally left blank.

Executive summary

Dynamic Defensive Posture for Computer Network Defence

Craig Burrell, Julie Lefebvre, Joanne Treurniet; DRDC Ottawa TM 2006-250;
Defence R&D Canada – Ottawa; December 2006.

Background: A network administrator needs to understand network security, and how network management decisions affect that security. The complexity of large networks, however, in which user activities, security policies, local configurations, and software vulnerabilities interact, makes it difficult for the administrator to know what is happening on the network, much less understand the significance of the activity. Ideally, the administrator should be aware not just of the low-level details of network events, but of their high-level impact on the operations and services which the network supports.

This paper introduces dynamic defensive posture in the context of computer network defence. The dynamic defensive posture of the network combines knowledge of those resources exposed to attack with those critical to the support of network services in order to indicate the extent to which attacks could affect the network's operation.

Principal results: We find that no existing technology meets all the requirements for dynamic defensive posture in computer networks. We propose a number of research problems for which the solutions would contribute to the goal of a working dynamic defensive posture system. We point out the need for improved network modeling, fast methods of discovering possible multi-stage attacks, algorithms for severity ranking of network attacks, and schemes for assigning value to network assets in a way that reflects their role in providing high priority network services. We also identify a need for methods of handling incomplete data in the network model, and for visualization technologies.

Significance of results: It is hoped that these research proposals will inspire further work on the subject, serving as a road-map to the goal of producing a working system for dynamically identifying a network's defensive posture.

Future work: The effort to build a system capable of determining dynamic defensive posture is one stage of a larger research program. Once successful, dynamic defensive posture could serve as a foundation for research into Course of Action technologies which not only identify security problems, but recommend ways of addressing them.

This, in turn, would make possible Automated Response systems which repair discovered security problems without the direct intervention of a network administrator, and may eventually lead to the design of Self-healing Networks.

Sommaire

Dynamic Defensive Posture for Computer Network Defence

Craig Burrell, Julie Lefebvre, Joanne Treurniet; DRDC Ottawa TM 2006-250;
R & D pour la défense Canada – Ottawa; décembre 2006.

Contexte: Un administrateur de réseau doit comprendre la sécurité du réseau et comment les décisions relatives à la gestion du réseau influent sur cette sécurité. Toutefois, la complexité des grands réseaux, dans lesquels interagissent les activités de l'utilisateur, les politiques en matière de sécurité, les configurations locales et les vulnérabilités des logiciels, fait en sorte qu'il est difficile pour l'administrateur de savoir ce qui se passe sur le réseau, et encore plus difficile de comprendre l'importance de l'activité. Idéalement, l'administrateur doit connaître non seulement les détails de niveau inférieur des événements de réseau, mais aussi l'incidence de haut niveau sur les opérations et les services soutenus par le réseau.

Cet article présente une position défensive dynamique dans le contexte de la défense des réseaux informatiques. La position défensive dynamique du réseau combine la connaissance des ressources exposées aux attaques avec celles qui sont essentielles au soutien des services du réseau afin de déterminer la mesure dans laquelle les attaques pourraient influencer sur les activités du réseau.

Résultats principaux: Nous constatons qu'aucune technologie existante ne satisfait à toutes les exigences d'une position défensive dynamique dans les réseaux informatiques. Nous proposons un certain nombre de problèmes de recherche pour lesquels les solutions contribueraient à la création d'un système de position défensive dynamique qui fonctionne. Nous soulignons le besoin d'avoir une modélisation de réseau améliorée, des méthodes rapides pour détecter des attaques possibles à plusieurs étapes, des algorithmes pour le classement de la sévérité des attaques sur le réseau et des mécanismes permettant d'attribuer une valeur aux éléments d'actif du réseau d'une manière qui correspond à leur rôle dans la prestation de services de réseau à priorité élevée. Nous identifions également la nécessité d'avoir des méthodes de traitement de données incomplètes dans le modèle du réseau et des technologies de visualisation.

Importance des résultats: Nous espérons que ces propositions de recherche inspireront des travaux plus poussés à cet égard, qui servent de carte routière pour produire un système qui fonctionne afin de déterminer de façon dynamique la position défensive d'un réseau.

Travaux futurs: Les efforts déployés en vue de bâtir un système capable de déterminer la position défensive dynamique ne constitue qu'une étape d'un programme de recherche plus important. Dès qu'elle donnera de bons résultats, la position défensive dynamique pourrait servir de base pour la recherche dans des technologies plan d'action qui non seulement détectent des problèmes de sécurité, mais aussi recommandent des façons de les régler. Il serait ainsi possible d'avoir des systèmes de réponse automatisée qui réparent les problèmes de sécurité détectés sans intervention directe d'un administrateur de réseau, ce qui pourrait entrainer la création de réseaux d'autorégénération.

Table of contents

Abstract	i
Résumé	i
Executive summary	iii
Sommaire	v
Table of contents	vii
List of figures	ix
Acknowledgements	x
1 Introduction	1
1.1 Situational awareness	1
1.2 Critical resources	2
1.3 Exposure	2
1.4 Defensive posture	3
1.5 Dynamic defensive posture	3
2 Requirements for Defensive Posture	4
2.1 Questions about critical resources	4
2.2 Questions about exposed resources	4
2.3 Questions about exposed, critical resources	6
3 Elements of Defensive Posture	7
3.1 Operations requirements	7
3.2 IT infrastructure (ITI)	8
3.3 Software vulnerabilities	9
3.4 Exploits	10
3.5 Safeguards and security policy	11

4	Related Work	12
4.1	Commercial projects	12
4.1.1	Skybox View	12
4.1.2	CycSecure	13
4.2	Research-grade projects	14
4.2.1	MulVAL	14
4.2.2	EDDY	16
5	Proposals for Research Directions	17
6	Beyond Defensive Posture	19
	References	20

List of figures

Figure 1: Flow diagram leading to defensive posture. 8

0

Acknowledgements

We wish to thank Eugen Bacic, Luc Beaudoin, Michael Froh, Marc Gregoire, Glen Henderson, and Reg Sawilla for helpful conversations.

1 Introduction

Modern networks are complex environments. The rapidity with which events occur, the wide variety of applications present, the interdependence of those applications, and other factors make networks difficult to understand, even if they have been carefully designed. This is particularly true with respect to the security of the network, for it is often not clear whether security is strengthened or weakened by changes in the network configuration, nor what network resources are threatened by existing security problems, nor even whether security problems exist in the network architecture at all.

This paper introduces the concept of dynamic network defensive posture. Roughly speaking, to know the dynamic defensive posture of the network is to know which elements of the network are exposed to potential attack by a malicious agent, and the extent to which such an attack would affect the network's operation. A defensive posture analysis system is an application that can determine the network's defensive posture.

A network administrator would benefit from having a current and thorough picture of the network's defensive status, and an awareness of how network events and management decisions affect security. What network resources are exposed to attack? For those exposed resources, what impact on network operations would a successful attack have? What would the consequences be, and how would the priority services offered by the network be affected?

The purpose of this 'way ahead' document for dynamic defensive posture is to discuss the concept of dynamic defensive posture, to describe the goals of a defensive posture analysis system, to discuss anticipated problems in achieving those goals, to identify sub-problems in need of clarification or solution, and to describe the long-term context into which the defensive posture project fits in the research program of the Network Information Operations (NIO) section of DRDC Ottawa.

1.1 Situational awareness

Situational awareness is a concept that has found application in many different areas (see [1] for a good review). In the context of network security management, situational awareness consists in a valid interpretation of the meaning of network activity, an understanding of its likely consequences for the provision of network services and enforcement of security policies, and the capacity to make informed network management decisions [2]. This requires a tightly coupled knowledge of network management and network security information, including knowledge of the network infrastructure and vulnerabilities. The primary aim is to be aware not just of the low-level details

of network events, but of their high-level impact on the operations and services which the network supports. To have situational awareness is to have a clear and correct picture of the network's state as it evolves in time and an understanding of how that state affects network services.

From a network management point of view, situational awareness requires that the network operator have a thorough and current model of the network's IT infrastructure, including network topology, available services, installed software, and host configurations. The operator must ensure that the network supports the missions and operations relying on it.

From a network security point of view, the network operator needs to know which network elements are exposed to attack by malicious users. This, too, requires knowledge of the network configuration, correlated with known software vulnerabilities. In fact, network security is closely related to network management, for any proposed network management decision must be evaluated in light of its security impact.

1.2 Critical resources

An essential element of situational awareness is a knowledge of the critical resources on the network. Critical resources are defined relative to the missions or operations which the network is supporting. They are the resources without which one or more important services or capabilities would be compromised. Although the assessment of the criticality of network elements is not the subject of this paper, we can remark that such an assessment would likely assign a value to each of the network elements, where the value reflects its importance in supporting the network mission.

1.3 Exposure

We have said that an element of situational awareness is knowledge of which network resources are vulnerable to attack. It is important to understand that attacks can consist of more than one stage: the attacker may move through multiple hosts and exploit multiple vulnerabilities or configuration problems before finally reaching the target. Also, such multi-stage attacks may in general originate from any point inside or outside the network. The set of network resources vulnerable to an attacker starting from a particular point in the network defines the *partial network exposure* relative to that point. The step-by-step description of how the attacker can carry out the attack is called an *attack path*. The *total exposure* is specified by giving all vulnerable targets and the possible attack paths to those targets.

One could also consider subsets of the total exposure, such as the outsider exposure, defined as the set of targets vulnerable to an attacker originating from outside the

network being defended, together with a catalogue of all possible attack paths to those targets.

1.4 Defensive posture

The defensive posture of a network is the set of *exposed, critical resources* on the network. It combines knowledge of the critical resources with the exposed resources. When the defensive posture is known, the network operator knows the degree to which critical network assets are vulnerable to attack. From an operational point of view, the defensive posture is a current, prioritized catalogue of existing security weaknesses requiring attention.

It is important to understand the relationship of defensive posture to risk, for they are not identical. Defensive posture states which of the organization's most critical assets are vulnerable to attack, but it does not say how *probable* an attack would be. Risk, on the other hand, folds into defensive posture an estimate of the probability of a particular attack being launched against the network.

Risk is often used to prioritize vulnerability instances. An attack could be possible against a critical asset, but if the probability of the attack is low the actual risk to the network is low, and the network operator may decide to focus his efforts on other problems. A problem with using risk to prioritize is that often the probability of an attack is difficult to specify precisely. There are cases where the probability of a particular attack is known to be high - such as when an automated worm is propagating on the Internet - and there may, less frequently, be cases where the probability is known to be low, but more often the probability is uncertain. As such, the use of risk introduces an intrinsic uncertainty to the prioritized list of threats against the network, possibly distorting that list if the probabilities are estimated wrongly. By contrast, defensive posture is assessed by combining knowledge of which assets are most important with an analysis of whether and how those assets could be attacked. In a well-designed system, neither of these elements should require guesswork, and the prioritized list of exposed, critical resources would be a reliable guide to security problems on the network.

1.5 Dynamic defensive posture

The network's critical resources may change with time in response to changing missions and operational priorities. At the same time, the network state can be altered by new software installations, the discovery of new vulnerabilities in existing software, changes to firewall rules, and other network events. Both types of changes affect the defensive posture. Consequently, defensive posture is a highly dynamic concept, evolving in response to changes in the network's structure and purpose.

The inherently dynamic nature of defensive posture distinguishes it sharply from a traditional network Threat and Risk Assessment (TRA). In a TRA, a one-time assessment of network security is made, either by inspection of the network architecture or by penetration testing. While perhaps valuable for identifying major architectural problems or security holes, a TRA is by its nature based on a static picture of the network. On a network for which configuration changes are frequent, and given that even a seemingly innocuous change can have serious unforeseen security consequences, the value of a TRA is doubtful even a short time after it is completed. Defensive posture improves on a TRA by updating frequently in response to changing conditions.

2 Requirements for Defensive Posture

To further clarify the scope and requirements of defensive posture, it is helpful to consider certain questions which the defensive posture should allow the network operator to answer. We first consider questions that pertain specifically to the main components of defensive posture - namely, criticality of resources and network exposure; a system that can determine the defensive posture should also be able to answer these questions. We then consider the questions that can be answered when knowledge of critical resources and network exposure are combined.

2.1 Questions about critical resources

Which assets are most critical to the network's mission?

The network exists to provide services and information in support of some set of goals or missions. For a particular set of goals, it is reasonable to expect that certain assets - whether services, information, or devices - will be critical to success. For instance, on a network supporting classified communications, encryption services are essential. The criticality of each asset is always defined relative to the goals and priorities of the mission. A significant challenge for a defensive posture analysis system is to accept a high-level, prioritized description of the mission's requirements, and map that description onto lower-level network services, information, and devices.

2.2 Questions about exposed resources

Which assets are exposed to an attacker placed at a particular location?

We have said in section 1.3 that the exposure is the set of assets that are vulnerable to attack, and is always defined relative to a particular starting point. After all, an attacker who starts with administrator privileges on a central network server can likely attack more targets than an unprivileged attacker on the Internet. We should be

able to specify which assets are accessible to an attacker who begins at any particular location, whether inside or outside the network. A full specification of ‘location’ may also include the attacker’s privileges, since privileges can affect the attacker’s reach. The exposure includes not only those assets directly accessible to the attacker, but also those which can be reached by multiple steps or stages.

From which network locations is asset X exposed to attack?

This question is the converse of the previous one, and identifies the set of users who could compromise the asset. Compromise might mean gaining unauthorized access in the case of a file, gaining execution privileges in the case of an application, or compromising availability in the case of a service. Again, the ‘network location’ might include a hostname, an account, and its associated privileges.

How could asset X be compromised by an attacker at a particular location?

Here we ask for more detail than in the previous questions. We want to know not only whether an attack is possible, but also *how* it is possible. The answer should include a set of attack paths from the starting position to the target. Each attack path is a step-by-step account of how the attacker could move from their initial location to the asset under consideration, possibly by exploiting vulnerabilities. Knowledge of the attack paths is important because it may help the network operator to mitigate the potential attack by making changes to the network configuration; simply knowing an attack is possible without knowing how does not provide such insight. It may even be possible to design a system that analyzes the attack paths and determines the course of action that should be taken to block the attack, while still providing all essential services (see section 6), but this is beyond the scope of this paper.

For any particular attack path, a number of subsidiary questions can be asked:

(a) How direct is the attack? The simplest measure would be the number of steps in the attack path, but a more nuanced approach might weight each step based on the perceived difficulty of completing it successfully. For instance, a step in an attack path which calls for exploiting a vulnerability using widely available exploit code could be considered more ‘direct’ than one which exploits a vulnerability for which no exploit code is known to exist. The former course is more probable, and therefore more direct. The directness of an attack might be a relevant consideration if one was trying to sort the attacks based on likelihood.

(b) What vulnerabilities make the attack possible? The rate at which software vulnerabilities are discovered makes it highly unlikely that the network will be entirely free of them, yet each vulnerability is a weakness in the network security. If the network operator knows which vulnerabilities are permitting an attack, he can focus his efforts on correcting the problem.

Which safeguards are protecting asset X?

The network operator may have an interest in knowing which safeguards are currently protecting a particular asset from attack. Safeguards in a network are devices, applications, policies, or configurations which prevent unauthorized activities. Examples are firewalls and routers, which can block unauthorized traffic, or authentication technologies, which prevent users from gaining unauthorized access to services or hosts.

Safeguards are obstacles between the attacker and the target asset, and, like exposure, these obstacles can only be identified with reference to a particular attacker's starting location. In fact, the question of safeguards is closely related to that of exposure, and we might rephrase our question in this way: Which safeguards, *were they absent*, would cause asset *X* to be exposed to an attacker at a particular location? Those safeguards are protecting the asset from that attacker.

This rephrased question also suggests a possible method for discovering the safeguards protecting asset *X*: assuming we have a model of the network, and an analysis technique that determines whether, given that model, the asset is vulnerable to attack, we need only run the analysis technique on an altered model in which one or more safeguards have been removed. Because of the potentially large number of permutations of safeguards in the model, the attack path discovery algorithm would need to be reasonably fast for this method to be practical.

A further refinement of this question is to ask about the *effectiveness* of the safeguards that are in place. The effectiveness of some safeguards could be affected by network conditions, as when high traffic volume causes a firewall to miss inspection of a malicious packet, while others could become less effective over time, as when passwords are vulnerable to brute force cracking. A general scheme for characterization of safeguard effectiveness is an outstanding problem.

2.3 Questions about exposed, critical resources

Defensive posture identifies the exposed, critical resources on the network. Having discussed the preliminary questions concerning critical resources and exposed resources independently, we are now in a position to consider the questions that arise when the two are brought together.

What are the most critical assets vulnerable to attack?

Given a list of network mission priorities, we (a) determine the most critical assets, (b) whether and how they are vulnerable to attack, and (c) produce an ordered list of exposed, critical assets. Steps (a) and (b) have already been discussed in sections 2.1 and 2.2 above, so the only new requirement is to rank the attacks with

respect to severity. The ranking may be ordered simply based on the criticality of the vulnerable asset, or according to some more elaborate scheme that also considers other factors, such as attack complexity, directness, probability of success, and so forth. The precise mixture of criteria by which a list of attacks should be ranked has yet to be determined, but certainly the criticality of the targeted asset will be the central, if not sole, consideration.

What are the expected consequences of a given attack?

If a particular asset is vulnerable to attack, we want to know what the impact of a successful attack would be. The impact is one important criterion relevant to assessing the severity of the attack. An attack that completely compromises an essential service, or even affects other services running from the same hardware, for instance, is more severe than one which partially compromises the service. An outstanding problem, however, is how to precisely specify both the type and extent of impact. One way to characterize the type is in terms of confidentiality, integrity, and availability. If the asset is a sensitive file, would the attacker be able to read it (thus compromising confidentiality), change it (compromising integrity), or even delete it (compromising availability)? It is often less clear, however, what compromising confidentiality or integrity means for a service or a device. As for the extent of impact, it is again not clear how this can be described precisely enough to be meaningful, but generically enough to be implemented in an automated system. This is a matter requiring further clarification.

3 Elements of Defensive Posture

In the previous section we considered defensive posture from the point of view of the network operator, exploring the questions that the defensive posture would answer. In this section, we turn to more technical considerations. In particular, we discuss what information must be gathered and combined to make a determination of defensive posture possible. We also discuss the problem of data collection.

To understand the various inputs required, refer to the flow chart in Figure 1 [2]. The output of triangle 4 in this diagram is defensive posture. We can see that we require as input knowledge of the operational requirements, IT infrastructure (ITI), network safeguards, known software vulnerabilities and exploits. We consider each of these items in turn.

3.1 Operations requirements

The operations requirements are high-level, prioritized descriptions of the services and information the network must support or provide. They are specified by a force

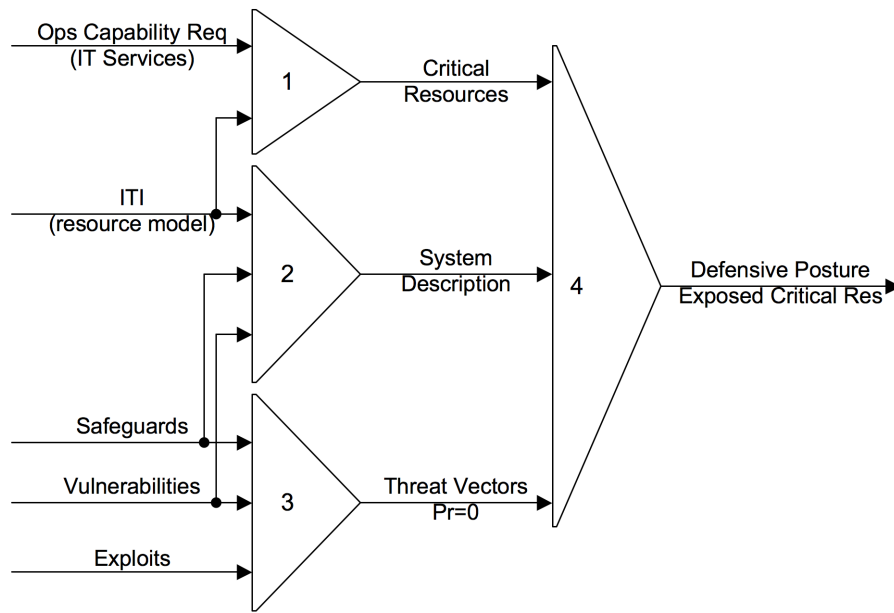


Figure 1: Flow diagram leading to defensive posture.

commander, and are derived from the mission of the organization or group using the network. They may be dynamic. For instance, the commander may require VoIP services and encrypted email services with high priority, and access to an image server with medium priority, for the duration of a mission, and also require video conferencing services with high priority on the morning of one particular day. In practice, the operations requirements would probably be defined by providing a template of available services to the force commander, and having him indicate which are needed, when they are needed, and with what priority. He, or the network commander who decides which network resources can best meet the requirements, should also indicate which services require redundancy in case of failure. The commander should be encouraged to be soberly realistic about prioritization of network services to avoid having everything marked as high priority.

3.2 IT infrastructure (ITI)

The IT infrastructure (ITI) is a crucial element needed to determine the defensive posture. To accurately identify the network's exposed resources, we require a thorough model of the network's structure and state. Knowing the ITI involves knowing how many hosts are on the network, their connectivity, what operating systems and software each host has installed, and how that software is configured; it involves knowing what servers are present on the network, the services they provide, and the interdependencies among these services; it includes knowing the access controls

present on files and applications, and the access permissions granted to users; it involves knowing the configuration of network firewalls and routers, so as to determine the network connectivity and allowed traffic flow. This information must be collected and used to construct a network model on which analyses, such as searches for attack paths, can be performed.

Automated collection of the data needed to populate a thorough and accurate network model is a challenging problem. Some of the data could be gathered by non-host-based scanners (examples are network scanners like nmap [3] or Nessus [4]), but much of it would only be accessible to host-based scanners. A host-based scanner is an application residing on the host which collects information about that host. For instance, a non-host-based scanner might determine that a particular machine is running an ftp server, but only a host-based scanner could identify with certainty the exact implementation and version number, the permissions with which the service runs, and the users who have permission to change the server's configuration. In an ideal situation, therefore, software would be installed on each host, server, and networking element that would collect information locally and securely transmit it to a central analysis machine, where it would be parsed and added to the network model. Regular scans would ensure that the network model is current with the actual network. Widespread deployment of host-based scanners, however, would require a management system to handle deployment, updates, and triggering of scans.

If one attempted to build a defensive posture analysis system using only non-host-based scanners to populate the network model, one would face many obstacles. Much information that would be valuable in evaluating the network's security would simply be unavailable. Since the conclusions of an analysis of the network are only as good as the information on which the analysis is based, these gaps in knowledge would seriously compromise the value of any network analysis. A useful and reliable determination of defensive posture requires current, accurate, and thorough knowledge of the network's state.

3.3 Software vulnerabilities

Some attacks against a network proceed by exploiting one or more vulnerabilities present in software on the network. These vulnerabilities are typically the result of defects in the design of the software, and, if exploited by a sufficiently capable attacker, may permit the attacker to obtain unauthorized access to hosts (as in a privilege escalation attack), adversely affect network performance (by crashing a server), or otherwise compromise the confidentiality, integrity, and availability of network resources.

Ideally, no such vulnerabilities would be present on the network, but realistically this ideal is unattainable. New vulnerabilities in deployed software are always being

discovered, and there will often be a gap between the announcement of a vulnerability and the availability of a patch. Even after a patch becomes available, administrators are sometimes prevented from applying it immediately. If the organization's policy permits users to install their own software, we must expect that users will frequently fail to keep their applications patched. In other words, we must always expect to have vulnerabilities on the network.

In recent years, organizations have been formed to catalogue and track known software vulnerabilities. Public databases such as the National Vulnerability Database (NVD) [5] maintain information about each known vulnerability, such as the affected software (and version), patch availability, exploit availability, the potential impact of a successful exploit, as well as various characteristics of the vulnerability itself, such as whether it is remotely exploitable, whether the attacker requires authentication, and so on.

This publicly available and maintained information about vulnerabilities is of great value in identifying exposed resources on a network. If the vulnerability information is correlated with the IT infrastructure data discussed above, we can identify the vulnerabilities present on the network. This is an important step toward discovering attacks against network targets.

3.4 Exploits

Anyone responsible for defending a network needs to be aware of the threats against them. Part of this awareness is provided by knowledge of the vulnerabilities in the network defences, as discussed above, but this awareness is complemented by knowledge of existing exploit code which targets those vulnerabilities. A particular exploit will typically be designed to target one or more vulnerabilities. Given the set of known vulnerabilities on a network, therefore, we are interested in the corresponding set of exploits which target those vulnerabilities.

Several projects, such as the Common Malware Enumeration project [6] and the Metasploit project [7], maintain public catalogues of information on exploits and exploit code.

It is perhaps debatable whether knowledge of exploits is essential when determining the defensive posture. When identifying the exposed resources, one is certainly interested in which vulnerabilities can be exploited to reach a target, but it is less clear what is gained by knowing about the specific program that does the exploit. One possible reason for wanting to know about specific exploit programs is that signatures of those programs could be identified and used by safeguards, like Intrusion Prevention Systems (IPS), to block the attacks.

Knowledge of specific exploits becomes more useful when we consider the problem of ranking attack paths with respect to risk, for here our knowledge of the properties of the exploit can be helpful. It may be important, for instance, to know that exploit code for a particular vulnerability is widely available on the Internet, since this increases the likelihood of an attack. The more prevalent or accessible is an exploit for a particular vulnerability, the more attacks which rely on that vulnerability will command the attention of the system administrator. Some thought needs to be given to the manner in which the availability and prevalence of exploits should affect the assessment of risk.

3.5 Safeguards and security policy

Safeguards are the network attributes which prevent attacks from being successfully carried out. They include authentication systems, access controls on files and applications, software patches, cryptographic algorithms, integrity checkers, firewall rules, and other network attributes.

It is perhaps surprising how many network attributes can be interpreted as safeguards. The configurations of the routers and firewalls, which determine the traffic patterns that can flow on the network - and so, in a sense, determine the shape of the network itself - are safeguards. A feature as pervasive as file access control is a safeguard. If we were to imagine the network without safeguards, it would be an environment in which all communication was allowed (within limits imposed by the network architecture) and all network entities, such as users and applications, could access any file. In other words, the safeguards so strongly determine the properties and behaviour of the network that it is, arguably, not profitable to consider them as a separate class of network attributes.

The set of network attributes we call 'safeguards' is closely related to that set of attributes which the network administrator can alter in order to affect the security of the network. Indeed, these two sets of attributes may be identical. Consider a network administrator confronted by an attack path against the network. There are potentially many courses of action which could reduce the risk of that attack: shutting down the vulnerable machine, patching the vulnerable service, blocking malicious traffic, adjusting file permissions, imposing an authentication requirement, and so on. Any course of action that improves the security of the network, however, could be interpreted as the introduction of a safeguard. Conversely, any course of action that degrades the security of the network can be interpreted as the removal of a safeguard. On this account, nearly any aspect of the network configuration could be interpreted as related to safeguards. Whether this is a sensible way to think about the problem is open to question.

The network's *security policy* abstracts the network's safeguards. The security policy

states, at a higher level, the principles or rules that the low level network safeguards enforce. A best practice for security policy is that critical assets be protected by dedicated safeguards, for instance. Policy might also specify the permissions granted to various groups of users. One technique that could be used to search for security weaknesses would be to search for violations of the security policy.

Safeguards remain a perplexing element in the defensive posture model. Accounting for them is clearly critical to the identification of exposed network resources, but how they should be included in the model is not entirely clear. Indeed, at some level it is difficult to say decisively what is, and what is not, a safeguard. For practical reasons, therefore, it may be necessary to sharpen the definition of a safeguard in order to limit the category to a more manageable scope. This topic is discussed again in section 5.

4 Related Work

We know of no existing project that is able to identify the defensive posture of a network. There are, however, a number of projects that address elements of the problem. In this section, we briefly discuss several of the most relevant and assess their suitability for application to the defensive posture problem.

4.1 Commercial projects

4.1.1 Skybox View

Skybox View is a commercial product that aims to evaluate network risk by identifying network resources that are vulnerable to attack [8]. It gathers data into a virtual model of the network, imports publicly available data about vulnerabilities, discovers multi-step attacks that could be launched against the network, and permits a network administrator to test the security consequences of network configuration changes before implementing them on the live network. In these respects, it is a tool for identifying exposed resources.

Skybox View also permits some reasoning about resource criticality. It is possible to classify resources according to monetary value, risk scale, or by indicating the confidentiality, integrity, and availability risk associated with them. This classification is done manually, and is therefore not suitable for dynamic re-valuation of resources. Given a classification, Skybox View will try to rank attack paths based on the perceived risk to the organization.

All of the information used to construct the network model is imported from external sources. Skybox View imports the network configuration, for instance, from firewalls and routers. It relies on network scanners to gather information about hosts, and

does not appear to use host-based scanners. Skybox View has a graphical interface, including network maps and overlaid visual representations of attack paths.

Shortcomings

A central question about Skybox View concerns the performance of its attack path discovery algorithm. The algorithm is not known. Performance data is not available, nor is it known how well the system scales to large networks.

Though it has the functionality to identify exposed resources by generating attack paths through a network, there is some question of how comprehensive the set of attack paths are. We noted that Skybox View uses network scanners, rather than host-based scanners, to gather information about the network; this limits the information that it can gain about each host, and it is unclear, therefore, whether it can reason about attack steps which are internal to a host or server [9].

Skybox View has some ability to reason about asset value, and uses assigned values to rank attack paths in order of severity. However, the assignment of asset value is a manual, labour-intensive process, and is therefore not suitable for environments in which asset values and network priorities are highly dynamic.

In addition, one must consider that since SkyBox View is a commercial product the source code is not available for customization or experimentation.

4.1.2 CycSecure

CycSecure is a commercial network risk management and network monitoring tool based on the Cyc Artificial Intelligence project [10, 11]. It builds a model of the network, and uses the Cyc knowledge base (KB) of basic facts and reasoning rules to deduce multi-step attacks. It permits network administrators to test proposed changes before making alterations to the live network, and can also model the interdependencies of network elements, which may be important for assessing the impact of an attack on other systems and services.

The reasoning model of CycSecure is very detailed, with, for example, over 350 classes of software faults defined, and nearly 700 classes of vulnerabilities. It is able to reason about a wide variety of attacks and scenarios, including power outages and social engineering attacks. Because it has such a large knowledge base on which to draw, it can reason very thoroughly about network attacks. The set of attacks it finds in a network is likely more comprehensive than for the other projects under discussion.

The system collects information about the network using a combination of network scanners and host-based scanners (called Sentinels), which gather data about in-

stalled and running applications, privileges, users, and so forth. Information about vulnerabilities is gathered from public sources.

Shortcomings

CycSecure's knowledge base is comprehensive, but this very comprehensiveness means that the system is quite complex to use. Although the tool obtains data about vulnerabilities from public sources, no public source gives information sufficiently fine-grained to suit CycSecure's detailed representation. Therefore new vulnerabilities require a trained CycSecure 'ontologist' to add them to the knowledge base. The company offers training in the Cyc data model [10].

It is not clear whether CycSecure allows value to be assigned to network assets [9]. If it does not, it may prevent one from specifying critical assets; if it does, it is not clear how the values can be derived from network priorities, nor how dynamic the asset values can be.

Some performance data for CycSecure is available [11]. Generation of attack paths takes between several minutes and several hours, depending on the complexity of the paths that are present. The performance is most sensitive to the number of discovered attack paths; it is not known how well the system scales to large networks.

In those areas where CycSecure falls short of the goal of dynamic defensive posture, it is not clear whether it would be possible to convince the developers to add the desired features. Several attempts to contact the company to obtain more information about the product have been unsuccessful.

4.2 Research-grade projects

4.2.1 MuIVAL

The Multi-Host, Multi-Stage Vulnerability Analysis (MuIVAL) project is a logic-based system for reasoning about network security [12, 13]. In many respects it is similar to Skybox View: it accepts a description of the network, including network connectivity, installed software applications and operating systems, users and privileges, and security policy, and, correlating that data with known software vulnerabilities, deduces all possible multi-host, multi-stage attack paths to specified targets in the network. Implemented in Datalog [14], it uses a small set of reasoning rules to test whether network resources are exposed to an attacker with certain initial network privileges (possibly none). For those resources which are exposed, the system outputs a set of attack paths detailing the means by which the attacker could compromise them.

An attractive feature of MuIVAL is its ability to reason about hypothetical situations.

The network configuration is described by a set of Datalog assertions, and it is possible to alter that set - whether by insertion, deletion, or replacement - and reason about the resulting set of assertions. This could allow a network administrator to evaluate the security implications of a proposed network change before making any changes on the live network. It could also allow the administrator to identify the systems which would become vulnerable if a new vulnerability with certain properties were discovered, making it possible to plan in advance for certain contingencies.

The data collection aspect of MulVAL is presently handled by the host-based Open Vulnerability Assessment Language (OVAL) scanner [15], an open-source application that collects detailed information about the software installed on its host. As discussed in section 3.2 above, host-based data collectors do require a significant effort to manage and maintain, but it is equally true that they can uniquely provide the precise, low-level data that is needed for a thorough and correct analysis.

A critical consideration for any defensive posture analysis system is performance. While previous efforts at identifying attack paths in networks executed in exponential time, MulVAL has been shown to execute in polynomial time [16]. The first stage of MulVAL analysis, which is the identification of valid attacks, has $\mathcal{O}(k^2)$ complexity, where k is the number of hosts in the network; the second stage of analysis, which is the generation of attack paths, executes in better than $\mathcal{O}(k^3)$. Network size is the main factor that affects performance; the system is relatively insensitive to the number of vulnerabilities in the network [16].

At the present time, MulVAL has a text-based interface, and can output attack paths in simple, text-based HTML.

As a candidate for an application to determine defensive posture, MulVAL is a strong contender. It can reason powerfully and efficiently about a network to identify exposed resources.

Shortcomings

MulVAL has not been designed to identify nor reason about criticality of network resources. As such, it does not currently support assignment of asset value or specification of essential services, nor does it provide a means to determine the impact of a successful attack (in terms of confidentiality, integrity, and availability) on the network mission. Though it produces a set of attack paths, it does not currently sort those paths in order of severity.

It should be remembered, however, that MulVAL is a young, open-source project that is still very much under development. The opportunity exists to contribute to the project in order to bring it closer to the goals of dynamic defensive posture. Indeed, collaboration is already taking place between the creators of MulVAL and DRDC

Ottawa.

There is one problem, discussed in section 2.2 above, that may present a particular difficulty for MulVAL - namely, how to identify the safeguards protecting a particular asset. Given a particular starting state (such as an attacker on the Internet) and a finishing state (such as an attacker gaining root access to a machine), MulVAL can efficiently search out all possible sequences of steps that move between the two states. It does this by iteratively searching all possible paths, and retaining only those that are successful. Typically, the number of successful paths is far fewer than the total number of paths attempted. There could be many reasons why a path fails to reach the goal, and while some such reasons will be related to the presence of safeguards, others will not. At present it is not clear how to effectively distinguish between the two cases. Moreover, because MulVAL's Datalog interpreter is designed to output successful paths, it is not clear that it can be made to indicate why paths fail to reach the goal. This is a technical problem which further work may resolve.

4.2.2 EDDY

The Event-Driven DiscoverY (EDDY) tool is a prototype developed at DRDC Ottawa for network exploration and monitoring [17]. The application collects information about the network using a variety of tools, and stores that information in a World Model. It is designed to respond intelligently when it discovers interesting new facts about the network; it may, for instance, upon discovering that a port has opened on a server, try to discover the version of the software listening on that port. It gradually builds up a picture of the network environment, and corrects that picture when it detects changes.

EDDY is able to translate its World Model into Datalog assertions and then invoke the MulVAL reasoning engine to identify possible attacks. Since it relies on MulVAL to identify exposed resources, EDDY inherits most of the strengths and weaknesses of MulVAL. It has the advantage of being able to gather information without relying on host-based scanners; however, when it operates without such scanners, the quality of its data suffers.

Shortcomings

Since EDDY relies heavily on MulVAL for the identification of exposed resources, it inherits most of the weaknesses of MulVAL that we have already identified. It is possible that some of the defensive posture requirements missing from MulVAL, such as asset valuation or attack path ranking, could be implemented by EDDY instead of being incorporated into MulVAL itself.

5 Proposals for Research Directions

In this section, we draw on the preceding discussion and try to give concrete proposals for research work that would contribute to the solution of the outstanding problems facing the development of a dynamic defensive posture model.

- *Network modeling.* The model of the network is the foundation for the analysis of exposed and critical resources and is, of course, a critical part of a defensive posture model. Several important questions must be addressed in the design of the model: What elements of the network must be modeled, and what elements can be ignored? At what level of granularity should the network be modeled? The answers to these and related questions will depend on the range of attacks that the model attempts to identify. A minimal model would likely include privilege escalation attacks; a more comprehensive model might also include denial of service attacks, eavesdropping and sniffing attacks on data in transit, or data tampering; more complex still would be models of social engineering attacks, or attacks on the physical network infrastructure. A good model would also permit the network administrator to evaluate the effects of changes to the network configuration before the changes are actually made; alterations could be made in the model, and the security consequences evaluated prior to deployment.
- *Asset valuation schema.* Essential to a quantitative assessment of asset criticality, and potentially also to assessment of a successful attack's impact on the network, is some means of assigning value to the network assets. The value in question, whether expressed numerically or categorically (for example, on a High-Medium-Low scale), should indicate the level of support the asset provides to the meeting of the network's current priorities. It may also indicate the nature of the support provided, or, equivalently, the negative impact on the network that would result from the asset's loss or compromise. This impact measure would likely be specified in terms of the conventional metrics of confidentiality, integrity, and availability.
- *Network priorities schema.* To identify the network's critical resources, we require some indication of the scale of importance of the various network services to the network's mission. It is therefore necessary to develop a standard means to describe the mission-critical network services and information. At what level of detail this description should be given - as a high-level quality of service requirement, for instance, or as a lower level technical requirement - is an open question. At whatever level of detail the priorities are specified, we require a technique for mapping those priorities down to the low-level network infrastructure (see the following item).

- *Rapid valuation of assets in response to network priorities.* Once a scheme for indicating the value of a network asset has been developed, one faces the problem of assigning value to network resources in response to a set of network priorities. This ‘trickle-down’ problem involves mapping a relatively high-level description of network priorities onto the network infrastructure that supports those priorities. For example, if secure e-mail service is a high priority, one would expect the e-mail server, its associated cryptographic applications, and the hardware on which the services run to be assigned a high value. The process of asset valuation must be automated, such that the network assets can be re-valued as changes occur in either the network priorities or the network infrastructure itself. The automated process should be rapid, at least compared to the pace at which network priorities and infrastructure change.
- *Competing network priorities.* We have said that the value of network resources should reflect their role in serving the network priorities, as specified by the network administrator or commander. A problem arises, however, when two or more sets of competing priorities are being supported by the network. Thought needs to be given to how the various claims on network resources are to be weighed, and a scheme developed for resolving conflicts.
- *Attack path ranking.* When a set of exposed network resources has been identified, together with the attack paths by which an attacker might reach each resource, it is desirable that the various attack paths be ranked in order of severity. Such a ranking can guide the system administrator’s efforts to secure the network. Research must investigate the appropriate criteria for ranking attack paths, and the manner in which those criteria should be combined to produce a quantitative ranking.
- *Safeguards and safeguard effectiveness.* As we have discussed, a good understanding of safeguards is essential because the safeguards in large measure determine which network resources are exposed. The wide variety of safeguard types makes it unlikely that a single safeguard model can be applied to all the kinds of safeguards present in a network. It is possible that a system of *safeguard classification* must be developed, into which safeguards of different kinds could be sorted, and each safeguard class would be modeled differently. Safeguards might be classified according to the kind of protection they provide (confidentiality, integrity, availability), the technical means they employ (encryption safeguards), the OSI layer at which they act, their effectiveness, the aspects of network activity they affect (traffic flow, file access), the vulnerability they mitigate, or according to some other criteria. Certain problems, such as how to determine which safeguards are protecting a network resource, require good modeling control over the safeguards. The effectiveness of safeguards, conceived either as the ability of a safeguard to withstand attack or as the extent

of the protection the safeguard provides to the network, is an issue that may prove relevant to defensive posture.

- *Rapid reassessment of defensive posture.* A system that can identify the exposed, critical network resources would be a major achievement. We also require, however, that this identification be completed quickly. The greater the complexity of the network model, and the wider the scope of attack types which the system can discover, the greater the challenge of producing a system with acceptable performance. To be of practical use, the system must reason about a broad enough scope of possibilities to give a reliable picture of the network's defensive status, it must be able to reason about relatively large networks, and it must do so within a reasonable time frame. A minimal requirement would be that the network analysis to determine the exposed, critical resources not take longer than the interval between data collection cycles. This requirement should influence the design of the algorithms for identifying both the critical and the exposed resources, and may exclude some algorithms from consideration.
- *Incomplete data.* The defensive posture model should reason on data gathered directly from the deployed network. In some circumstances, such as when a system is prevented from reporting or a new system appears on the network without a host-based scanner installed, all of the desired information may not be available. In such cases, one option is to simply use what one has, aware that some possible network attacks could be missed. Another option is to make an educated guess, based on past experience or convention, about the missing data. It may be possible to formalize the treatment of tentative data using the theory of evidence [18, 19].
- *Visualization of defensive posture.* Recognizing that a clear, informative presentation of the information about exposed, critical resources to the network administrator is important if the tool is to be useful, possible schemes for visualizing the network, the criticality of the various resources, and the attack paths against vulnerable resources should be investigated.

6 Beyond Defensive Posture

The effort to build a system capable of determining dynamic defensive posture is one stage of a larger research program. In this section, we briefly discuss several longer-term research goals, all of which rely to some extent on the success of the dynamic defensive posture project.

In addition to producing an ordered list of exposed, critical resources, it may be desirable to produce a corresponding *risk assessment*. A risk assessment takes the

set of possible attacks against the network combined with the probability that each attack will occur, and distills from it a simple indicator of overall risk, such as a numerical value or an alert level. The method by which this simplification is carried out is a matter for future research.

It is one thing to give a network administrator a list of possible attacks against exposed, critical resources; it is another to give the administrator *Course of Action advice*. It may be possible, by analyzing the attack paths, to determine what changes should be made to the network configuration in order to block severe attacks, and to thereby reduce the network risk. For instance, perhaps a set of attacks all exploit a particular vulnerability; a Course of Action advice application might propose patching that vulnerability, or taking the vulnerable system off-line. Such a system could be of considerable value to a network administrator.

Even more ambitious would be to design a system that could identify and repair security problems without the manual intervention of the network administrator. Such an *Automated Response* system would build on the Course of Action research, but have the recommended action implemented automatically. Needless to say, granting to an application permission to automatically alter the network configuration is potentially dangerous, and the conditions under which an automated response could be permitted would need to be carefully specified. A sufficiently developed Automated Response technology would make the design of *Self-healing Networks* a genuine possibility.

References

- [1] Breton, Richard and Rousseau, Robert (2003), Situation Awareness: A Review of the Concept and its Measurement, (DRDC Valcartier TR 2001-220), Defence R&D Canada – Ottawa.
- [2] Lefebvre, Julie, Gregoire, Marc, Beaudoin, Luc, and Froh, Micheal (2005), Computer Network Defence Situational Awareness Information Requirements, (DRDC Ottawa TM 2005-254), Defence R&D Canada – Ottawa.
- [3] Nmap - Free Security Scanner for Network Exploration and Security Audits (Online), <http://insecure.org/nmap/> (Access Date: October 19, 2006).
- [4] Nessus (Online), <http://www.nessus.org/> (Access Date: October 19, 2006).
- [5] National Vulnerability Database (Online), <http://nvd.nist.gov> (Access Date: October 19, 2006).
- [6] Common Malware Enumeration (CME) (Online), <http://cme.mitre.org/> (Access Date: October 19, 2006).

- [7] The Metasploit Project (Online), <http://www.metasploit.com/> (Access Date: October 19, 2006).
- [8] Skybox Security (Online), <http://www.skyboxsecurity.com/> (Access Date: October 19, 2006).
- [9] Bacic, Eugen, Froh, Michael, and Henderson, Glen (2006), MulVAL Extensions for Dynamic Asset Protection, (DRDC Ottawa CR 2006-251), Defence R&D Canada – Ottawa.
- [10] CycSecure Homepage (Online), <http://www.cyc.com/applications/cycsecure> (Access Date: October 19, 2006).
- [11] Shepard, Blake, Matuszek, Cynthia, Fraser, C. Bruce, Wechtenhiser, William, Crabbe, David, Güngördü, Zelal, Jantos, John, Hughes, Todd, Lefkowitz, Larry, Witbrock, Michael, Lenat, Doug, and Larson, Erik (2005), A Knowledge-Based Approach to Network Security: Applying Cyc in the Domain of Network Risk Assessment, In *17th Innovative Applications of Artificial Intelligence Conference*, Menlo Park, CA, USA.
- [12] Ou, Xinming, Govindavajhala, Sudhakar, and Appel, Andrew W. (2005), MULVAL: A logic-based network security analyzer, In *14th USENIX Security Symposium*, Baltimore, MD, USA.
- [13] Ou, Xinming (2005), A Logic-Programming Approach to Network Security Analysis, Ph.D. thesis, Princeton University.
- [14] Ceri, Stephano, Gottlog, Georg, and Tanca, Letizia (1989), Everything You Always Wanted to Know About Datalog (But Never Dared Ask), *IEEE Transactions on Knowledge and Data Engineering*, 1(1), 146–166.
- [15] Open Vulnerability and Assessment Language: OVAL Interpreter (Online), <http://oval.mitre.org/language/download/interpreter/index.html> (Access Date: October 19, 2006).
- [16] Ou, Xinming, Boyer, Wayne F., and McQueen, Miles A. (2006), A Scalable Approach to Attack Graph Generation, In *13th ACM Conference on Computer and Communications Security*, Alexandria, VA, USA.
- [17] (2006), Dynamic Asset Protection - Prototype Software Development Report, Cinnabar Networks Inc.
- [18] Dempster, Arthur P. (1968), A generalization of Bayesian inference, *Journal of the Royal Statistical Society, Series B*, Vol. 30, 205–247.

- [19] Shafer, Glenn (1976), A Mathematical Theory of Evidence, Princeton University Press.

Distribution list

DRDC Ottawa TM 2006-250

Internal distribution

- 3 Author
- 1 Section Display
- 4 Library electronic
- 1 electronic Marc Grégoire electronic
- 1 electronic Reg Sawilla

Total internal copies: 10

External distribution

- 1 Lt. Col. Chevalier (CO CFNOC), CFS Leitrim, 101 Colonel By Dr., Ottawa, ON K1A 0Z4
- 1 Maj. Castonguay (DCO CFNOC), CFS Leitrim, 101 Colonel By Dr., Ottawa, ON K1A 0Z4
- 1 Maj. Torrington-Smith (NOS CFNOC), CFS Leitrim, 101 Colonel By Dr., Ottawa, ON K1A 0Z4
- 1 Maj. Nickerson (L/CIRT CFNOC), CFS Leitrim, 101 Colonel By Dr., Ottawa, ON K1A 0Z4
- 1 Capt. Bendelier (L/NVAT CFNOC), CFS Leitrim, 101 Colonel By Dr., Ottawa, ON K1A 0Z4
- 1 Capt. Messous (L/AAT CFNOC), CFS Leitrim, 101 Colonel By Dr., Ottawa, ON K1A 0Z4
- 1 Sgt. Arndt (AAT CFNOC), CFS Leitrim, 101 Colonel By Dr., Ottawa, ON K1A 0Z4
- 1 Capt. Scheurkogel (A/IO CDI), 101 Colonel By Dr., Ottawa, ON K1A 0Z4
- 1 Luc Dandurand, Cyber Defence Futures (CSE), Edward Drake Building, 1500 Bronson Ave. Ottawa, ON K1G 3Z4

- 1 Annie DeMontigny-Leboeuf, Cyber Defence Futures (CSE), Edward Drake Building, 1500 Bronson Ave. Ottawa, ON K1G 3Z4

Total external copies: 10

Total copies: 20

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.) Defence R&D Canada – Ottawa 3701 Carling Avenue, Ottawa, Ontario, Canada K1A 0Z4		2. SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable). UNCLASSIFIED	
3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title). Dynamic Defensive Posture for Computer Network Defence			
4. AUTHORS (last name, first name, middle initial) Burrell, C.; Lefebvre, J.; Treurniet, J.			
5. DATE OF PUBLICATION (month and year of publication of document) December 2006	6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc). 36	6b. NO. OF REFS (total cited in document) 19	
7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered). Technical Memorandum			
8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include address). Defence R&D Canada – Ottawa 3701 Carling Avenue, Ottawa, Ontario, Canada K1A 0Z4			
9a. PROJECT NO. (the applicable research and development project number under which the document was written. Specify whether project). 15BO03		9b. GRANT OR CONTRACT NO. (if appropriate, the applicable number under which the document was written).	
10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique.) DRDC Ottawa TM 2006-250		10b. OTHER DOCUMENT NOs. (Any other numbers which may be assigned this document either by the originator or by the sponsor.)	
11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification) (X) Unlimited distribution () Defence departments and defence contractors; further distribution only as approved () Defence departments and Canadian defence contractors; further distribution only as approved () Government departments and agencies; further distribution only as approved () Defence departments; further distribution only as approved () Other (please specify):			
12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution beyond the audience specified in (11) is possible, a wider announcement audience may be selected).			

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

This paper examines the concept of dynamic defensive posture in computer networks. We propose that defensive posture is provided by knowledge of whether and how critical network resources are vulnerable to attack. After introducing the basic concept, we discuss the constituent elements of defensive posture. We then review relevant technologies, finding that current technology addresses only aspects of the problem. Finally, we propose a variety of research problems for which the solutions would contribute significantly to our ability to identify a network's defensive posture.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title).

Defensive Posture
Network Security
Situational Awareness

Defence R&D Canada

Canada's leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale



www.drdc-rddc.gc.ca