THE NEED FOR CENTRALISED, CROSS PLATFORM INFORMATION AGGREGATION

Fred Otten, Barry Irwin, Hannah Slay

SNRG (Security and Networks Research Group)
Computer Science Department
Rhodes University
Grahamstown
South Africa

g05o5894@campus.ru.ac.za, b.irwin@ru.ac.za, h.slay@ru.ac.za

ABSTRACT
With the move towards global and multi-national companies, information technology infrastructure requirements are increasing. As the size of these computer networks increases, it becomes more and more difficult to monitor, control, and secure them. Network security involves the creation of large amounts of information in the form of logs and messages from a number of diverse devices, sensors, and gateways which are often spread over large geographical areas. This makes the monitoring and control difficult, and hence poses security problems. The aggregation of information is necessary in information audits, intrusion detection, network monitoring and management. The use of different platforms and devices complicates the problem, and makes aggregation more difficult. Network security administrators and security researchers require aggregation to simplify the analysis and comprehension of activity across the entire network. Centralised information aggregation will help deal with redundancy, analysis, monitoring and control. This aids the detection of wide spread attacks on global organisational networks, improving intrusion detection and mitigation. This paper discusses and motivates the need for centralised, cross platform information aggregation in greater detail. It also suggests methods which may be used, discusses the security issues, and gives the advantages and disadvantages of aggregation.

KEY WORDS
Data Fusion, Denial of Service, Intrusion Detection, Information Aggregation

THE NEED FOR CENTRALISED, CROSS PLATFORM INFORMATION AGGREGATION

1 INTRODUCTION

The scope, size and spread of network infrastructures have been rapidly increasing over the recent years. This has resulted in many changes within Information Technology, particularly within information security. Corporations now acknowledge the need for Information Security because they realise the vulnerability and economic value of their data. This change in attitude, coupled with regulations which require accountability and storage of records, and the potential to enhance business decisions has meant that the role and management of data and logs has changed. The rise in Internet business activity such as online shopping and online exchanges also necessitates records of customer activity to keep a competitive edge, thus raising a new use for logs [11].

Besides security logs and traffic captures, there is also a lot of other related data and information spread throughout the network. These include employment, financial and telephone records. All of this data and information needs to be brought together at some stage for auditing and management purposes. The confidentiality, integrity and availability (CIA) of data, information and knowledge within the network are the key concerns in providing security [1]. A good understanding and overview of the network, its users, and its operation is key to being able to keep the network secure. The distributed nature and the geographical spread of networks pose the first problem and the use of a myriad of heterogeneous devices and platforms poses the second problem.

Network backbones involve the use of many different types of devices, sensors and gateways made by different manufacturers, each with their own logging structure and storage formats. Different platforms such as smartphones, Voice over IP, wireless networks, and a variety of applications such as web servers, databases, mail servers and proxies are involved in the network infrastructure. Often, different applications are used within different sections of the business for the same purpose. To maintain good security, all this data needs to be kept secure, and there need to be facilities to bring related data together from their different platforms for management, decision making and audits.

The problem is compounded with the explosive growth of the data and information involved. The quantity, value, amount of detail and the amount of systems in place are all increasing with business demands. There is a constant demand for new technologies, while maintaining backwards compatibility and security. “For most networks and businesses, the most important requirement is to keep the network running at an acceptable risk level without downtime” [9]. This means that Denial of Service (DoS) and other attacks that lead to downtime needs to be avoided. “Most organisations deal with literally millions of messages daily from these incompatible security technologies, resulting in security information overload, in turn contributes to high overhead, duplication of effort, weak security models and failed audits.” [9]. And this is not getting any better.

In order to maintain CIA, situational awareness is essential. “Situational awareness can be described broadly as a person’s state of knowledge or mental model of the situation around him or her” [2].

Cognitive Psychology defines three levels of situational awareness [2]:

1. Perceiving critical factors in the environment
2. Understanding what those factors mean, particularly when integrated together in relation to the decision maker's goals
3. Understanding what will happen with the system in the near future.

Essentially, security administrators require at least the first level of situational awareness in order to secure the network. If critical devices are not able to work with the security information management products they can lead to dangerous blind spots in the network [9]. This is the primary motivation for situational awareness. With the distributed approach, security administrators are often isolated, and don't have access to the whole picture of the network. "Data is the measurements and observations. Information is the data placed in context, indexed and organised. Knowledge or intelligence is information explained and understood" [1]. Security Administrators desire the best information possible so that they can make good diagnoses, avoid attacks and maintain CIA within the network. The use of centralised, cross platform information aggregation can provide administrators with a real-time view into a network's security status, making a proactive approach to security a reality via automated alerts, detailed reports, and remediation [9]. This makes monitoring, control and management easier, reducing redundancy, and making it possible to detect complex attacks. The more data that can be gathered and correlated, the more accurate intelligence you have to mitigate and resolve the event [9].

This paper investigates the problem presented above. It starts by looking at monitoring and control, discussing the problems facing security administrators and the requirements. It then contrasts the difference between distributed and centralised monitoring and control. Next it discusses the need for information aggregation, and what it adds to the approach, giving useful applications. It then details intrusion detection, and how information aggregation and multisensor data fusion add to the existing approaches. Finally, it then discusses our suggested architecture for information aggregation, which uses XML, and concludes with a summary and a few comments.

2 MONITORING AND CONTROL

In order to achieve CIA in a network, it is essential to be able to monitor and control the network and its perimeter. There are a myriad of heterogeneous devices, sensors and gateways involved in the perimeter of the network. Regulations in many countries require that logs are kept of network activity, and may be called up during legal investigations. Security administrators are required to keep track of the logs and activity on the network, investigating previous incidents and attempting to avoid new incidents from occurring. In a large, distributed network, this is a mammoth task. For most networks and businesses, the most important requirement is to keep the network running without downtime, with little regard to the risk level [9]. Most of the time, security administrators spend their time deploying patches, performing investigations and researching exploits. This means that the monitoring of logs in their area of the network is not a high priority, and often left to Perl scripts using regular expressions which are written hastily after CIA has already been violated.

The aim of monitoring and control is essentially situational awareness."The phrase 'forewarned is forearmed' sums up the value situational awareness. Simply put, being aware is about being prepared to act and respond" [9]. Situational awareness is a constant health check of the network. This is essential to maintaining CIA. Early diagnosis and response to a situation potentially saves the life of the system, and protects the data and information involved. The more data and information that can be gathered and correlated in the process of monitoring and control, the more data and information you have to mitigate and resolve problematic events. Security information management essentially involves information and knowledge found during monitoring and control. The current trends show that security information management is converging with network and systems management [9].

"In the typical security environment, businesses rely on a multitude of disparate point solutions to prevent viruses, worms, spam and malicious content from infiltrating their networks, as well as to ensure that business data and private information are not compromised" [10]. This distributed approach has its limitations when it comes to situational awareness in large networks. With a distributed approach there are a number of system administrators each manning their own section of the network, while the centralised approach demands resources and opens the door to loss of information should the location be compromised and suitable redundancy not available. Both centralised and distributed approaches have their advantages and disadvantages, which lead to debate as to which is more applicable.

3 DISTRIBUTED VS CENTRALISED MONITORING AND CONTROL

The question remains whether a distributed approach or a centralised approach is a better methodology for monitoring and control. In this section we take a look at the issues involved. It must be emphasised that the solution largely depends on the particular organisation at hand. Our major focus is on large geographically disparate organisations.

We have established that there is plenty of data and information distributed throughout the network that is useful for monitoring and control. Our goal is to be able to use as much of this data and information to get better information and knowledge to provide better security. The question is whether the results will be better obtained from a distributed or a centralised approach, and whether the extent is worth the cost.

The vast quantity of data and information means that it is a battle to compile the resources needed to review the data coming from all these systems . There are millions of alerts and messages generated by each individual system, such as the intrusion detection systems, anti-virus systems, firewalls, operating system logs and access control systems [9]. This is overwhelming. To employ a distributed approach, the workload is great for few administrators, and this could lead to sections being overlooked, which could in turn lead to large holes in the perimeter for intruders to exploit. Most of the devices display related activity, and complement each other in analysis. A distributed approach could mean that administrators work in isolation and do not see the bigger picture of the network. Collaboration is possible with a distributed approach, however in large geographically spread networks, it is often not practical because administrators are quite busy keeping their segment of the network secure, as discussed previously. "Consolidating all of the reports from all of these devices and tying

the information together into a coherent visual artifact closes the window of risk" [9]. This motivates for the use of the centralised approach, however sometimes it is not practical, and collaboration using email and other means suffices. This largely depends on the size and spread of the organisation and the number of administrators employed.

When defining his new architecture for managing enterprise log data, Adam Sah [11] gives examples of the amount of log traffic produced. The amounts range from 7 GB per day for security related syslogs to 200 GB per day of records for people viewing online ads. Sending these logs to a central location may seem impractical. Let us say we are monitoring and controlling 10 servers, each producing 7 GB, this would mean that we would be looking at 70 GB of logs rather than ten people each looking at 7 GB of logs. 7 GB of data is also a lot of data to be sending to a central location, and would require a lot of processing power and storage space. Two solutions may be proposed: The clustering of the central location, to allow for greater incoming bandwidth potential and greater processing power; and pre-processing by each server before sending to the central location. The problem with pre-processing is that samples and summaries are not useful for security applications or satisfying regulatory requirements [11]. Centralised logs are necessary for data mining to find new attacks and to detect anomalies. By accumulating statistics about what constitutes "normal" activity, log monitors may be able to recognize anomalous behaviors that a human system administrator might at first overlook, such as the cessation of events which normally occur with a given frequency [4]. The information retrieved from logs is also useful for security researchers in their search to define "normal activity", which is still an open and difficult research problem [12]. The more data that can be gathered and correlated, the more accurate intelligence you have to mitigate and resolve the event. [9]. Forensic and historical data provide maps of past activity. Through analysis of these maps, we can get a better picture of what happened during the attacks and can gain a better understanding of the attack's operation and the path along which it travels. This may lead to the discovery of defense strategies [9].

A good question to ask, is whether it is worth the cost to install a centralised monitoring and control system. New infrastructures will need to be put into place and off-site redundancy needs to be provided to avoid the fatal results should the central facility go down, such as the situation which occurred on September 11 [5]. The value of the information and the cost of the loss of information need to be considered. Although the new infrastructure would cost the organisation millions of dollars, the long term benefits and the value of the information being protected must be considered. As noted earlier, the current security methodologies place a large burden on security administrators. They rarely have time to monitor logs, and spend most of the time trying to repair and defend against attacks which have already occurred, rather than monitoring and researching to prevent attacks. The attack implications, particularly to large organisations, are drastic. DoS leads to a large loss of revenue, and organisations could be sued, should customer information be obtained. Centralisation can result in fewer administrators who achieve more, providing better utilisation of resources. It also makes automation easier, as management and control can be deployed from the central location [10]. Essentially, the quality of administrator's work is improved. This improvement is necessary since the frequency and complexity of attacks are increasing, as new technologies are developed and there is an expansion of networks and services. With the added pressure of regulatory compliance, and the

pressures of security audits, security administrators are too busy to use bad systems and spending time doing tasks which, ideally, should be automated and centralised. It helps to have a central point where a few administrators may view the whole network and concentrate on maintaining the CIA within the network. Security administrators can then be assigned jobs that are more worthy of their talents, such as diagnosing anomalies and creating prevention schemes [4].

A disadvantage of centralisation is that it opens an exploitation target in a single location. In terms of all the logs going to a central location, it opens the possibility of attackers sending false logs, to distract the administrators or to make it seem like nothing is going on, if not correctly implemented. This can however be easily solved by signing and / or encrypting the traffic sent to the central location in the same mold as a service such as syslog-ng. Other issues that need to be addressed are: how to structure the monitoring and control network; whether to use existing networks, or install new private networks; and then making sure these new networks are secure to avoid exploits which take advantage of the centralised structure.

Essentially, a centralised architecture makes situational awareness clearer and has many advantages over a distributed approach. The extent of the separation of systems, however, still remains an issue.

4 THE NEED FOR INFORMATION AGGREGATION

Mere centralisation of monitoring and control is not enough to achieve a higher level of situational awareness. This requires the information from related disparate point solutions to be aggregated and available for further analysis. This will also decrease the burden on system administrators, by dealing with redundancy and providing facilities to focus on understanding their systems rather than all dealing with the same problem, and only relying on their own observations of their network sphere.

The typical security environment consists of a multitude of disparate point solutions to prevent viruses, worms, spam and malicious content from infiltrating their networks. Their role also entails ensuring that business data and private information are not compromised [10]. As mentioned earlier, situational awareness is required to effectively achieve CIA within the network. The incessant flood of data, literally millions of messages daily, from incompatible security technologies results in security information overload and redundancy in separate analysis. With no way to manage and integrate information, this fragmented approach often leads to duplication of effort, high overhead for the system administrators, weak security models and failed audits [10]. If an attack or vulnerability occurs on the network and action is taken and the problem eliminated, it is essential that the process doesn't stop there. It is necessary to learn how to prevent this kind of attack from occurring again in the future and stop it from spreading further. Forensic and historical data are maps of what happened. They can divulge the working and the path of an attack [9]. Without information aggregation, we have fragmented data and we are not able to easily assess the working path of the attack. Determining how the attack occurred is essential to developing a response or countermeasure [8].

Tim Bass [1] takes a look at using multisensor data fusion in Intrusion Detection Systems (IDS) to achieve situational awareness in cyberspace. In his paper he highlights that in order to take our level

of inference from low (existence of intrusion) to high (threat analysis), we require the the use of aggregated information. He also describes the classical decision system used in the military, observe - orient - decide - act (OODA). In order to achieve our goal of being able to act from a place of knowledge, a standard metalanguage is required for object refinement, data storage, cleansing and primitive correlation. “Data is the measurements and observations. Information is the data placed in context, indexed and organised. Knowledge or intelligence is information explained and understood” [1]. It is challenging to turn data into information & knowledge. Situational awareness makes this task much easier. We can bring a lot of information to a central point, but without bringing it together for analysis, we lose the essential advantage.

Figure 1: Typical data fusion system architecture [13]

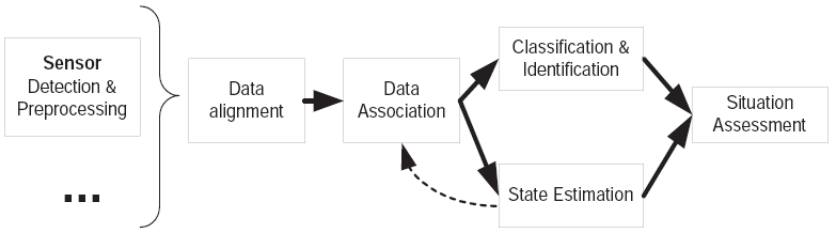

Siateris, *et al*. [13] defines a typical data fusion system architecture. This is illustrated in Figure 1. As this picture shows, the sensors detect the information and perform simple processing before passing the data on. This preprocessing could include data reduction and/or data transformation to a common format. The data is then aligned in terms of time, space and measurement units with other data, and associated data pooled to reduce redundancy. From this, data analysis is performed to estimate the system state. This process leads to information aggregation and thus situational awareness through assessment of aggregated data and information.

Information aggregation provides the framework necessary to deploy central management and achieve situational awareness and better diagnosis. There are a number of useful applications for aggregated information. These include: load balancing; easing and improving daily monitoring and control; easing the task of day-to-day debugging and the generation of reports; improved forensic capabilities, intrusion detection and situational awareness. More information is processed and redundancy is reduced with the use of information aggregation. Aggregated information provides a tools for researchers to get a better idea of what “normal” traffic constitutes. The legal requirements for storing data are also satisfied in the process. Audits and management are all improved and made easier. DoS poses a huge threat to large organisations, and detection and mitigation is aided by intrusion detection systems using aggregated data and information [1].

5 INTRUSION DETECTION AND MITIGATION

Traditional IDSs implement approaches such as known pattern templates, threatening behavior templates, traffic analysis, statistical-anomaly detection, and state-based detection for intrusion detection and mitigation [1]. These techniques are limited, and have difficulty detecting complex attacks such as DoS reliably. DoS poses a pressing issue to organisations and their customers.

There are an increasing number of DoS attacks on organisations. Between 2000 and 2003 there were several examples of Distributed Denial of Service (DDoS) attacks [13]. These included attacks against root name servers, attacks against a spam black-list company and attacks against Yahoo!, Amazon, eBay, CNN, ZDNet, E*Trade and Excite [3]. "Judging from the latest trend to use worms as DDoS attack agents, the future looks bleak" [13]. Intruders are attacking from geographically dispersed networks [1], spoofing IP addresses, changing IPs and using complex attacks such as TFN2K and Stacheldraht [6]. They are targeting organisations with the aim of taking down their networks. These challenges make it difficult to defend. "Defensive information operations and computer ID systems are primarily designed to protect the CIA of critical information infrastructures. These operations protect information infrastructures against DoS attacks, unauthorised disclosure of information, and the modification or destruction of data. The automated detection and immediate reporting of these events is required to respond to information attacks against networks and computers" [1]. False positives (or false alarms) and false negatives are a problem for intrusion detection systems. False alarms lead to undermined confidence in the systems which leads to poor maintenance and under utilisation. False alarms also lead to financial loss and DoS when there is no attack because of investigation [1]. "Today high false alarm rates and successful detection only when damage is already done (near the vicinity of the victim where the available bandwidth has already been consumed in the upstream path) are the main problems that hinder the automatic deployment and the effectiveness of countermeasures like firewall filtering, rate limiting or route blackholes" [13] False negatives mean that intruders are missed, and no defense is thus offered to their attacks. This is a large problem.

Using information aggregation we will be able to detect and deploy patches for viruses and worms, and improve the detection and mitigation of DoS. Information aggregation can also reduce the number of false positives and false negatives through combining information from multiple intrusion detection systems. There are many point solutions which provide useful information for intrusion detection and mitigation and may be used as sensors. These include: various IDSs; custom DDoS detection programs; SNMP-based network monitoring systems; active measurements or accounting systems like Cisco's Netflow; commands and *a priori* data from established databases; distributed packet sniffers; system log files; user profile databases; system messages; and operator commands. Other sensor types might make active measurements like round trip time or packet loss estimation or provide flow level information about network traffic [13] "Next generation cyberspace intrusion detection (ID) systems will require the fusion of data from myriad heterogeneous distributed network sensors to effectively create cyberspace situational awareness. The vast majority of security professionals would agree that real-time ID systems are not technically advanced enough to detect sophisticated cyber attacks by trained professionals." [1]

Multisensor data fusion methodologies include many established mathematical models such as: Clas-

sical Inference; Bayesian Inference; Dempster-Shafer Method; Generalized EPT; and Heuristic Methods [1]. The discussion of these methods is beyond the scope of this paper. Their purpose is to take data and information from the many point solutions and combine them into a single metric which may be used to make decisions.

Siaterlis, *et al*. [12, 13] provides a multisensor data fusion proof-of-concept. It is built using the Dempster-Shafer Method, based on the Dempster-Shafer theory of evidence. "Network engineers know empirically, that there are often signs of flooding attacks but these are not always accurate or definite indication." [12]. Their research shows the value of multisensor data fusion. Even though one of the sensors failed to detect an attack, the combined knowledge resulted in detection of the UDP flood. In their paper, they give details of the DDoS detection tool they built. They used the number of active flows and the ratio of incoming and outgoing UDP traffic as metrics on their two sensors. The sensor using the number of active flows failed to detect the UDP flood, while the sensor using the ratio detected the attack. The combined knowledge indicated correctly that there was a UDP flood. One might argue that in a research environment it is easier to determine as there is not real traffic present. Their prototype evaluation, however, was performed in a real world environment, with the attack coming along a busy link from the ISP. This illustrates the value of information aggregation in intrusion detection.

Information aggregation also offers great potential for mitigation after attacks. Performing data mining on the aggregated data results in an improved picture of the situation. It makes learning new trends possible, and aids the understanding of "normal" activity within a network. In order to cope with advances in attack technologies and distributed attacks, it is necessary to involve information aggregation in intrusion detection and mitigation.

6 AN ARCHITECTURE FOR CENTRALISED INFORMATION AGGREGATION

We have established the need for information aggregation and the need for a standard meta-language. XML provides users with a great option for transporting data between our sensors and our central point. Our architecture involves the use of collectors which are scripts which transport relevant data and information to the central point. Figure 2 shows the architecture using a router, a firewall, an access point and a server as sensors. The collectors retrieve the log information from these sources, and send them to the central location where system administrators are monitoring and controlling these devices.

Since the problem is system specific, a generic model needs to be able to cater for as many cases as possible. The collectors determine what is sent to the central location, which provides flexibility. The collector scripts may perform pre-processing and only send important information to the central location, or it may simply send all data and information to the central location. This also provides flexibility to the point where we perform data fusion. The scripts may take two or more pieces of data from the sensor and perform data fusion, sending the resulting information to the central location. As noted earlier, the problem with pre-processing is that samples and summaries are not useful for

Figure 2: Our architecture for central aggregation

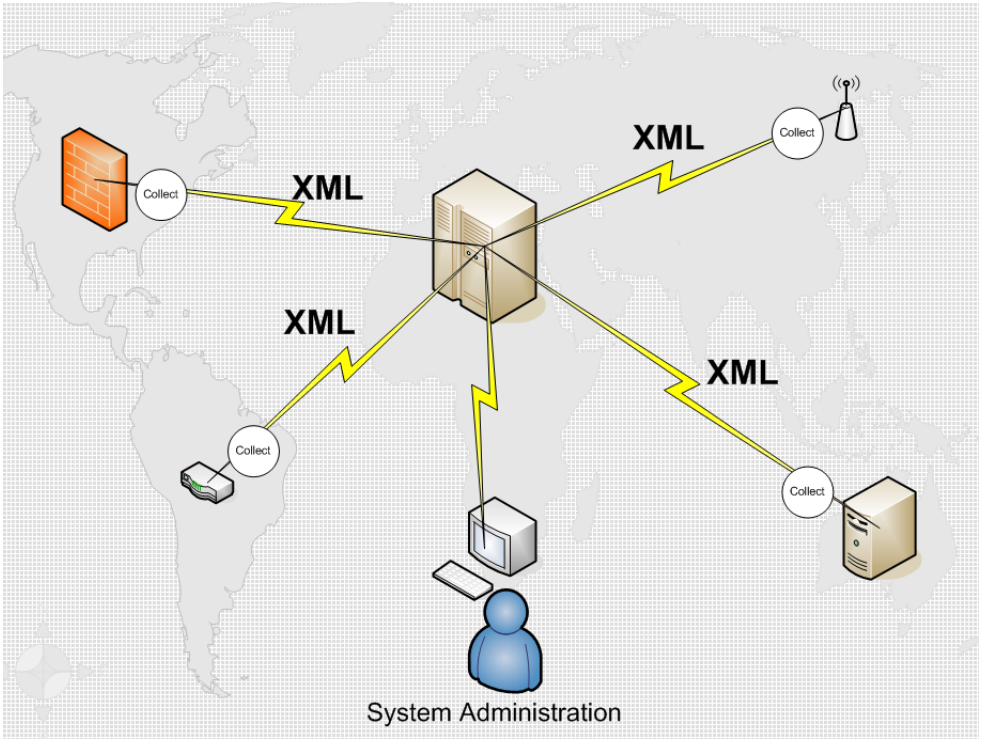

security applications or satisfying regulatory reasons [11]. Hence, pre-processing is not advisable, but the facility is there should the particular system require.

To deal with the quantity and reliability of data, we will need to introduce a further step. The collectors require acknowledgment of receipt and may also keep a buffer of messages to be passed to the central location, preventing loss of data and maintaining near real-time monitoring and control. Time stamps are also necessary to allow for accurate correlation in information aggregation. The various collectors should all be time-synchronised. This means that within our network, we should deploy a centralised time server which may be used as a point of reference, and to which the sensors regularly synchronise themselves.

The role of the collector is largely determined by the system, but is limited by the XML. In order to maintain a generalised model, it might be necessary to alter the XML schema according to the system. For instance, in a telecommunication monitoring system, the logs will be dealing with different metrics than in intrusion detection. They do, however, both involve the establishment and teardown or sessions (connections or calls). "Correlation in cyberspace requires the comparison of observations based on a different set of parameters such as source (IP address), network path, session flow, or behavior." [1]. We need an XML schema which can deal with this. Below is a list of generalised aspects which need to be sent to the central location via XML:

- The collector identity

- Connection information involved
- Observation details
- Threat level
- Specific detail (used later in data mining, and to satisfy regulations)

As mentioned earlier, the collectors are scripts. For these scripts, we need a fast, simple language with good regular expression compatibility and the ability to easily establish connections and send XML. Python is our chosen language as it satisfies all of these requirements, and coupled with the Twisted framework is well adept at establishing connections and sending XML.

Read-Miller defines three functions for a centralised monitoring and control center employing information aggregation: It must operate in real-time; have facilities to provide further insight (through forensics); and provide good communication of information and knowledge generated [10]. Our architecture satisfies these criteria. Our architecture allows for a "central console for network and security situational awareness allowing organisations to quickly identify, respond and mitigate security events across the organisation" [9].

There are a few security concerns that may be raised with this architecture. How do we prevent an intruder injecting logs? How do we know that the information is coming from the collector? What about man in the middle attacks? By employing standard security practices, we can resolve these issues. Messages may be signed in the mold of syslog-ng, which is a centralised log monitor. We can also encrypt the communication between the endpoints and the central location, and use certificates to prevent man in the middle attacks and accurately establish identity.

7 CONCLUSION

This paper has highlighted the benefits of information aggregation and presented a general purpose architecture which may be used for information aggregation. The use of centralised, cross platform information aggregation can provide administrators with a real-time view into a network's security status, making a proactive approach to security a reality via automated alerts, detailed reports, and remediation [9]. Information aggregation leads to the establishment of knowledge and the removal of redundancy. It increases the possibility of detecting and mitigating more complex attacks, and provides better intelligence for mitigation and prevention of attacks.

Information aggregation has a number of security related applications. These include: load balancing; improving daily monitoring and control; easing the task of day-to-day debugging and the generation of reports; improved forensic capabilities and intrusion detection. Information aggregation essentially helps to achieve situational awareness, leading to a better understanding of what is going on within the network. Aggregated information provides a tools for researchers to get a better idea of what "normal" traffic constitutes. Legal requirements for storing data are satisfied, audits and management improved and the threat of DoS and DDoS attacks is minimised through the use of information aggregation.

In summary, information aggregation provides the basis for the framework of taking distributed data and information and turning it into knowledge which may be used to improve and maintain the CIA within the network and security infrastructures.

References

[1] T. Bass. *Intrusion detection systems and multisensor data fusion*. Communications of the ACM, Vol 43, No. 4, pg 99-105. April 2000.

[2] M. Endsley. *Theoretical underpinnings of situational awareness: A critical review.* SA Technologies Publications. 2000.

[3] L. Garber. *Denial-of-Service Attacks Rip the Internet.* Computer pg 12-17. April 2000.

[4] B. Glass. *Log Monitors in BSD UNIX.* Proceedings of the BSDCon 2002 Conference, San Francisco, California, USA. USENIX Association. February 2002.

[5] S. Hanning. *Recovering From Disaster: Implementing Disaster Recovery Plans Following Terrorism.* SANS Security Essentials. September 2001.

[6] F. Lau. S. Rubin, M. Smith and L. Trajkovic. *Distributed Denial of Service Attacks*. Proceedings of 2000 IEEE International Conference on Systems, Management, and Cybernetics, pg 2275-2280. October 2000.

[7] J. McCray. *A Roadmap to Becoming Security Conscious*. Proceedings of the 2003 IEEE Workshop on Information Assurance. United States Military Academy, West Point, New York. June 2003.

[8] M. Ranum, K. Landfield, M. Stolarchuk, M. Sienkiewicz, A. Lambeth and E. Wall. *Implementing a Generalized Tool for Network Monitoring.* Proceedings of LISA '97, San Diego, California, USA. USENIX Association. October 1997

[9] S. Read-Miller and R. Rosenthal. *Best Practices for Building a Security Operations Center*. Computer Associates White Paper. April 2005.

[10] S. Read-Miller. *Security Management: A new model to align security with business needs.* Computer Associates White Paper. April 2005.

[11] A. Sah. *A New Architecture for Managing Enterprise Log Data.* Proceedings of LISA '02, Berkeley, California, USA. USENIX Association. November 2002.

[12] C. Siaterlis, B. Maglaris and P. Roris. *A novel approach for a Distributed Denial of Service Detection Engine.* 2003.

[13] C. Siaterlis and B. Maglaris. *Towards Multisensor Data Fusion for DoS detection.* SAC '04, Nicosia, Cyrus. March 2004.

This work was undertaken in the Distributed Multimedia Centre of Excellence at Rhodes University, with financial support from Telkom SA, Business Connexion, Comverse, Verso Technologies, Tellabs and StorTech THRIP, and the National Research Foundation.