Enhancing response in intrusion detection systems

M.Papadaki, S.M.Furnell, S.J.Lee, B.M.Lines and P.L.Reynolds

Network Research Group, Department of Communication & Electronic Engineering,
University of Plymouth, Plymouth, United Kingdom
info@network-research-group.org

Abstract. *With rising levels of attacks and misuse, intrusion detection systems are an increasingly important security technology for IT environments. However, while intrusion detection has been the focus of significant research, the issue of response has received relatively little attention. The majority of systems focus response efforts towards passive methods, which serve to notify and warn, but cannot prevent or contain an intrusion. Where more active responses are available, they typically rely upon manual initiation. The paper examines the reasons for this, and argues that a more comprehensive and reliable response framework is required in order to facilitate further automation of active responses. A range of factors are identified that a software-based responder agent could assess in order to improve response selection, and thereby increase trust in automated solutions.*

Keywords: Intrusion Detection, Intrusion Response, Active Response, Automated Response.

INTRODUCTION

An increasing level of attacks upon IT systems represents a seemingly unavoidable reality of the Internet revolution. From the malicious activities of external hackers to deliberate misuse by organisational insiders, no sector has shown itself to be immune from attack, and the provision of a public-facing server is effectively all that an organisation needs to do in order to establish itself as a potential target. Evidence of the problem is provided by results from the annual CSI/FBI Computer Crime and Security Survey, which has seen the percentage of respondents reporting incidents rise from 42%, in 1996, to 64% in 2001 (Power 2001),

having reached an interim highpoint of 70% in 2000. The associated financial losses have also increased, and the 2001 survey results reported total losses approaching $378 million (from 186 respondents who were willing and able to quantify the financial impacts of their incidents). In fact, since 1997, the annual damage from security breaches has been increasing by an average of 52% every year.

Statistics such as those above emphasize the need for security in networked systems, and a key technique for combating attacks is provided by Intrusion Detection Systems (IDS). The concept of intrusion detection was originally proposed by Denning (1987), with the underlying rationale that the complicated infrastructures of computer and network systems are inherently insecure, and thus may be under attack. Pursuing the aim of creating totally secure systems may not be feasible or cost effective, so an intrusion detection system should be able to detect such attacks, preferably in real time. Since Denning's original work, the concept has received wide acceptance in the computer security domain, and several efforts have concentrated on the development of practical intrusion detection systems. There are many challenges in that process and, to date, the focus of research has been on the detection capability of systems (Mukherjee et al. 1994; 2000). However, the issue of response to detected incidents is another significant issue, but has so far been largely overlooked (Schneier 2000) and therefore requires further research in its own right.

The paper begins by introducing the concept of intrusion response, and considering the different approaches by which it may be realised. The response capabilities of current intrusion detection systems are then analysed, from the perspective of both commercial products and ongoing research projects. The need for further enhancement is identified, leading to the proposal of a broader response framework, and the identification of various contextual factors that need to be considered in order to select appropriate responses.

THE CONCEPT OF INTRUSION RESPONSE

Intrusion response can be defined as the process of counteracting the effects of an intrusion. In the context of intrusion detection it includes the series of actions taken by an IDS following the detection of a security-related event. It is important to note that consideration is

not only given to taking action after a full-scale breach has been detected, but also when events of interest take place and raise the alert level of the system (i.e. the early stages of an attack, when the system is suspecting the occurrence of an intrusion, but is not yet sufficiently confident to take action).

In general, the aims of response actions can be classified into one of the following categories:

1. Notification about the occurrence of an intrusion.

2. Protection of system resources:
   - in the short term, this will include mechanisms to contain the intrusion, as well as to recover and restore the system to a well known state.
   - in the longer term, it includes learning from the intrusion, using this knowledge to remove identified vulnerabilities of the system, and to enhance the detection and response capability. The objective here is to prevent reoccurrence of the intrusion.

3. Identification of the perpetrator of the intrusion.

At the highest level, there are two main approaches to intrusion response; namely human/organisational approaches and technical methods. The former are those that involve human processes and organisational structures, and may include actions such as reporting an incident to the police or invoking disciplinary procedures (e.g. in cases where internal personnel are responsible). From the list above, the process of identifying the perpetrator often requires further investigation and co-operation with other parties, such as Incident Response Teams, and thus it naturally falls under the human/organisational aspect of response. By contrast, technical responses involve the use of functional techniques and software-based methods. These technical actions can themselves be further sub-classified, into either passive or active forms of response (Bace and Mell 2001). Technical response actions can also be characterised as either manual or automated, according to the way they are initiated (Amoroso 1999). The main distinctions will now be considered in more detail.

Passive and Active Responses

Passive responses aim to notify other parties about the occurrence of an incident, relying on them to take further action. Passive responses may include methods such as:

- Recording details for later inspection (e.g. adding an entry in a log file);
- Alerting an administrator, by displaying a pop-up window on the console, or generating an email, pager or mobile phone message;
- Generating alarms and alerts to report to a central network management console by using SNMP traps and messages.

Passive responses, in the form of notifications and alerts, have been used by IDSs since their initial development, primarily as an indicator of their detection effectiveness. Hence they are still present in every intrusion detection product, offering the standard level of response, and making them the most common response option in commercial IDS systems. The obvious disadvantage here is that they do nothing to impede the intruder, and rely upon someone to manually respond at some later point (by which time it may be too late to avert a more significant security breach).

In contrast to the passive approaches, active responses are the actions taken to counter the incident that has occurred. Such actions might include the following approaches:

- collecting more information about the incident (e.g. issuing an authentication challenge, increasing the monitoring level);
- limiting permitted user behaviour or process activity;
- blocking network traffic through firewalls and routers;
- terminating network connections;
- introducing delay on network connections.

Active responses can have a more significant impact upon a system, and thus they engender the danger of causing unwanted effects, in the event that they are falsely initiated. In order to overcome this danger, careful consideration should firstly be given to the thoroughness and extensiveness of the response options available. It is also important to study the conditions under which the selection of appropriate responses is made. This requires consideration of the

factors that can influence the response decision process and assessment of their weighting upon that process.

Not surprisingly, active responses have mainly been used in research prototype systems. Although there are some commercial systems utilising active response methods, especially ones that involve blocking of network traffic and termination of network connections, their application is still in an early stage and their effectiveness has not yet been conclusively proven.

Manual and Automated Responses

The detection of a suspected intrusion typically triggers a manual intervention by a system administrator, after having received an alert message from the intrusion detection system. The IDS can additionally assist the incident response process, by providing the details of the attack, saved in a log file (Bace et al. 2001). However, responding manually to intrusions is not necessarily an easy task, as it can represent a significant administrative overhead. That may involve dealing with a high number of alerts and notifications from the IDS, ensuring awareness of security bulletins and advisories from incident response teams, and taking appropriate actions to resolve each of the alerts reported. From the system administrator's perspective, the main requirement is to ensure that the system remains operational and available – this is what the users expect and complaints will quickly occur if this is not the case. Unless resolving a reported incident is explicitly required to ensure that this is the case, then the task is likely to be given a lower priority.

The ability to mount a rapid response to an attack is, however, extremely important. The effect of reaction time on the success rate of attacks was demonstrated by Cohen, who carried out a simulation of attacks, defences and their consequences in complex cyber systems (Cohen 1999). The results indicated that if skilled attackers are given 10 hours after they are detected, and before a response is generated, then they will be successful 80% of the time. If they are given 20 hours, they will succeed 95% of the time. At 30 hours, the attacker almost never fails. The results also indicate that if a skilled attacker is given more than 30 hours, the skill of the system administrator will make no difference, as the attacker will irrespective of that succeed. On the other hand, if the response is instant, the probability of a successful

attack against a skilled system administrator is almost zero. This proves that there is a relationship between the effectiveness of response and the time it is issued, and that there is a window of opportunity for an attacker if response is not issued on time.

Another factor that highlights the need for automated response is the changing nature of the techniques employed by attackers, including the widespread use of automated scripts to generate attacks of distributed nature (Cheung and Levitt 1997). These can further diminish the ability to respond manually, since there is practically no time available to do so.

At the time of writing, the degree of automation in current intrusion detection systems is very low, being largely limited to the automation of passive responses. Nonetheless, a feasibility-level research study has estimated that 33% of available response actions have the potential to be safely automated, without having to further enhance the detection capability of the IDS (Lee 2001). As such, this would have the potential to significantly reduce the burden upon system administrators.

RESPONSE CAPABILITIES OF CURRENT IDS

A literature search was carried out to investigate the response capability of current IDSs, focusing upon the systems with more interesting approaches to intrusion response. The aim of this task was to cover the most representative set of response options available, rather than reflect the degree to which automated active response has been adopted in the intrusion detection domain. Thus, consideration is given to the more significant response features of commercial IDS products available (Table 1) followed by the systems under research.

		Passive Response			Active Response				
IDS name	**NB / HB***	**Pop-up window alerting**	**E-mail, Pager alerting**	**SNMP (trap) alerting**	**Collect more information**	**Limit permitted user behaviour**	**Terminate / reset network connection**	**Block network traffic**	**Introduce Delay**
Axent Technologies NetProwler (Hurwitz Group 2000)	NB	✓					✓	✓	
Axent Technologies Intruder Alert (Shipley 1999)	HB	✓	✓	✓					
CentraxICE ICEpac, BlackICE (Gilliom 2001)	HB / NB	✓	✓	✓		✓	✓		
Cisco Secure (Cisco Syst. 1998)	NB	✓	✓		✓		✓	✓	
eTrust IDS (Computer Associates 1999; 2001)	NB	✓	✓	✓	✓			✓**	
ISS RealSecure 6.0 (ISS 2001)	HB / NB	✓	✓	✓		✓	✓	✓	

* NB / HB: Network-based / Host-based

** Blocking of network traffic is not done via reconfiguration of a router or firewall, but by using patent pending "***unobtrusive blocking***" based on pre-defined rules (*stateful dynamic blocking*) or in response to a specific alert.

Table 1: Response capabilities of current IDSs

As expected, nearly all Intrusion Detection Systems offer a wide range of passive responses, but the situation regarding active responses is somewhat more varied. There are systems that do not offer any active methods, while others that seem to offer a wide range of options. However, it seems that the active responses (terminate/reset network connections, block network traffic) available for network-based IDSs are more widely adopted than the ones fitted for host-based systems (limit permitted user behaviour), suggesting an opportunity for further enhancement.

In addition to commercial products, there is also a significant amount of active research in the IDS domain. As such, it is relevant to consider whether these have more advanced approaches in the context of response.

EMERALD

Event Monitoring Enabling Responses to Anomalous Live Disturbances (EMERALD) is an intrusion detection project being pursued within the Systems Design Laboratory at SRI International. Within the Emerald architecture, the Resolver is a countermeasure decision engine capable of handling the alerts from its associated analysis engines and invoking response handlers to counter malicious activity. The Resolver is an expert system that receives the intrusion and suspicion reports from the profiler and signature engines. Based on the combination of these with reports from other peer monitors, it decides which response to invoke, and how to invoke it. Possible responses may include direct countermeasures such as closing connections, terminating processes, or the dispatching of integrity-checking handlers to verify the operating state of the analysis target (Porras and Neumann 1997). EMERALD provides an interesting architectural approach, focusing on the co-operation of distributed response elements. However, although elements of the detection functionality have been realised in practice, the eResponder (the countermeasure invocation system) is still under development at the time of writing (SRI 2001), and the conceptual definition that has been published does not specify in detail the response mechanism, the actions available, or the operation of the decision engine.

Response and Detection Project

This research project is being pursued as a collaboration between Boeing Corporation, Trusted Information Systems, and UC Davis University (UC Davis 2000). It is an effort to combine state of the art intrusion detection systems with intelligent firewalls and routers to form an Intranet wide automated response system. The basic idea is to enable cooperation among response components in a virtual security network, where security components alert each other of an attack, and a component will be selected to initiate an automated response. Automated responses mainly examine network-based attacks, and at present are restrained to filtering network packets. However, focus has been given to the extension of response options, including options such as the introduction of delay into network connections and the replacement of sensitive files with look-alikes.

Adaptive, Agent-based, Intrusion Response System (AAIRS)

The AAIRS project is being pursued within the Computer Science Department of Texas A&M University (Carver et al. 2001). It focuses on the response decision mechanism, proposing a methodology for adaptive automated response using intelligent agents. In AAIRS, a new *Analysis agent* is created every time a new attack is reported by the IDS. In collaboration with other agents (for response and policy) an abstract response plan is developed, taking account of any appropriate legal, ethical, institutional, or resource constraints specified in the policy. The plan is then passed to a *Tactics agent*, which decomposes it into specific actions and invokes the appropriate components of a *Response Toolkit*. Both the *Analysis and Tactics agents* employ adaptive decision-making based on the success of previous responses, which serves to limit uncertainty in the response decision process, and facilitate adaptation of the system based on the effectiveness of its detection and response capability in the past (Ragsdale et al. 2001). For the response decision process, the following factors have additionally been taken into account (Carver and Pooch 2000):

- Timing of the attack (pre-emptive, during attack, after attack)
- Type of attack (threat to confidentiality, integrity, availability)
- Type of attacker (Cyber-gangs, Economic Rivals, Military Organisations, …)
- Degree of suspicion (high, …, low)
- Attack implications (Low implications, …, Critical implications)
- Environmental Constraints (No offensive responses, …, No router Resets)

Of the response-oriented work described in published literature, the AAIRS project is considered to offer the most comprehensive treatment of the issue to date, giving considerable focus to the mechanism, and the influencing factors, of the response decision process.

LIMITATIONS OF CURRENT INTRUSION RESPONSE METHODS

An important issue that was not reflected in Table 1 was the degree to which the response methods are automated. Although it is technically feasible to implement and automate many

forms of response in software, it is not as straightforward a solution as it might seem. Whereas passive responses have already been automated to a large extent, active responses are largely initiated manually. The reason for this is that passive actions have little impact upon a system, and thus there is no danger of causing damage if a response is initiated in a false positive alarm scenario. On the contrary active responses could cause disruption to legitimate users, affect their access level to the system or even cause an unintentional denial of service attack to the system itself. Hence we need to make sure that when response actions are launched automatically - without prior human authorisation - they have the intended effects and do not put the system in a greater risk than it currently is. This requires confidence in the detection capability of a system (i.e. that its assessment of a scenario as being intrusive is accurate), as well as its subsequent ability to choose appropriate countermeasures in response.

Even if automated response capabilities were to be made widely available, there is a question of whether they would actually have practical value in the eyes of security administrators, who may prefer to place trust in their own abilities rather than those of the system. This viewpoint is reflected in the results of an email survey conducted by Lee (2001), in which various IDS vendors and intrusion detection specialists were asked to comment upon the automated response issue. A number of relevant responses are presented below (having been anonymised to remove details of specific individuals or products). It can be seen that whereas some vendors do not support automated response, but at the same time do not exclude the option of doing so, other commentators can see a fundamental risk in the concept:

> *'We do not have any self-imposed automated proactive responses; we enable the creation of policies for response.'*

> *'We have the functionality for automated responses but we haven't yet explored fully what we should do to pro-actively deal with suspected intrusions.'*

> *'Our product at this time does not provide a particular response mechanism out of the box,… What I have found is that, at least in North America, many of the security professional prefer not to use automated response systems.'*

'We think it is dangerous to put all your faith in automatic responses believing that you are protected. Attackers are very smart and know how to use your own equipment against you if it will benefit their attack. …We believe that the ability to see if an attack was successful or not and then have a human acting on that is better for the overall health and security of the network. …most people in the security product industry agree that automatic responses can be very dangerous and should not be relied upon to make important decisions about your networks.'

'Proactive measures are a reasonable idea, unless they can be subverted. For instance, if you decide to shut down your network connections as a proactive approach, then an intrusion attempt can be used as a denial of service…'

'Retaliation (DoS for instance) is out of the question since one can't ignore the impact on innocent users coming from the same network (say with an ISP). Therefore, [DELETED] will continue to only support detection and reporting.'

'Right now, in its current form, I don't believe that the current products are mature enough to be performing active response. … any device that is re-configuring infrastructure equipment (shunning) could easily be turned into a denial of service tool.'

Although the negative comments above can be considered to offer a valid perspective, what they tend to overlook is the previous argument that, in many circumstances, manual response may not represent a viable alternative (e.g. in the context of attack via automated scripts). As such, it can be concluded that, in spite of the difficulties, efforts are required to improve the prospects of automated methods.

EXTENDING AUTOMATED RESPONSE

In order to address the automated response issue, it is considered that further attention is required in two main areas; namely the broadening of possible response options, and the

assessment of factors that influence the suitable response to be taken. Appropriate attention to these aspects will help to address the problem of reliability in relation to automated responses.

Broadening of response options

Further consideration should be given to novel response actions, which will possibly have less significant impact upon the system and its legitimate users, causing at the same time the desired effect upon the attack, and preserving the uncompromised state of the system. Intuitively, however, these requirements may be mutually exclusive, in the sense that responses that have minimal effects upon legitimate users may also have limited potential for safeguarding the system, and vice versa. For example, an attack against the confidentiality of the system could potentially be addressed by:

- delaying the disclosure of information, until an authentication challenge is issued;
- denying access to sensitive information by limiting permitted user behaviour;
- providing false information instead.

Each of these actions has different impact upon users and attacks. Delaying the disclosure of information, in order to issue an authentication challenge in the meantime, does not have significant impact upon the system, as the introduced delay could easily be disguised as usual system overload. However, the effect of the delay upon the attack is not significant either, as no action is taken to actually eliminate it. If the authentication challenge reinforces the suspicion of the system about the occurrence of an attack, it is possible to either deny access to the requested information, by limiting user behaviour, or provide false information instead. In both cases, the impact upon the system is the same, as the requested service is denied by the system. However the effect on the attack might be different, as the attacker who unsuccessfully attempts to compromise a system is likely to try again using another method. If false information is provided instead, the attacker is led to believe that the attack was successful, and thus the likelihood of attempting to break-in again is limited. That saves more time for the administrator and the IDS to counter any future attack, by patching vulnerabilities, increasing the monitoring level of the system, and developing defence mechanisms based on the security policy. Of course, false information provision could have

significant adverse effects in a false positive scenario, as legitimate users could make decisions or act upon the false data. Thus the employment of such a method could be meaningful only for attacks with significant low false rejection rate.

It is clear that the issue of selecting appropriate responses to specific types of incident demands more structured analysis. It is necessary to consider the different classes of attacks, and their distinct characteristics that will influence the appropriateness of a response. To this end, the authors have designed a response-oriented taxonomy of IT system intrusions, which can be used as the basis for such analysis. Details of the taxonomy are presented in Papadaki et al (2002).

Assessment of influencing factors in intrusion response

There are numerous individual response actions that could be pursued in order to counter an intrusion, and some decision-making ability is required when a suspected incident presents itself. As previously identified, the AAIRS project has already conducted some research in this direction, identifying a number of factors that influence the response, along with the requirement for adaptive decision making. However, the authors consider that the range of contextual factors influencing response selection can be established in more detail than the AAIRS taxonomy has currently considered, and a number of further dimensions can be identified. This is illustrated in Figure 1, with the factors split according to whether they relate to the incident or the IDS.

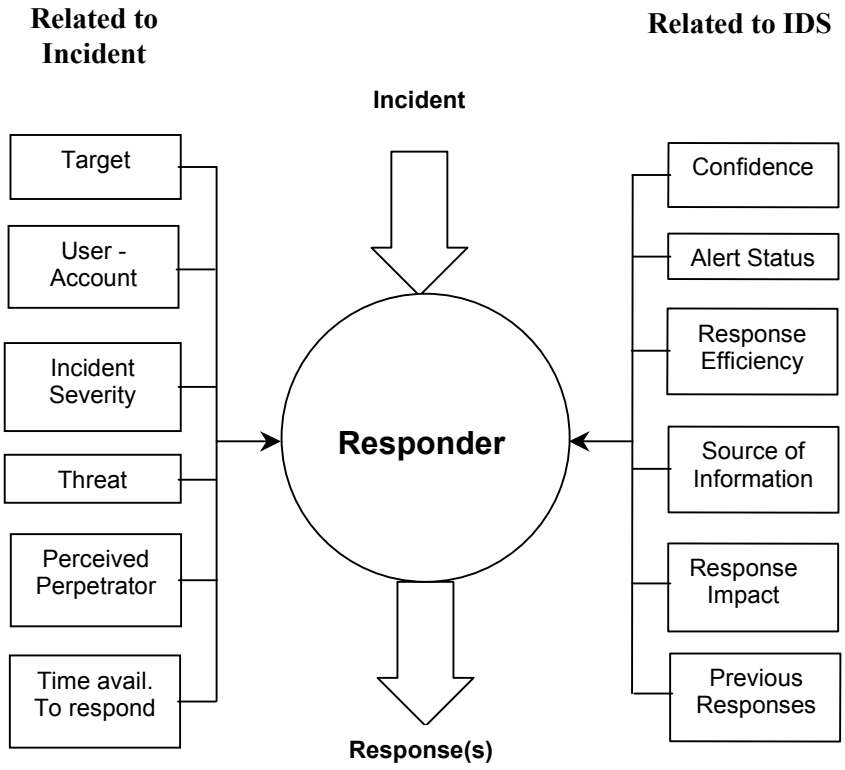

Figure 1: Contextual factors influencing intrusion response

As the figure shows, the *incident* is the trigger for the response and still represents the principal influence over what should be done. However, assessment of the other influencing factors enables the responder to establish the context in which the incident has occurred, and therefore select appropriate responses accordingly. The various factors are defined as follows:

Factors related to the incident

- **Target:** what system, resource or data appears to be the focus of the attack? What assets are at risk if the incident continues or is able to be repeated? How important is that resource for the continuation of the system operation?

- **User account:** if the attack is being conducted through the suspected compromise of a user account, what privileges are associated with that account? What risk do those privileges put the system on?

- **Incident severity:** what impact has the incident already had upon the confidentiality, integrity or availability of the system and its data? How strong a response is required at this stage? For example, the detection of a severe incident could warrant the initiation of correspondingly severe responses, in order to protect system resources.

- **Threat posed by incident:** how serious is the threat to the system, after the occurrence of the incident? Which attacks are more likely to follow, after that incident?

- **Perceived perpetrator:** does the evidence collected suggest that the perpetrator is an external party or an insider? Is there any history associated with that person/account?

- **Time available to respond:** How urgently is a response needed? This factor will be influenced by several of the other factors.

Factors related to the IDS

- **Confidence:** how many monitored characteristics within the system are suggestive of an intrusion having occurred?

- **Alert status:** what is the current status of the monitoring system, both on the suspect account / process and in the system overall?

- **Response efficiency:** what has the efficiency of a specific response proven to be under specific conditions? The IDS will gradually update the efficiency rating of a specific response, after considering its efficiency in previous incidents. For example, for some types of attacks, targets, or attackers, some responses might be more efficient than others.

- **Source of Information:** what is the detecting capability of the source of information about the incident? Some sources or IDS metrics might be more efficient in detecting attacks than others, generating less false positive alarms (e.g. anomaly detectors tend

to generate more false positive alarms than misuse detectors (Bace and Bell 2001), and some monitoring sensors produce less false alarms than others, depending on their location and configuration). The IDS should be able to determine the credibility of sources over time and adjust the confidence of the system on the probability of an intrusion.

- **Response impact:** what would be the impact of initiating a particular form of response? How would it affect a legitimate user if the suspected intrusion were, in fact a false alarm? Would there be any adverse impacts upon other system users if a particular response action were taken? Would it be possible to eliminate any adverse impacts and return the system to its initial state?

- **Previous Responses:** have any responses already been issued as a result of this incident? If one or more responses have already been issued and been unsuccessful in countering the intrusion, it would be relevant to consider this before determining the acceptable impact of the next action. The failure of previously issued responses might lead to the selection of more severe response actions (or an increase of the overall alert status of the system).

At the heart of Figure 1 was an entity referred to as the *Responder*. This is the element that will assess the various factors in order to select and invoke the required response(s). In current systems, this role is most likely to be fulfilled by a system administrator. However, in the context of an automated approach, the role would be assumed by a software-based agent, which itself would be an element of a wider intrusion monitoring system (Furnell and Dowland 2000).

Although Figure 1 highlights a number of factors, more thought should be given to the effect of different factors on the response decision mechanism. Indeed, identifying the factors that can influence response is only one part of the problem. The way in which one factor can influence others (i.e. the interrelationships between them) must also be analysed in order to determine the mechanism of the responder. For example, the type of target can influence the severity of the incident, as the more important the targeted system is (or the more vulnerable it is to specific attacks) the more severe the incident can become. In its turn, the severity of

an intrusion can also influence several factors, such as the urgency to respond, and the acceptable impact of response (the more severe an intrusion is, the less transparent the response can be). All such ways in which factors can influence one another need to be identified and analysed in order to proceed with the conceptual design of the response framework.

Other outstanding issues at this stage include the relative weightings that should be assigned to the different factors in the response decision-making process. Some factors are likely to exert more significant influence than others, and the modelling of inter-relationships will clarify this to some extent. In addition, however, it is anticipated that weightings may alter according to the type of incident involved.

A final, yet crucial, aspect that requires investigation is the extent to which the various contextual factors can actually be measured in practice. Whilst all of them make sense at a conceptual level, obtaining the necessary information to quantify them in an operational system may be non-trivial.

CONCLUSION

The paper has established the importance of intrusion response within the context of IDS systems. Although the concept is represented to some extent within current systems, the most prevalent approaches are of a passive nature, aiming only to notify other parties about the occurrence of an attack, and then relying on them to take appropriate action. Automated active responses have the potential to offer a greater level of protection, since they can include actions to actually counter attacks. However, fears are currently expressed in the security community that automated responses introduce the danger of causing negative effects on a system, in case of a false positive alarm scenario. Nonetheless, an automated capability is desirable in that it will ease the administrative workload, and can protect systems from automated attack tools around the clock.

A broadening of response methods is necessary to extend the possibilities beyond the largely passive options that exist at present. Where possible, responses must be identified that have

the potential for maximum impact upon an intruder, whilst minimising the effects upon legitimate users.

In order to increase confidence in the ability of automated response systems, the decision making process that underpins the selection of responses must be enhanced. This paper has summarised a range of factors that can influence the decision process. However, a deeper level of analysis is required in order to determine the relative importance, and consequent weightings, of factors in different scenarios, as well as potential inter-relationships between them. These aspects represent ongoing elements of research, and further findings will be documented in later publications.

REFERENCES

Allen J., Christie A., Fithen W., McHugh J., Pickel J. and Stoner E. (2000), "State of the Practice of Intrusion Detection Technologies", Carnegie Mellon University, Technical Report CMU/SEI-99-TR-028, http://www.sei.cmu.edu/publications/documents/99.reports/99tr028/99tr028abstract.html

Amoroso E. (1999), "Intrusion Detection: An Introduction to Internet Surveillance, Correlation, Traps, Trace Back, and Response", Second Printing, Intrusion.Net Books, New Jersey, June 1999.

Bace R. and Mell P. (2001), "NIST Special Publication on Intrusion Detection Systems", National Institute of Standards and Technology (NIST), http://csrc.nist.gov/publications/nistpubs/800-31/sp800-31.pdf

Carver C.A.Jr., Hill J.M.D., Surdu J.R., and Pooch U.W (2000), "A Methodology for Using Intelligent Agents to provide Automated Intrusion Response", IEEE Systems, Man, and Cybernetics Information Assurance and Security Workshop, West Point, New York, June 6-7 2000.

Carver C.A.Jr., and Pooch U.W. (2000), “An Intrusion Response Taxonomy and its Role in Automatic Intrusion Response”, IEEE Systems, Man, and Cybernetics Information Assurance and Security Workshop, West Point, New York, June 6-7 2000.

Cheung S. and Levitt K.N. (1997), “Protecting Routing Infrastructures from Denial of Service Using Cooperative Intrusion Detection”, Proceedings of the New Security Paradigms Workshop, Langdale,Cumbria UK, September 23 - 26, 1997, http://riss.keris.or.kr:8080/pubs/contents/proceedings/commsec/283699/

Cisco Systems (1998), “Cisco Secure Intrusion Detection System: Technical Overview”, Cisco Systems Inc., http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/tech/ntran_tc.htm

Cohen F.B. (1999), “Simulating Cyber Attacks, Defenses, and Consequences”, Infosec Baseline studies, http://all.net/journal/ntb/simulate/simulate.html

Computer Associates (1999), “White Paper – eTrust Intrusion Detection”, September 6, 1999, http://www.cai.com/solutions/enterprise/etrust/intrusion_detection/product_info/sw3_whitepaper.htm

Computer Associates (2001), “SessionWall3 – Intrusion Detection”, http://ca.com/solutions/enterprise/etrust/sw_intrusion_detection/product_info/sw3_intrusion.htm

Denning D.E. (1987), “An Intrusion-Detection Model”, IEEE Transactions on Software Engineering, Vol. SE-13, No. 2 pp.222-232.

Furnell S.M. and Dowland P.S. (2000), "A conceptual architecture for real-time intrusion monitoring", Information Management & Computer Security, Vol. 8, No. 2, pp65-74.

Gilliom G., Proctor P. E. (2001), “The Case for Centrac ICE Hybrid Security Solution”, NetworkICE Corporation – CyberSafe Corporation, March 2001, http://www.centraxice.com/centrax/content/CentraxICE_whitepaper.pdf/

Hurwitz Group (2000), “Hurwitz Report: AXENT Technologies’ NetProwler™ and Intruder Alert™”, Hurwitz Group Inc., September 2000, http://www.hurwitz.com/

Internet Security Systems (2001), “RealSecure 6.0: Frequently Asked Questions”, June 14, 2001, pp. 5-6, http://documents.iss.net/literature/RealSecure/rs60_faq.pdf

Lee S.Y.J. (2001), “Methods of response to IT system intrusions”, MSc thesis, University of Plymouth, Plymouth, UK, September 2001.

Mukherjee B., Heberlein L.T.; Levitt K.N. (1994), “Network Intrusion Detection”, IEEE Networks 8, No.3. pp26-41.

Papadaki, M., Furnell, S.M., Lines, B.M. and Reynolds, P.L. (2002). “A response-oriented taxonomy of IT system intrusions”, in M.Roccetti (ed.), *Proceedings of EUROMEDIA 2002*, Modena, Italy, April 2002. pp87-95.

Porras P.A. and Neumann P.G. (1997), “EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances”, 20th National Information Systems Security Conference, October 9, 1997, http://www.sdl.sri.com/projects/emerald/emerald-niss97.html

Power R. (2001), “2001 CSI/FBI Computer Crime and Security Survey”, Computer Security Issues and Trends, Vol. VII, No. 1.

Ragsdale J.D., Carver C.A. Jr., Humphries J.W., and Pooch U.W. (2001), “Adaptation Techniques for Intrusion Detection and Intrusion Response Systems”, 2nd Annual IEEE Systems, Man, and Cybernetics Information Assurance and Security Workshop, West Point New York, June 5-6, 2001.

Schneier B. (2000), “Secrets and Lies: Digital Security in a Networked World”, John Wiley & Sons, ISBN 0-471-25311-1.

Shipley G. (1999), “Intrusion Detection, Take Two”, Network Computing online journal, November 15, 1999, http://www.nwc.com/

SRI (2000), “EMERALD expert-BSM: Capabilities”, SRI International. http://www.sdl.sri.com/projects/emerald/releases/eXpert-BSM/cap.html

UC Davis (2000), "UC Davis Response and Detection Project Overview", December 2000, http://seclab.cs.ucdavis.edu/response/