Analysis of GMU Firewall Rule Set Configuration Management and Auditing Practices

SYST 699 - Final Report

Chris Grubb, Ray Kepler

Contents

Table of Figures

Acknowledgements

Ray and I would like to thank Ben Allen for the opportunity to work on this project, and his open-mindedness and flexibility as a project sponsor and client; Dr. Karla Hoffman for her guidance through a few significant project challenges; Larry Song in GMU NET for his time and willingness to candidly discuss GMU NET firewall operations; Jon Goldman in VSE for his time explaining firewall operations from a user's perspective; and Ankit Shah for some early, sage advice on preparing a successful project. We would also like to thank our families for their support throughout our time in school.

1 Project Background

1.1 Client

George Mason University (GMU) provides internet and network-based services for the students, faculty and staff, and other personnel within the university. These services include hardware and software to run email servers, education services such as Blackboard, administrative functions including registration and parking pass purchases, and a myriad of other functions that enable GMU to operate efficiently.

GMU is also responsible for providing these services securely. One tool that GMU Network Engineering uses to provide secure networked services is a firewall. A firewall is "…a security structure that creates a barrier – or firewall – between a secure network and another network that is not known to be secure."[1] Firewalls use multiple techniques to filter traffic by ultimately determining whether internet traffic is dropped or allowed.

Within GMU, the Information Technology Services (ITS) provides IT services within the university. Within ITS, Network Engineering & Technology (NET) operates and manages the network firewall. Our sponsor for this project is the Director of GMU NET, who has overall responsibility of the management and operation of the firewall, its rule set, change management of the rule set, and auditing of the firewall rule set (among many other responsibilities).

GMU NET works closely with other organizations within ITS, including the IT Security Office (ITSO), who are responsible for setting security policy, reviewing some firewall rules or proposed firewall rule changes, and generally providing as-needed support to ITS.

1.2 Problem Statements

Currently, GMU NET has no documented processes or procedures to manage their network firewall rules' configuration. Additionally, contrary to industry standards for configuration management, GMU NET has no documented firewall rule set auditing processes or procedures.

1.3 Approach

Our team developed a primary and secondary approach. The primary approach is intended to address the problems of a lack of documented processes and procedures related to firewall rule configuration management and auditing. The secondary approach is related to assessing

[1] Computing Technology Industry Association (CompTIA), A+ Certification Glossary of Terms

firewall rule sets and the impacts of introducing new firewall rules, changing existing rules, or deleting rules from the rule set.

The primary approach is composed of interviews with our client, GMU NET engineers who are responsible for managing firewall rules, and users of GMU NET support services related to firewall rule management (e.g. proposing a new computer resource to be added to the GMU network).

We also review literature relevant to firewall rule set configuration management and auditing, with an emphasis on identifying industry standards. From these standards we develop assessment criteria that we use to evaluate existing configuration management and auditing procedures. Then, on the basis of our findings, we propose recommendations that address our findings.

Our secondary approach is centered on impact analysis and was contingent on receiving data from GMU NET related to the operation of the firewall. Obtaining this data was challenging, but was obtained in November. Due the point in the semester at which we received data, this approach is intended to demonstrate methods that can be used to assess the impact of firewall rule changes. In particular, we use basic descriptive statistics, basic network/graph analysis, and a graphical modeling language to demonstrate methods for impact analysis of firewall rule additions, deletions or changes.

2 Background

2.1 Network Firewall Basics

As described above, firewalls create barriers between computer networks by filtering communications based on the features of that communication.[2] Internet communication flows in the form of packets that can be specified by values for the following kinds of communication attributes:

- a source and destination internet protocol (IP) address;
- the ports used in the communication (i.e. the source and destination ports);
- the communication protocol being used such as hypertext transfer protocol (HTTP), transmission control protocol (TCP)/IP, file transfer protocol (FTP), post office protocol (POP), and internet message access protocol (IMAP);
- and the traffic content associated with the applications being used.

[2] Since the scope of this paper is focused on the management of firewall rules, the description of internet communications will only be as detailed as is necessary to communicate firewall rule management.

IP addresses can be thought of in much the same way as physical addresses: they specify a location from which a communication (such as a letter or package) originates and to which it is intended to be sent. In the case of internet communications, IP addresses can be associated with equipment that generate and receive internet communications. Examples of this hardware include personal computers and servers, printers, and firewalls among many others.

IP addresses in the IPv4 format are four 8-bit values separated by a '.'. So, for example, an IPv4 address such as 199.27.76.73 (which resolves to cnn.com) includes the four 8-bit numbers described above. Recently, internet communications have also needed to be compliant with six 8-bit numbers, due to the scarcity of unique IPv4 addresses.

Ports are the specific hardware endpoints[3] at an IP address that are listening for communications to that address. To extend the physical metaphor of addresses above, an address describes a location and a mailbox describes the hardware endpoint of a letter (the analog to internet communication). Ports are like channels on a computer (e.g. a network server) that are expecting internet communications. For example, if one were to navigate in their browser to 'cnn.com' the computer sends a request to that address which is listening for specific kinds of requests on specific ports. When the 'cnn.com' server receives that request on the port that is listening, then it will respond with the content. Ports can be configured to listen for only certain kinds of communication protocols.

We provided several common internet communication protocols above (e.g. HTTP or TCP/IP). Communication protocols are standards that specify the syntax and attributes of communication between the sender and receiver in a number of communications scenarios. For the purposes of this paper, it is sufficient to know that internet communications follow these standards and that firewalls are capable of filtering communication traffic based on the format of the communication, i.e. the source IP, destination IP, source port, destination port, or communication protocol.

Finally, modern firewalls are also capable of filtering traffic on the basis of their content. For example, while a firewall may allow general HTTP traffic to pass, it may not allow Facebook traffic to pass. We avoid discussing this kind of capability in many of our examples because it adds a distracting level of complexity. There is no reason that these attributes couldn't in principle also be considered.

[3] Again, it is important to recognize that this definition is not strictly speaking accurate because of arrangements such as virtualized computers, which are able to emulate hardware. In such a case, the port would be virtualized hardware but not hardware *per se*. For the purposes of this project, identification and discussion of source and destination ports are sufficient.

2.2 Firewall Policies/Rules Basics

Firewall rules can be thought of as filters. These filters will either allow or deny internet communications based on properties of the traffic. We focus exclusively on the following properties: source and destination IP addresses, source and destination ports, and the communication protocol. A rule describes whether an allow or deny action is taken on the basis of the value of any subset of these properties.

Consider a simplified representation of internet traffic and firewall action in the following figure:

Source Address	Destination Address	Action
192.168.1.1, 192.168.1.2	*	Allow
192.168.1.2, 192.168.1.3	*	Block

Figure 1: Simplified Firewall Rule

Using the first row's information for an example, traffic that is coming from either of the two source addresses heading to any destination address (represented by a '*') will be allowed. It is important to note that 'allow' means that the traffic will not be filtered by additional rules. If the traffic's state remains undetermined – that is, if it hasn't been either dropped or allowed – it will continue to pass through each rule, until it's state is determined. Therefore, one can see that rules are applied sequentially. Also, at the end of these sequences of rules, there is typically a default behavior to block or allow the traffic.

It is also important to note that the sequence in which the rules are applied may affect the outcome. Using the table above, if the first row's rule is applied first, then traffic coming from '192.168.1.2' will be allowed. However, if the second row's rule is applied first, then it will be blocked.

Finally, firewall rules can have 'anomalies' that arise out of the definitions of the rules and the sequence in which they are applied. "Firewall policy anomaly occurs when rule list contains two or more rules which match the same packet. Also, it can occur when a rule never matches any packet that goes through the firewall."[4] Following Katić and Pale, we group firewall rule anomalies into the following five groups of anomalies and use definitions from Katić and Pale:

1. Shadowing: "rule is shadowed by a preceding rule which matches all the packets that would match this shadowed rule. Thus the shadowed rule will never be activated."
2. Correlation: "two rules are correlated if they have different filtering actions, and the first rule matches some packets that match the second rule and the second rule matches some packets that match the first rule."

[4] Katić, P. and Pale, P. "Optimization of Firewall Rules." Proceedings of the ITI 2007 29th Int. Conf. on Information Technology Interfaces, June 25-28, 2007

3. Generalization: “rule is a generalization of a preceding rule if they have different actions, and if the first rule can match all the packets that match the second rule.”
4. Redundancy: “rule is redundant if there is another rule that produces the same matching and action such that if the redundant rule is removed, the security policy will not be affected.”
5. Irrelevance: “filtering rule in a firewall is irrelevant if this rule does not match any traffic that may flow through this firewall. This exists when both the source address and the destination address fields of the rule do not match any domain reachable through this firewall.”

Firewall rule anomalies can also be shown using Venn diagrams, where one rule stands in relation to a preceding rule in the rule set, as can be seen in the following figure:[5,6]

[5] This figure does not show Irrelevance anomalies because these anomalies do not arise due to one rule’s relationship to another. Irrelevance rules describe cases where a rule filters traffic based on a value that is not present in the traffic. For example, if a rule were to drop traffic with a source IP address as ‘1.1.1.1’ but such a source address never occurs on the network, then the rule would be an example of an Irrelevance anomaly.

[6] Al-Shaer, E.S. and Hamed, H.; “Firewall Policy Advisor for Anomaly Discovery and Rule Editing” Integrated Network Management VIII, pp. 17-30

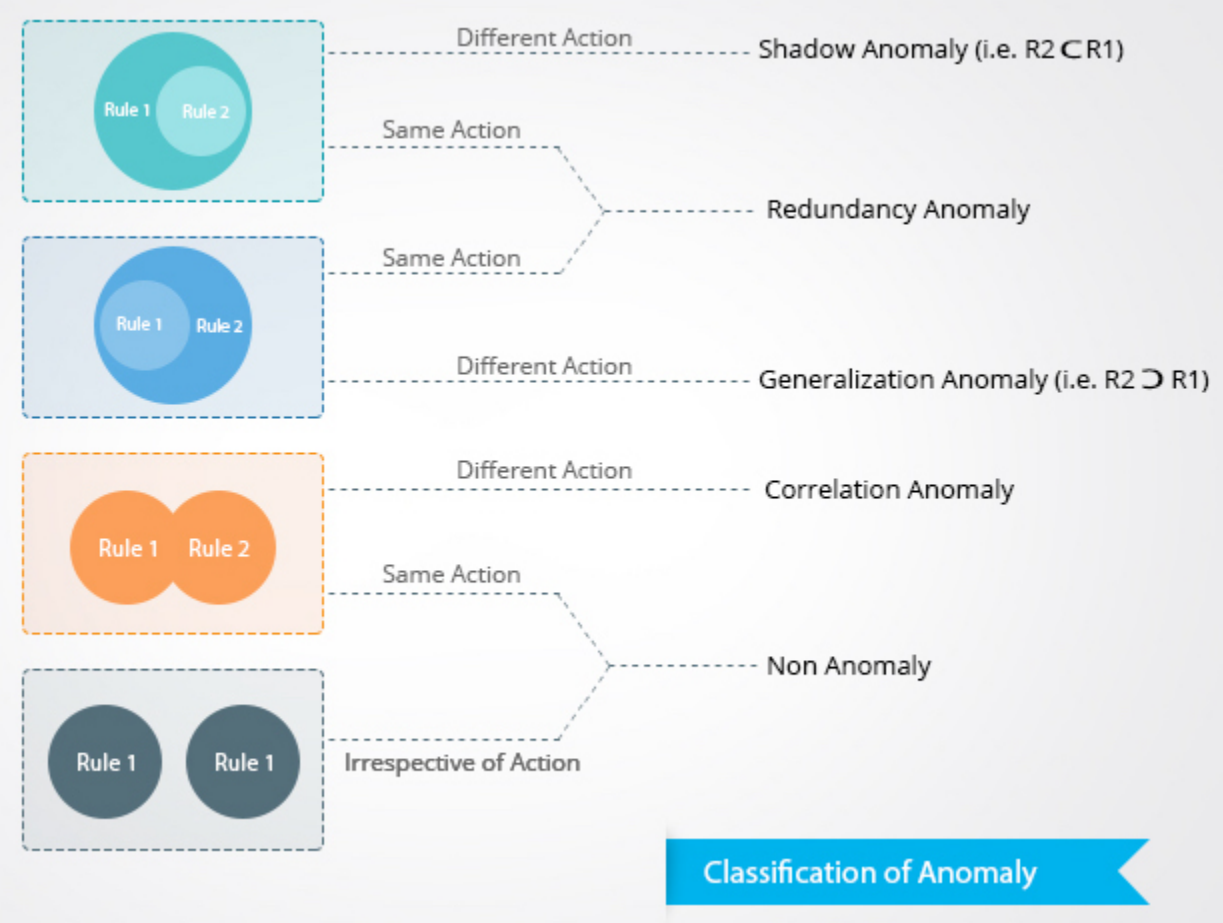

Figure 2: Firewall rule anomalies represented as Venn diagrams.

With the exception of Irrelevance anomalies, all anomalies require a sequence of two rules. In other words, a single firewall rule cannot be referred to as an 'anomaly.'

2.3 Palo Alto Networks Firewall Capabilities

GMU operates a Palo Alto Networks (PAN) firewall, model 5060. In this section, we describe some of the major features of the firewall and associated policy management software that are relevant to the task of assessing GMU NET's firewall policy.

2.3.1 Firewall Security Features

PAN advertises the following security features (among many others) that can be used to filter network traffic. We will describe policy management features in the next section.

- ***Application and User Visibility***: this feature allows the firewall to surface "information about the applications, users and content"[7] moving through the network in which it is installed. The advertised benefit of this feature is to provide filtering capabilities beyond those based only on IP addresses, port numbers, and communication protocols. For example, these features allow the ability to reject traffic based on the types of files that are being sent or the applications that are sending them.
- ***Advanced Persistent Threat Prevention***: PAN describes this feature as providing the ability to automatically detect and begin protecting against suspicious network activity without having to manually create a firewall rule.
- ***Intrusion Prevention Systems***: these technologies are intended to detect and prevent sophisticated attacks using statistical, heuristic, or IP address-based methods used by adversaries with well-developed attack capabilities.
- ***URL Filtering***: PAN offers the ability to filter network traffic based on URLs, which, for example, allows network engineers to 'whitelist' or 'blacklist' specific internet addresses.

2.3.2 Firewall Management Features

In addition to features that are directly responsible for the filtering of internet traffic, many firewalls include software to help administrators manage their firewall and its policies. This section describes some of PAN's firewall management features.

- ***Web and client-based management applications***: PAN's management software can be accessed via web- and client-based applications. This management software allows administrators to create, edit, delete firewall policies; review logs and reports; monitor network traffic; and view network traffic analysis produced by PAN firewall analysis applications.
- ***Logging and reporting***: PAN management software provides automatically generated reporting of anomalous and suspicious activity, in addition to traditional logging of traffic activity and firewall responses
- ***Firewall rule qualification***: similar to modern software development environments, PAN firewalls allow administrators to qualify their rules before implementing them in production environments to observe effect and prevent unintended consequences.
- ***Real-traffic analysis***: PAN provides the ability to managers to monitor traffic through the firewall in near-real-time.
- ***Automated correlation analysis***: PAN firewalls are also capable of correlation-based analysis methods that are intended to algorithmically identify potential threatening activity, thereby alleviating manual data mining of network traffic. For example, if managers define a threat signature in terms of a high number of requests on a specific

[7] https://www.paloaltonetworks.com/products/features/application-visibility.html, accessed on November 14, 2015.

IP address, then, when a sufficient amount of this behavior is detected in traffic logs, the firewall will automatically generate a notification indicating that a threat was detected.

2.4 Roles and Responsibilities

In this section we will describe the roles and responsibilities of those people and organizations connected to and influenced by GMU's firewall policy management activities.

2.4.1 Stakeholders

Those people who may be affected by changes to GMU network firewall policies are considered 'stakeholders.' Stakeholders in this sense have a direct interest in the outcomes of the management of GMU network policies.

Defining the scope of stakeholders can be challenging in the case of internet-based systems because the ramifications of internet security are so far reaching. For example, some of the student and staff information that GMU maintains is considered sensitive. Unauthorized access to such data could result in very serious consequences for those whose data is entrusted to GMU. It is conceivable that GMU could be held liable for the breach of that data if it were to be shown that GMU failed to adequately protect it. In this sense, GMU as an institution, anyone who has data stored on GMU information systems, and all those who would be affected by major breaches in GMU security could arguably be considered stakeholders. However, for the purposes of this project, we define 'stakeholders' in a very narrow sense in order to focus our discussion on those most directly affected by GMU's firewall policy management activities.

Examples of direct stakeholders include, GMU Network Engineering (NET) office and its staff, GMU Information Technology Security (ITS), network zone points-of-contact and administrators, internet application owners, application requestors, and application users. For this document, it is these stakeholders only that will be meant whenever we use the term "stakeholders" within this document. Exceptions to this list will be identified where they occur.

2.4.1.1 GMU NET Office, especially Network Engineering and its Staff

GMU NET's mission is "NET's mission is to provide reliable, consistent and effective communications throughout George Mason University's data, voice, and video networks and to evaluate and implement new networking technologies that support and enhance the goals of the University."[8]

In order to meet their mission, GMU NET Network Engineering operates a service that allows users to submit help 'tickets' to keep track of requested activities, which we will commonly as refer to 'GMU NET's ticketing system,' their 'help desk ticketing system,' or similar descriptions.

[8] http://tsdnet.gmu.edu, accessed on November 14, 2015.

The basic features of GMU NET's ticketing system are likely familiar to many people who use computers. In short, users submit requests for services to be provided by GMU NET and its staff; these requests take the form of tickets that are queued in a first-in-first-out system; tickets are handled by GMU NET staff. Once the service has been delivered satisfactorily the request is 'closed,' indicating that the work has been completed.

2.4.1.2 GMU IT Security Office and Staff

GMU ITS Office[9] provides IT security services to GMU departments, employees, students, and those otherwise affiliated with the university. One of the services that security analysts in this office provide is an assessment of the security risk associated with requests for applications and recommendations for changes to firewall policy.

2.4.1.3 Network Zone Points-of-contact and administrators

As described above, GMU largely manages their firewall rules using zones. We assume for this project that they are used universally. Zones typically align well with major colleges or schools, but not always. For example, the Volgenau School of Engineering (VSE) has an administrator and staff for its computing resources, including its network[10] and would largely have responsibility for the networked resources in his or her zone. Exceptions to this rule include cases where departments operate networked computers in buildings owned and operated by another department, or when external organizations have access to university computing facilities remotely from machines off-campus or connected to the campus via some other means.

However, for simplicity, we generally assume for this project that zone points-of-contact correspond to departments, schools, and other distinguishable groups within the university.

2.4.1.4 Application Owners

Generally, the cases that this project is most directly concerned with are those where new or existing networked applications or hardware are placed on GMU's computer communications network or where existing applications or hardware require modifications to to the way they operate within the GMU's computer network. For example, email servers receive and send internet communications to other email servers. If a new email server were to be stood up, then the owner of that server, specifically the email application on it, would be required to register that application with GMU NET. This registration may result in the creation of new firewall policies or the modification of existing ones.

Application owners typically submit service requests through the network zone points-of-contact in which their particular application will operate. Application owners are responsible for

[9] http://itsecurity.gmu.edu, accessed on November 14, 2015.
[10] https://volgenau.gmu.edu/administration/director-computing-resources, accessed November 14, 2015.

ensuring that their application delivers the intended service and will work with zone POCs and GMU NET to ensure that it does, including ensuring that firewall policies permit the intended traffic.

2.4.1.5 Application Requestors

As alluded to above, in some cases, those who make requests to register applications on GMU's network are not those who own the application. Application requestors must have the ability to submit requests through GMU NET's help desk ticketing system. Application requestors are usually Zone POCs or application owners.

2.4.1.6 Application Users

Application users are those who are intended to use the networked application. Application users may exist within or outside of GMU network, or may exist within or outside of the same zone as the application location zone.

3 Configuration Management Processes

This section describes fundamental concepts of configuration management (CM), including definitions provided by standards bodies. This section also describes common features of CM procedures and the value of CM processes to those who apply them. Finally, we discuss why CM practices are relevant to firewall rule sets and some prominent risks associated with not following following CM 'best-practices.'

3.1 Configuration Management Fundamentals

CM is widely used in engineering disciplines to ensure that changes made to systems are well-managed. Managing configurations of systems ensures that changes are not only well-documented, authorized, and auditable, but more importantly, that those changes are beneficial and deliberate. Many engineering societies such as the Institute of Electrical and Electronics Engineers (IEEE), the International Organization of Standardization (ISO), the Software Engineering Institute (SEI), and the American National Standards Institute (ANSI), and Electronics and Information Technology Association (EIA) offer their own definitions of CM, often written as functions that have inputs, activities, and outputs. We review prominent definitions of CM below.

IEEE STD 828-2012 governs software configuration management (SCM) plans, and describes SCM activities with the following categories:

> "SCM activities include the identification and establishment of baselines; the review, approval, and control of changes; the tracking and reporting of such changes; the audits and reviews of the evolving software product; the management of software release and delivery activities, and the control of interface documentation and project supplier SCM."

Whether managing a firewall rule set should be considered a valid application of the SCM activities described in the standard above is debatable; however, we believe that in the absence of an IEEE standard governing firewall rule-set management, SCM standards are applicable and adequate for managing firewall rules. Of the activities mentioned above, all of them are applicable to the

management of firewall policies, with the possible exception of 'software release and delivery activities.' However, even in this case, the creation or modification of firewall policies could have similar effects as would the deployment of new software capability for users. For example, if a web application developer were to make changes to the software that governs the web application server, then that change could impact the functionality of the web application for some users. Similarly, deployment of new or modified firewall rules could also impact services for some users. Examples like these support the idea that firewall rules can be managed in a way similar to traditional software.

SEI's definition of CM is broader but similar in many ways to the IEEE description of activities:

> "Configuration management (CM) refers to a discipline for evaluating, coordinating, approving or disapproving, and implementing changes in artifacts that are used to construct and maintain software systems. An artifact may be a piece of hardware or software or documentation. CM enables the management of artifacts from the initial concept through design, implementation, testing, baselining, building, release, and maintenance." [11]

This definition is applicable to both the software and hardware features of 'software systems.' Because internet communication systems rely on both hardware and software to enable the infrastructure, this definition applies.

Both the IEEE and SEI definitions focus on activities associated with CM. Both definitions include the general concepts of the deliberate review and approval of changes to software systems. The IEEE definition includes references to the documentation and communication of the changes that are being made to the software system, whereas SEI's definition focuses more on the processes associated with the changes to the software system themselves, and less on documentation and communication.

EIA provides a fairly generic definition of CM:

> "A process that establishes and maintains consistency of a product's attributes with its requirements and product configuration information throughout the product's life cycle." [12]

EIA's definition can be applied to products of many types, including, but not limited to software products. All of the definitions include some reference to the maintenance over time of the system under configuration management, but the EIA's definition use of the term 'product life cycle' emphasizing that configuration management is an ongoing activity over the course of the system's existence, including the design and sustainment/maintenance phases of the system.

3.2 Purposes and Benefits of Configuration Management

Managers who are considering the implementation of specific CM processes and techniques rightly want to understand the costs and benefits that alternative approaches will have on their system. In this section we describe the purposes and benefits of CM processes, and the costs in time, manpower, and effort they entail.

---

[11] http://www.sei.cmu.edu/productlines/frame_report/config.man.htm, accessed November 7, 2015

[12] ANSI/EIA-649B (2011)

3.2.1 Purposes of Configuration Management

As implied by the definitions of CM provided above, there are many purposes of CM. We will discuss some of the most significant ones below.

3.2.1.1 Understand and know the state of the system

Perhaps most fundamentally, CM techniques are designed to help system managers and their stakeholders understand and manage their systems. This is a very difficult task for any system of even modest complexity, but especially challenging for large-scale internet communications systems such as GMU's diverse collection of computing systems. These systems exist across geographically different campuses, link to constituencies world-wide, and require flexibility since collaboration is a critical attribute of university services. Increasingly, collaboration depends on the delivery of internet services.

As the number of applications on a university's network grows, the complexity of the firewall policies is also likely to increase. As the complexity of firewall policies increases, it is increasingly difficult to know the state of the overall network, and, more relevant in this project's case, whether the firewall policies are meeting the goals of the network and application stakeholders.

3.2.1.2 Manage changes to the system

CM activities as applied to firewall policies are intended to help firewall managers know the state of their policies and whether they are meeting the required and intended design, operation, and maintenance goals of their network system.

3.3 Description of GMU's Service Ticketing System

GMU NET uses a help desk ticketing system to help manage how they provide service to those requesting it. GMU NET's responsibilities extend beyond those related to the creation or modification of firewall rules to include video and voice data services, hardware installation, and other services. However, this project is focused on their management of firewall rules from the point at which GMU customers request service to the point at which the service is delivered. This section will describe how the ticketing system is used, definitions for roles in the ticketing process, decisions that affect the creation or modification of firewall rules, and how tickets have been reported to be closed.

3.3.1 Sources of information for GMU NET's firewall rule management processes

Most of the information presented in this section has been acquired through interviews with GMU staff. Several factors have contributed to this fact. First, GMU NET has no documentation that describes the processes it follows in the administration of its help desk ticketing system,

though the development of documentation is reported by those we interviewed to be in development.

Second, many of the documents or information that would be helpful to understand GMU's network, specific firewall rules, or network zones were not available to review due to the sensitivity of the documents. This project attempted several methods for mitigating risks associated with exposing this information to the project team, such as non-disclosure agreements, a virtual computer environment, or a reading room. However, in all cases, except for the eventual acquisition of firewall traffic logs, we were unable acquire documentation describing processes and policies related to the configuration management of GMU network firewall rules.

Third, our project client was very involved in the initial scoping project because of the variety of project directions available to be pursued and the relative lack of management documentation associated with the GMU firewall policies and their management. As part of that scoping, our project client provided helpful background information on internet network communications and configurations to help orient our efforts in useful directions, given that the project team does not have expertise in the specifics of computer and internet network engineering.

3.3.2 GMU NET's Help Desk Ticketing Processes

Since GMU's ticketing system must be flexible enough to receive many kinds of requests and they are not equipped to develop purpose-built ticketing systems, there are very few data entry fields in the ticketing systems that are specifically tailored to each kind of request. As a result, many of the data-entry fields are generic and therefore not specific to requests related to firewalls.

Figure 3 is a SysML Activity Diagram which represents GMU's help desk ticketing process related to the creation, modification, and (conceivably) deletion of firewall rules, although we haven't observed any instances of firewall rules being deleted using this process.

Those requests that are identified by GMU NET's engineers as impacting firewall rules will initiate the actions that are described in Figure 3. Our interviews with NET staff uncovered no formal procedures for determining whether a request would impact firewall rules. For NET staff, it may be obvious which requests could result in firewall actions, such as connecting new web applications or email servers; however, a lack of documentation prevents this team from representing help desk processes in accordance with established procedures.

Based on our interviews, firewall-related requests are reviewed to determine which existing firewall rules (if any) are potentially affected. NET engineers conduct a manual search of the existing firewall policies to determine whether inbound or outbound traffic is included in any existing policies. In cases where insufficient information was provided by the service-requestor, the NET engineer will communicate with the requestor to get relevant information such as IP addresses, communication protocols, and access limitations.

If the NET engineer's manual search reveals that no existing firewall policies match the search criteria, then he or she will create a candidate rule or rules to meet the requestors' needs. The NET engineer may need to communicate with the requestor to get information required to create the rule if the request ticket did not contain enough relevant information. If the NET engineer can use his or her experience to create a candidate rule, then they may do so.

Once a candidate rule is created, then in some cases, the rule will undergo a review from the IT Security office. However, it is not clear, based on interviews, the criteria used to determine if an ITSO review is required. As far we could determine, there are no documented criteria that, if met, would require ITSO review. GMU NET is likely relying on the expertise of its engineers to make such a decision. In interviews, topics such as the sensitivity of the zone in which the new application would exist or the nature of the data might be such triggering criteria; however, we were not able to verify those potential criteria against GMU NET or ITSO policy or procedures.

In the event that ITSO is determined to need to be consulted for a review, then either a candidate rule or similar proposal will be sent to their office for a determination of security risk. GMU NET expects to receive a security recommendation related to the candidate firewall policy. If required, changes to the rule will be made prior to implementation that reflect GMU ITSO judgments.

Once a finalized rule is determined, it will be implemented into the firewall policies. GMU NET engineers will work with either the requestor or the application owner to determine whether the application is behaving properly. In some cases, errors may prevent the application from working properly. In such cases, GME NET engineers will troubleshoot to find the cause of the problems before closing the service request. In cases where firewall rules are preventing the application from functioning properly, changes to firewall rules will be made (e.g. changing the order of firewall rule application) to support the new application, unless firewall rule modifications would trigger a security review.

Once the firewall rule policy or polices are implemented properly and the application is functioning as required, then the request ticket will be closed. If subsequent service related to the application is required, then either the old ticket could be re-opened or a new service request could be created.

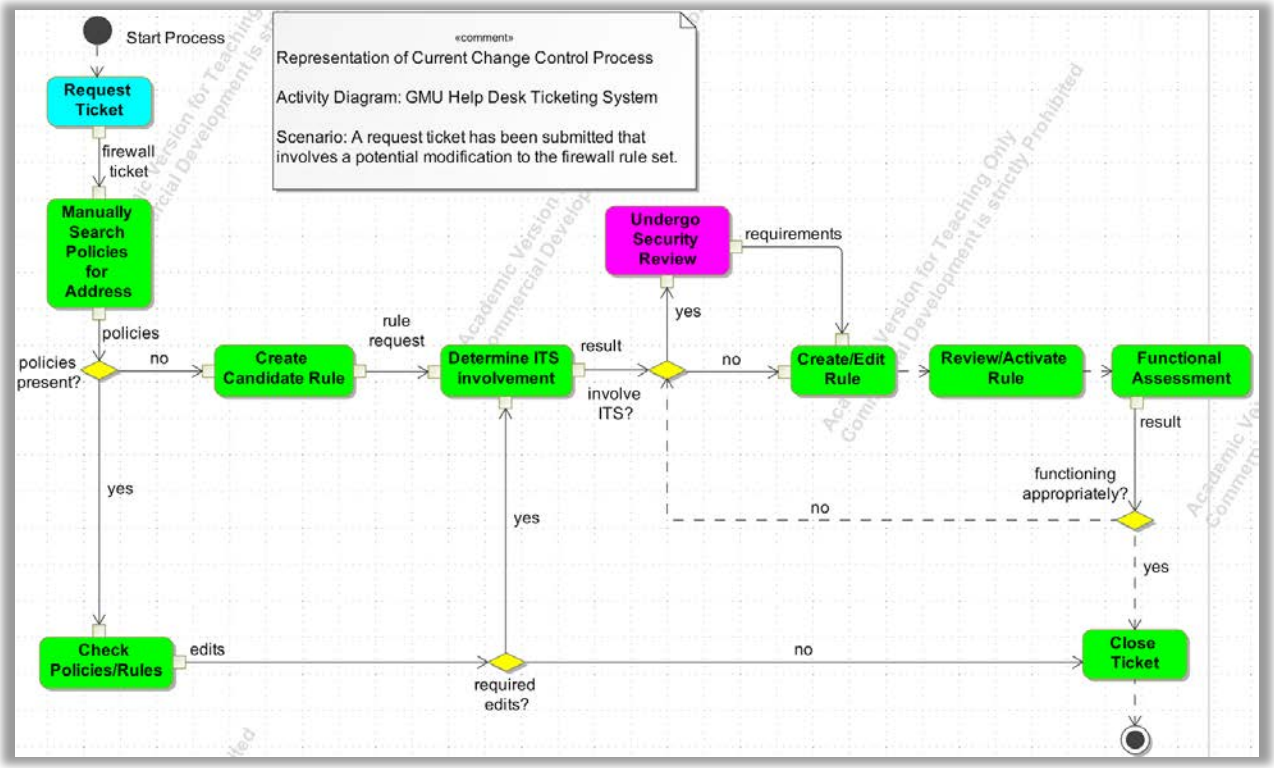

Figure 3: SysML model representation of help-desk tickets that affect firewall rules

4 Configuration Management Analysis

4.1 Criteria Used to Assess GMU Configuration Management Activities

The criteria used to evaluate the GMU management activities was based on the IEEE Standard 828-2012, NIST 800-128, and IEEE Journal SE-10. These documents present best practices for configuration management. The criteria were selected by comparing the processes of the research and selecting the items that were most relevant to our client's needs, which at a minimum included configuration change control and configuration auditing. These criteria were included into the areas of configuration management processes that are required for establishing and protecting the integrity of the configured product from the determination of need to it's eventual retirement. The processes include configuration identification, configuration change control, configuration status accounting, configuration auditing, and configuration release management.

The identified processes of configuration change control and configuration auditing were then decomposed to address our client's needs as detailed below. Figure 4 illustrates the areas that our work was focused. The green colored text reflects areas that were investigated, researched, and analyzed. The red colored text reflects areas that were not able to be addressed. Those topics can be used for future research as detailed later in this paper.

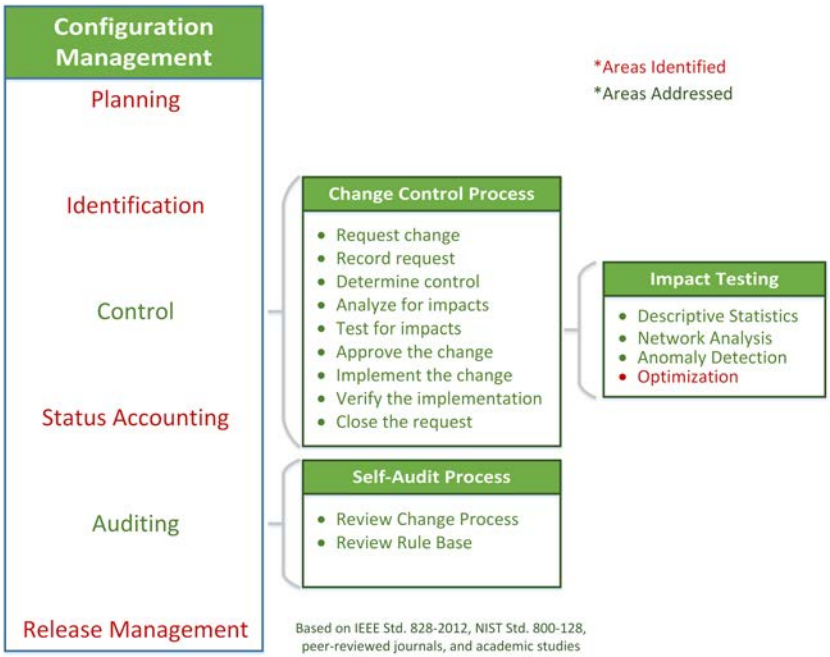

Figure 4. Breakdown of configuration management work areas

4.1.1 Configuration Management

In order to develop an audit process for the firewall rule set and assess the change control process of firewall change requests, research for best practices and standards was undertaken to build on preexisting academic studies. Upon further review, it became clear that configuration control processes and audits are fundamental processes included in the broader subject of configuration management. Broadening our scope to researching configuration management standards, two standards were subjectively reviewed to be pertinent. One standard is IEEE 828-2012[13], "IEEE Standard for Configuration Management in Systems and Software Engineering". The other standard is NIST 800-128[14], "Guide for Security-Focused Configuration Management of Information Systems." Additionally, peer-reviewed literature was used to develop aspects pertinent for auditing a set of firewall rules. The combination of these sources comprised the backbone of our assessment, analysis, and final recommendations.

4.1.2 Configuration Control

The following sections detail sub-processes used to control changes to designated resources. These aim to protect the integrity of the resources from requirements inception to retirement and provide documentation to the processes and responsibilities for those involved.

[13] IEEE Standard for Configuration Management in Systems and Software Engineering," in IEEE Std 828-2012 (Revision of IEEE Std 828-2005) , vol., no., pp.1-71, March 16 2012
[14] Special Publication (NIST SP) - 800-128

4.1.2.1 Establish Change Control Infrastructure

The infrastructure needs to be established that details the configuration items, processes to be undertaken, desired outputs, and roles and responsibilities of those involved.

4.1.2.1.1 Designate items subject to change control

The configuration items should be designated to determine their level of configuration control. This can be an approach to target each control change, or designate items on a priority basis based on metrics such as time to completion, security risk, or regulatory compliance.

4.1.2.2 Establish change evaluation criteria and authorities

To ensure consistency in the evaluation and disposition of change requests, evaluation criteria should be pre-established to evaluate change requests. Those would be affected by configuration changes should be identified as authorities, and they should be represented in the control process. Documentation should also be included that empowers individuals to approve or reject changes.

4.1.2.3 Establish Change Request Form

A change request form should be used for formal change management. It should be incrementally populated as the request continues through it's lifecycle from conception to final disposition. The following fields are recommended as a baseline:

a. Description of proposed change and rationale/purpose
b. State of the change (e.g., open, approved/rejected, implemented, tested)
c. Affected baseline
d. Outcome of analysis of impact on the project
e. Resolution
f. Approvals

4.1.2.4 Control changes to constituent configuration items/baselines

Changes should be controlled with a formal process. This process should follow that change requests are formally requested, evaluated for their security impact, tested for effectiveness, and approved before they are implemented. The general steps are detailed in the sections below, but can vary based on the organizations risk tolerance and impact level.

4.1.2.4.1 Request the Change

Change requests should come from pre-defined persons or organizations. It is to be evaluated later if the request undergoes further change control scrutiny.

4.1.2.4.2 Record the request

A change request should be recorded in accordance with organizational procedures. They should then be routed based on the pre-determined workflow, and also allow for electronic notifications to provide acknowledgements, updates, and approvals.

4.1.2.4.3 Determine if change control is required

Some types of changes may be exempt from configuration change control or pre-approved. If the change is exempt or pre-approved, this should be noted on the change request. The change should be made without further analysis or approval; however, system documentation may still require updating.

4.1.2.4.4 Analyze the proposed change

For types of changes that are not exempt or pre-approved, the proposed change need to undergo an analysis. This analysis can consist of a lookup of existing resources or feasibility in implementation, as well as predict security and functional impacts.

4.1.2.4.5 Test the change for security and functional impacts

After the change was analyzed, the change should be tested to confirm the impacts identified during analysis, and possibly reveal additional impacts. The impacts of the change are presented to the authorized users for evaluation.

4.1.2.4.6 Approve the change

The authorizers can approve the change, and they may require the implementation of additional controls if the change is necessary but has a negative impact on the security of the system and organization.

4.1.2.4.7 Implement the approved change

Once approved, the appropriate staff makes the change. Stakeholders are notified about the change, especially if the change implementation requires a service interruption or alters the functionality of the information system.

4.1.2.4.8 Verify that the change was implemented correctly

Configuration change control is not complete and a change request not closed until it has been confirmed that the change was deployed without issues. Although the initial security impact analysis and testing may have found no impact from the change, an improperly implemented change can cause its own security issues.

4.1.2.4.9 Close out the change request

With completion of the above steps, the change request is closed out in accordance with organizational procedures.

4.1.2.4.10 Record and Archive

Once the change has been analyzed, approved, tested, implemented, and verified, the organization ensures that updates have been made to supporting documents such as technical designs and baseline configurations. As changes are made to baseline configurations, the new baseline becomes the current version, and the previous baseline is no longer valid but is retained for historical purposes. If there are issues with a change to the baseline, retention of previous versions allows for a restoration to a previous version of the baseline configuration.

Additionally, archiving previous baseline configurations is useful for incident response and traceability support during formal audits.

4.1.3 Auditing

Firewall auditing procedures assist in proving the security of a network by verifying the integrity and adherence to dedicated policies and procedures. Firewall audits in particular help ensure compliance with regulatory standards as well as internal and external security policies in terms of the resource configuration integrity and administrative processes. The Higher Education Compliance Alliance outlines several areas of compliance required by higher education. Particular standards that are more widely known are FERP, FISMA, HIPAA, DMCA, FACTA[15]. Without an institutions compliance, it could incur fines and penalties.

In 2013, Idaho State University incurred a $400,000 penalty for not complying with HIPAA standards. The trigger was that 17,500 patient records were breached and went undetected for at least 10 months. Additionally, the university was not compliant with the first requirements for Health Insurance Privacy and Accountability Act (HIPAA), and hadn't been for over 5 years[16].

By undergoing a self-audit, institutions can verify their compliance and identify areas for improvement. Two of the most important aspects of conducting a firewall audit are to review the change process and the rule base. These aspects are described in the following sections, and a sample self-audit questionnaire is provided in Appendix A: Firewall Rule-set Audit Checklist[17].

4.1.3.1 Review Change Process

A firewall audit should include an examination of the firewall change control process. The goal of this step is to make sure that requested changes were properly approved, implemented, and documented.

4.1.3.2 Review Rule Base

The audit should also review the firewall rule base. For these questions a ranking should be given based on the type of firewall and its placement in the infrastructure. For example, a firewall not connected to the Internet does not have the same risk as one that is connected to the Internet.

[15] http://www.nacua.org/documents/costofcompliance.pdf
[16] http://www.4medapproved.com/hitsecurity/new-hipaa-penalty-issued/
[17] http://www.crn.com/blogs-op-ed/231903353/how-to-conduct-a-firewall-audit.htm

4.1.3.2.1 Review Policy Maintenance and Design Practices

The rule base should be reviewed for policy maintenance and design practices that emphasize security. For example, this review should cover rules that allow permissive actions such as rules with 'Any' for a field, redundant rules, uncommented rules, and the quantity of rules over time.

4.1.3.2.2 Risk and Compliance

The rule base should also be reviewed for risk and compliance. This includes reviewing compliance with regulatory standards, access within rules to sensitive resources, and risky services communicating with the Internet. It is imperative to understand your compliance requirements and security posture to properly complete this review.

4.1.4 Additional Areas

Other areas of Configuration Management include Planning, Configuration Identification, Configuration Status Accounting, Release Management. These additional processes within configuration management were identified but not addressed in this analysis. These processes are recommended to be developed towards a configuration management strategy and documented in a formal Change Management Plan for GMU NET.

4.2 Evaluation of GMU Configuration Management Activities

Using the criteria developed in section 4.1 and the descriptions of GMU's network firewall configuration management activities distilled from interviews of GMU staff, we present an assessment GMU's configuration management activities. This assessment is presented in two parts. The first part will focus on change control activities and the second will center on auditing.

4.2.1 Establishing context for change management and auditing activities

GMU NET's documented change control processes are very limited, and recognized by GMU NET as such. As a result, our project tried to focus assessment on those areas that would allow GMU to make the most progress toward implementing a change management process. For example, taking a very broad perspective, change management is centered around knowing the state of a system (especially complex ones) and tracking the changes from known states in order have confidence in one's understanding of the system over time. From an administrative or management perspective, this often entails the development of baselines and keeping documentation of changes that have occurred since a baseline condition (or known state) has been established. Figure 5 represents a distillation of configuration management perspectives into two axes. First, as described above, configuration through time involves establishing an understanding of the state of a system and tracking changes to that known state over time (here, from left to right).

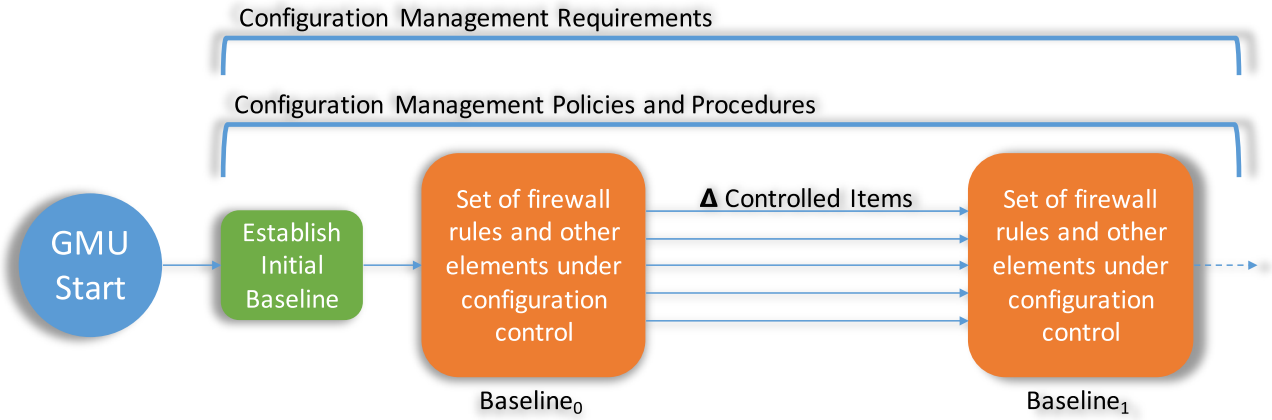

Figure 5: Overview of configuration management through time (horizontal) and from requirements to implementation (vertical)

The vertical axis refers to the hierarchy of management requirements, policies, and procedures that guide an organization's configuration management activities. Starting at the top, one must develop requirements that will inform policies and procedures. Configuration management policies, procedures and activities are context dependent. For example, the requirements of the health care industry are likely to be influenced by the Health Insurance Portability and Accountability Act[18] (HIPAA) and perhaps insurance regulations. Whereas in the financial sector, legislation such as Sarbanes-Oxley or Frank-Dodd might be the foundational, requirements-generating documents. To take an example on a much smaller scale, consider one's checking account and check register. In that case, the requirements are quite minimal: known one's account balances and outstanding payments.

Once requirements are established, policies and procedures can be developed to ensure the appropriate activities are undertaken to ensure the configuration of the system will be managed. Policies and procedures should document the prescribed activities of those executing a change management system. Since activities tend to impose a burden on an organization, only those activities that are essential to meeting requirements should be undertaken.

Additionally, these policies and procedures form the bedrock of auditing activities. Audits are generally conducted in order to ensure compliance with requirements, policies and procedures. If requirements, policies, and procedures are not established and documented, then auditing is of very limited value.

In the following sections we provide an assessment of GMU's firewall rule set change control activities. One theme that emerges is the tension between providing timely, responsive service, and developing, implementing, and adhering to change management practices that will almost certainly slow service responsiveness. Striking a prudent balance between the responsiveness and flexibility the absence of policies and procedures provide, and the knowledge, repeatability, and clarity those same policies and procedures could provide is difficult.

[18] http://www.educause.edu/library/health-insurance-portability-and-accountability-act-hipaa

However, failure to strike the appropriate balance can have notable consequences. Currently, GMU NET is managing with no documented requirements, policies, or procedures, which is likely not the ideal balance between the goals described above.

4.2.2 Assessment of change control activities

Change control generally starts with the establishment of a change control infrastructure, which begins with an identification of entities or systems whose state must be known and, therefore, whose changes must also be tracked closely. GMU NET (our client sponsor) identified firewall rules as the objects requiring configuration control.

Next, the literature recommends establishing the roles of those who will exercise configuration control and be responsible for configuration control. In some cases, configuration management literature will recommend at 'change control board' composed of stakeholders in a system. In other cases, configuration management can be implemented with smaller management structures. In GMU's case, firewall rules must be malleable enough to support a responsive internet communications service, but managed closely enough so that network managers can be confident in their knowledge of the state of the firewall policies. This tension is one between responsiveness and bureaucracy, and one's management approach can be thought of as falling on a continuum between those two poles.

Currently, GMU NET has no documented policies or procedures for the management of firewall rules. While this affords maximal flexibility to the NET engineering staff, it likely entails substantial risk to GMU in that one cannot be sure that the system is in fact functioning as intended. More specifically, GMU does not maintain a baseline rule set, does not seem to require documented approval of firewall rule changes, does not seem to identify change management roles (such as those permitted to change firewall rules, those whose approval must be received under specific conditions, etc.), or any other foundational or procedural change management infrastructure.

Recommendation: Develop a 'tiger team' (i.e. a small team focused on producing results quickly) to produce draft documentation of GMU's firewall rule set management processes. The initial effort should be focused on developing descriptions of 'as-is' behaviors and recommendations for 'to-be' procedures. The 'tiger-team' should be a cross-functional team composed of GMU NET, IT Security, major network administrators' staff, and other key stakeholders. They should be given an aggressive timeline to complete draft materials so as minimize the amount of time GMU will continue to operate with no documented firewall rule management procedures.

4.2.2.1 Establishment of a change request form

GMU NET maintains and uses a help desk ticketing system, which provides the means by which requests are tracked to completion. Although this system does provide a means to process requests, the system's 'general-purpose' nature does not seem to support important change

control functions, such as documenting those instances where firewall rules are created, changed or deleted in a way that can be easily reviewed or audited.

4.2.2.2 Control changes to all configuration items

The scope of our project was limited to firewall rules, which are the configuration items. Configuration items are those entities that whose configuration is being managed. There are some items whose configuration may not be managed. In each of the following subsections, we assess GMU NET's control processes in each of the areas.

4.2.2.2.1 Request the Change

Although GMU does have a mechanism to record service requests, the ticketing system doesn't to our knowledge allow requestors to set flags in the request that would allow auditors, managers, or others to easily find those requests that resulted in new or changed firewall rules.

Recommendation: Allow GMU NET engineers to indicate in the service ticket that a firewall rule was created, edited, or deleted. Similarly, allow application owners, or those submitting the service request to indicate that a firewall rule addition or change may be required.

4.2.2.2.2 Record the request

GMU does have a method for submitting, tracking and recording service requests that may result in firewall rules. However, because the service request does not capture if the service resulted in firewall rules, use of saved service tickets for auditing or review purposes is limited.

Recommendation: Ensure service tickets that resulted in a firewall addition or edit can be easily retrieved. For example, a properly constructed database should easily allow the creation of a field with a Boolean indicator of whether the service required modification of firewall rules.

4.2.2.2.3 Analyze the proposed change

In the cases where it is determined that the firewall rules need to be changed, then some degree of analysis should be conducted to ensure engineers and mangers understand the change. In our interviews, this process was handled entirely by the GMU NET engineers in the course of completing a service request. Analysis was very limited, and was in effect a verification that the application owner' service was functioning properly.

Recommendation: Develop a minimal firewall change analysis checklist that should be conducted whenever a firewall rule needs to be created or edited. Also, ensure that the service ticket reflects that the analysis was complete as part of the procedure of closing a service ticket that impacts firewall rules.

4.2.2.2.4 Test the change for security and functional impacts

According to interviews conducted with GMU NET engineers, little to no impact analysis of proposed firewall rules is conducted. None of our interviewees indicated that rules are tested in a testing environment or that any network analysis is conducted before completing the service

requests. There are many potential reasons for this lack of testing. For example, testing may require significant time and effort that would delay completing service requests an unacceptable amount.

However, even in cases where substantial firewall modifications occurred (as in the case of the installation of GMU's current firewall hardware solution) no evidence was presented that indicated that an analysis of the impact of adding or changing firewall rules was conducted.

In Section 5, we present several methods for conducting impact analysis. The methods we present are likely too time-consuming to be conducted for each or every service request, and therefore are likely more appropriate for periodic (e.g. semi-annual or annual) reviews.

Recommendation: develop and implement methods for analyzing the impact of firewall rules to be used in periodic reviews of firewall rules or during the process for providing service.

4.2.2.2.5 Approve the change

Except in those cases where an IT Security Office review is required (e.g. in those cases where service affects a network zone with higher security attention) no approval is currently required to make changes to firewall rules that are more restrictive than current rules. More generally, there are no documented and defined roles and responsibilities, or specific criteria that warrant escalation of approval.

Recommendation: Define roles and responsibilities for those who are involved in the addition or change of firewall rules. Specifically, define the authorities that each role has or doesn't have and under what conditions additional authority is required.

4.2.2.2.6 Implement the approved change

GMU NET engineers implement firewall rule changes. All GMU NET engineers who can provide service have access to the firewall policy administration tools provided by Palo Alto Networks, and are therefore eligible to make firewall rule-set changes. In those cases, where firewall rules need to be created or edited, network engineers work with the application owners to ensure that the application is working properly and that internet traffic is flowing to and from the application per the application owner's specification and in accordance with existing security requirements.

4.2.2.2.7 Verify that the change was implemented correctly

Verification that the changes to GMU's network firewall rules are largely limited to a functional verification conducted with an application owner that the application is operating. Our interviews revealed no indications that firewall rules changes were 'second-checked,' even in those cases where a firewall rule security review would be required.

Recommendation: In cases where extra caution is warranted, implement a second-check policy, to ensure that critical firewall rules are added or changed correctly. Although 'second-checks' do not always avoid errors, they do reduce their likelihood.

4.2.2.2.8 Close out the change request

After a functional verification of the application is performed and the application owner is satisfied that their application was placed on the network correctly, or is otherwise functioning, then the service ticket is closed.

Our interviews produced no evidence that network engineers follow a service request closing procedure. A closeout procedure may include documenting any system changes, problems that were encountered while providing the service, or any other relevant information that might be useful during subsequent reviews or service requests related to the original service request. Closing records related to service activities ensure that the records contain all required information per documented policies. When auditing these records, auditors should use existing policies, procedures, and other known requirements when assessing whether service requests contain all the information that is required.

Checklists are often used when closing out records. Checklists are generally updated as required to reflect changes in requirements, policies, and procedures. They also frequently include sections for secondary review and management approval of the work performed. However, because GMU has no documented procedures, none of these features of a document closeout procedure cold be assessed.

Recommendation: Develop close-out procedures based on firewall change management requirements, GMU NET policies and procedures. Additionally, consider using close-out checklists to ensure that service documentation contains required elements. Also, consider requiring supervisory or management approval of high-impact services or randomly selected service activities.

4.2.2.2.9 Record and Archive

The primary purpose of recording and archive service-related documents are so that they may be referred to at some point in the future as a self-sufficient representation and documentation of the work performed. Records should be kept in a format that supports subsequent review, including auditing. When a high-volume of records are being kept, then a method for retrieving relevant records should also be implemented. In this case, a method should exist for retrieving all service records that resulted in firewall rule changes or additions.

Recommendation: Provide methods for identifying and retrieving all service records that resulted in new or changed firewall rules.

4.2.2.2.10 Assessment of Auditing Procedures

GMU NET does not employ formal self-auditing procedures. This places risk on the resources in terms of network security and regulatory compliance. This also results in lack of insight into the

firewall rule set, which in turn grows into an unbound system. This unbound system can result in decreased firewall rule efficiency, decreased filtering performance, a larger rule set to manage, and an increase in unexpected behaviors. Auditing the rule set would at a minimum establish periodic baselines of the network security posture, analyze rule performance for integrity and effectiveness, and provide documentation pertinent for audit compliance. A sample of a self-auditing questionnaire is provided in Appendix 7.

Recommendation: Use the self-audit questionnaire and develop the necessary processes and authorizations in order to routinely audit the firewall rule set. Additionally, develop a cross-matrix that maps compliance requirements to existing firewall rules.

5 Impact Analysis

Analyzing the impact of changes to firewall rules is an important part of the change control process. This section demonstrates three approaches to conducting an analysis of the effects of changes to firewall rules: basic descriptive statistics, network analysis, and graphical modeling language – specifically, Petri Nets. Basic descriptive statistics of network traffic may provide network analysts insights into network traffic patters. For example, data may show the fraction of traffic currently being dropped in a zone and how much traffic might be dropped given a rule change. Second, network analysis could help analysts understand the structure of their network communications, as well as communication flow patterns. Finally, graphical models such as Petri nets, provide a visual representation of rules that could be used to detect firewall rule anomalies.

One the primary challenges associated with working with data and content associated with GMU network traffic is the sensitivity of the data. The project team worked with GMU NET and ITS to find common ground on access to data that would be helpful to the project but that would not divulge sensitive network security information. Given that the project is focused on the management of GMU's firewall rule policies, the project would have preferred to have access to the firewall rules for analysis. Unfortunately, such an arrangement was not achievable. Receiving data took considerable amounts of time due to several factors: identifying data that would potentially be useful, identifying a data provider, and ensuring that data that was provided didn't compromise GMU network security. Due to the security concerns, among other challenges of getting data, we didn't receive data until November, substantially limiting our options for the kinds and depth of analysis we could complete.

GMU NET and ITSO provided the project with a sample of network traffic as logged by the firewall. The sample was taken from 5 security zones (out of a total of approximately 250) for approximately 10 minutes of network activity. Additionally, in order to ensure that the data we used, this report does not expose sensitive information, we agreed to conceal actual IP addresses. In order to meet this requirement, we simply replaced unique IP addresses with unique integers. This approach, although straightforward, conceals potentially relevant information included in the IP address structure.

In order to demonstrate some of the potential drawbacks of our approach, it is important to understand the details of IP addresses. IP Addresses are often represented in 'dotted decimal' notation that have four 8-bit numbers (i.e. numbers between 0 and 255) separated by a decimal point. An example of IPV4 address is '192.10.10.80' IP addresses can be understood using postal addresses as an analogue. Postal addresses describe a unique physical mailing address just as IP addresses describe unique digital ones. Just like we could collect all physical mailing addresses in area code '22003' we could also collect all IP addresses in a range, such as '192.x.y.z' Our approach to conceal IP addresses potentially conceals activity that applies to ranges to IP addresses. Our team recognized this limitation and elected to both conceal IP addresses based on whether they are unique and work with their integer labels for two reasons. First, given the stage in the project at which we received data, we didn't have much time to devise more complicated methods for concealing IP addresses that preserve the IP address structure. Second, even if our team did have sufficient time, we wanted to adhere to the spirit of the arrangement we made that was a condition of using this data and not risk exposing sensitive network information.

The team also cautions against over-generalizing results gleaned from this data, since it represents such a small fraction of the total data. As a result, we emphasize analysis approaches that might be helpful to those who have access to all the data. Our hope is that even if our specific findings do not generalize to all data, our approaches do.

The purpose of this section is to describe the kinds of analysis approaches that we think would be helpful to GMU network engineers who want to understand not only their network traffic, but also the impact that new or changed firewall rules might have on traffic. In the first section, we show summaries of our data and provide basic descriptive statistics. In the next section, we provide examples of network analysis of the data, and in the third we demonstrate the use of Petri Nets to uncover firewall rule anomalies.

5.1 Descriptive Analysis

For this section, in addition to the data provided by GMU NET and ITSO, we use tableau software to develop the column charts and scatter plots. Under the license that was used, all data remained private. Also, we used Python do develop the box-and-whisker plots.

As mentioned above, we received data from 5 security zones. The total records for this activity represent communications that result in a 'drop' or 'allow' firewall action. Figure 6 shows the total amount of traffic per zone, as well as the total number of records (i.e. the 'grand total' on the farthest right of the figure).

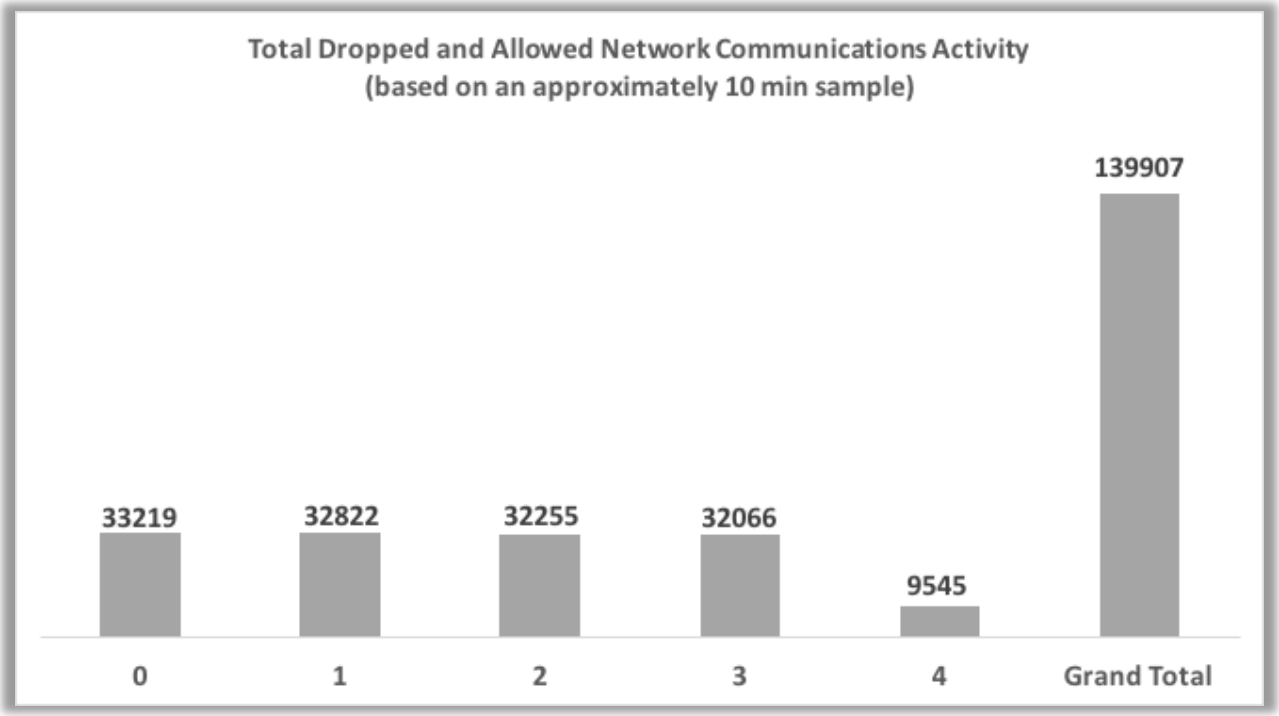

Figure 6: Total Traffic by Zone

Figure 7 shows what fraction of the sample traffic is dropped in each zone. The purpose of this graphic is to show two aspects of the sampled data: (1) the percent of dropped traffic for each zone, and (2) how much variability in dropped traffic each zone experienced over the sampled time period.

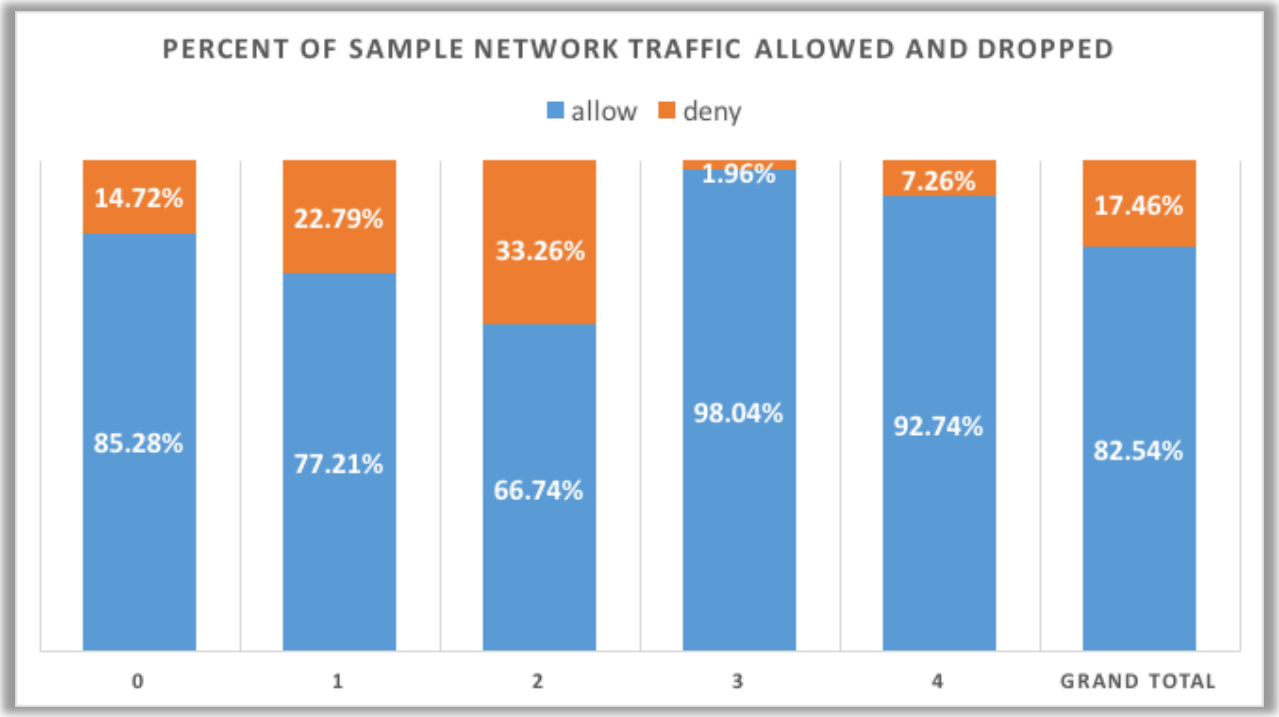

Figure 7: Percent of total traffic dropped and allowed

For another perspective of the variability in the amount of traffic that has been dropped, Figure 8 shows a box and whisker plot with median, mode, and expected values from the sample. As mentioned above, given the exceptionally small size of the sample, both from a perspective of the amount of time and number of zones sampled, we can't justify broader inferences. However, this approach may help administrators understand communication patterns within each of the zones and what impact the total firewall rule-set has on traffic.

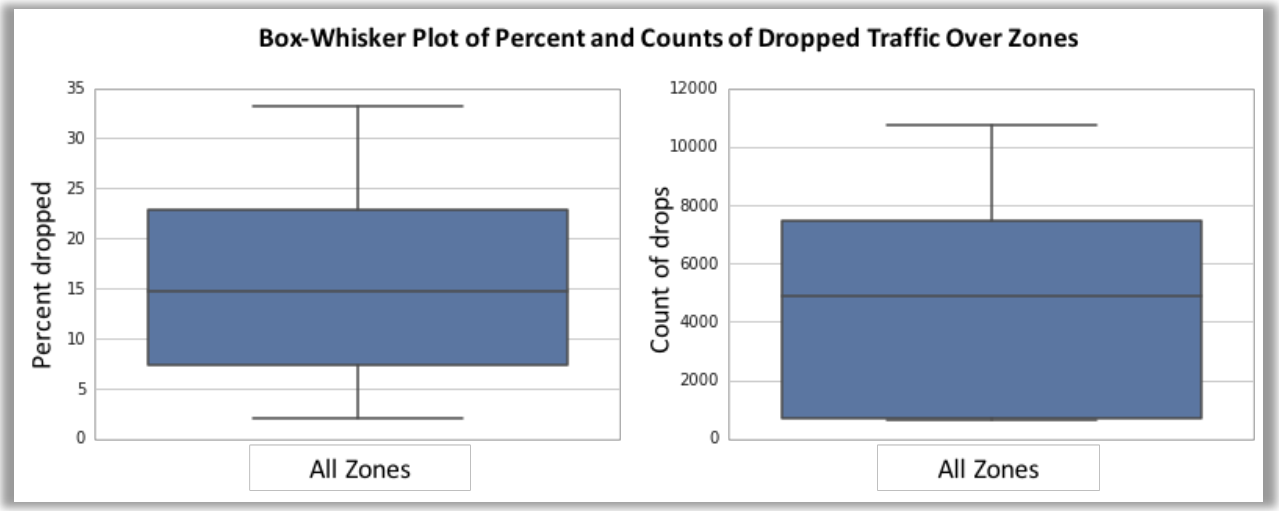

Figure 8: Box and whisker plots of dropped traffic

Analysts may also be interested in understanding the number of distinct source and destination IP addresses in each zone as a way of assessing the number or scope of firewall policies that either may be required to filter traffic. In the following (Figure 9) we show a basic scatter plot of the count of distinct destination vs. source IP addresses in each of the zones.

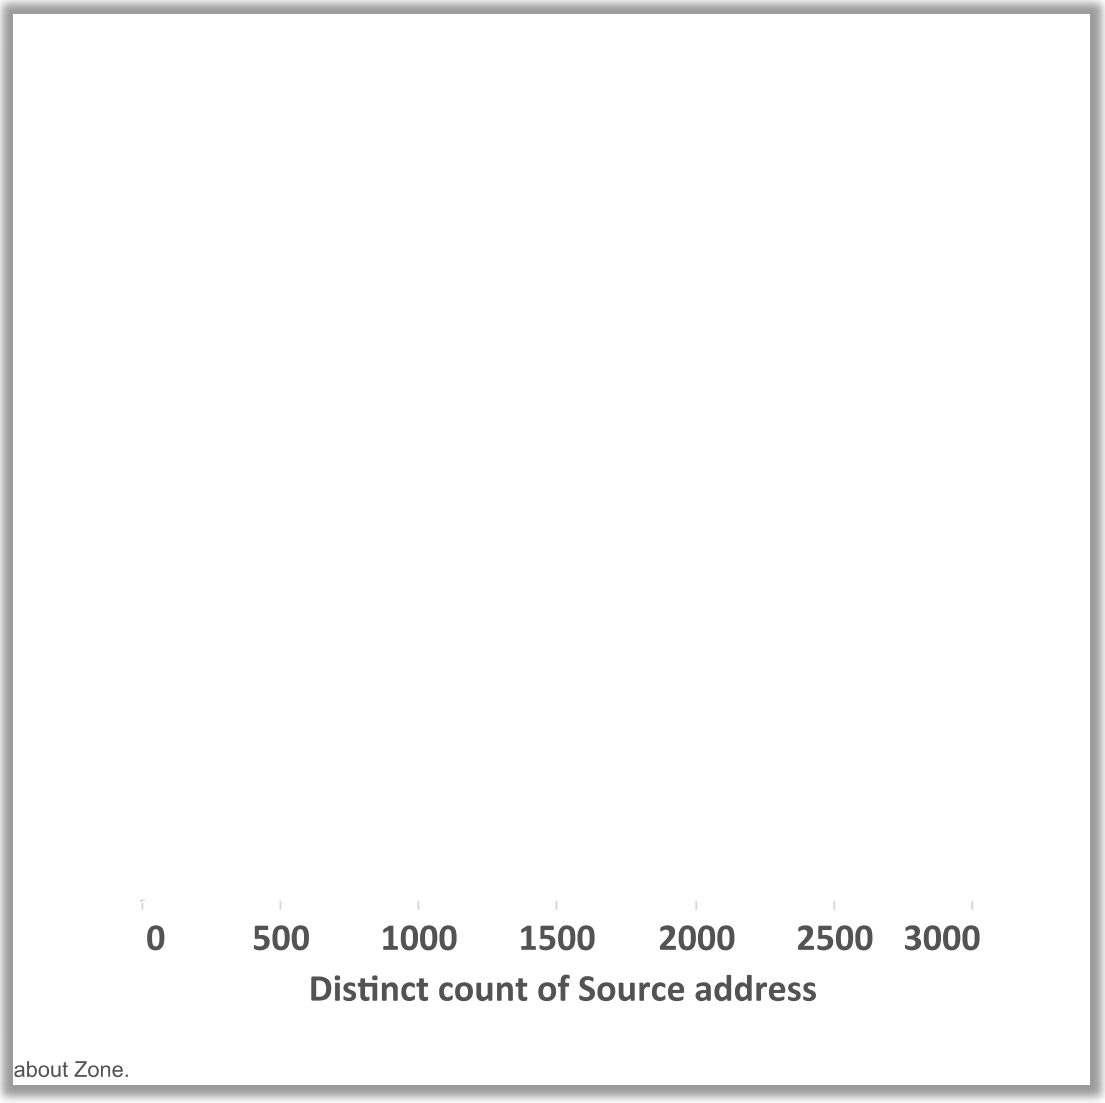

Figure 9: Counts of distinct source and destination IP addresses for each sample zone

Somewhat unsurprisingly, zone 4 has the fewest number of distinct addresses likely because it has the fewest number of records. However, zone 0 has far more unique source and destination addresses than zones 1, 2, or 3 even though all four zones have roughly the same number of records. In the following section, we will show how a network analysis may help develop insight into the structure of these communications.

Similar insights can be gained by looking at source and destination communication ports. Figure 10 shows a plot of the count of distinct destination communication ports vs. source ports.

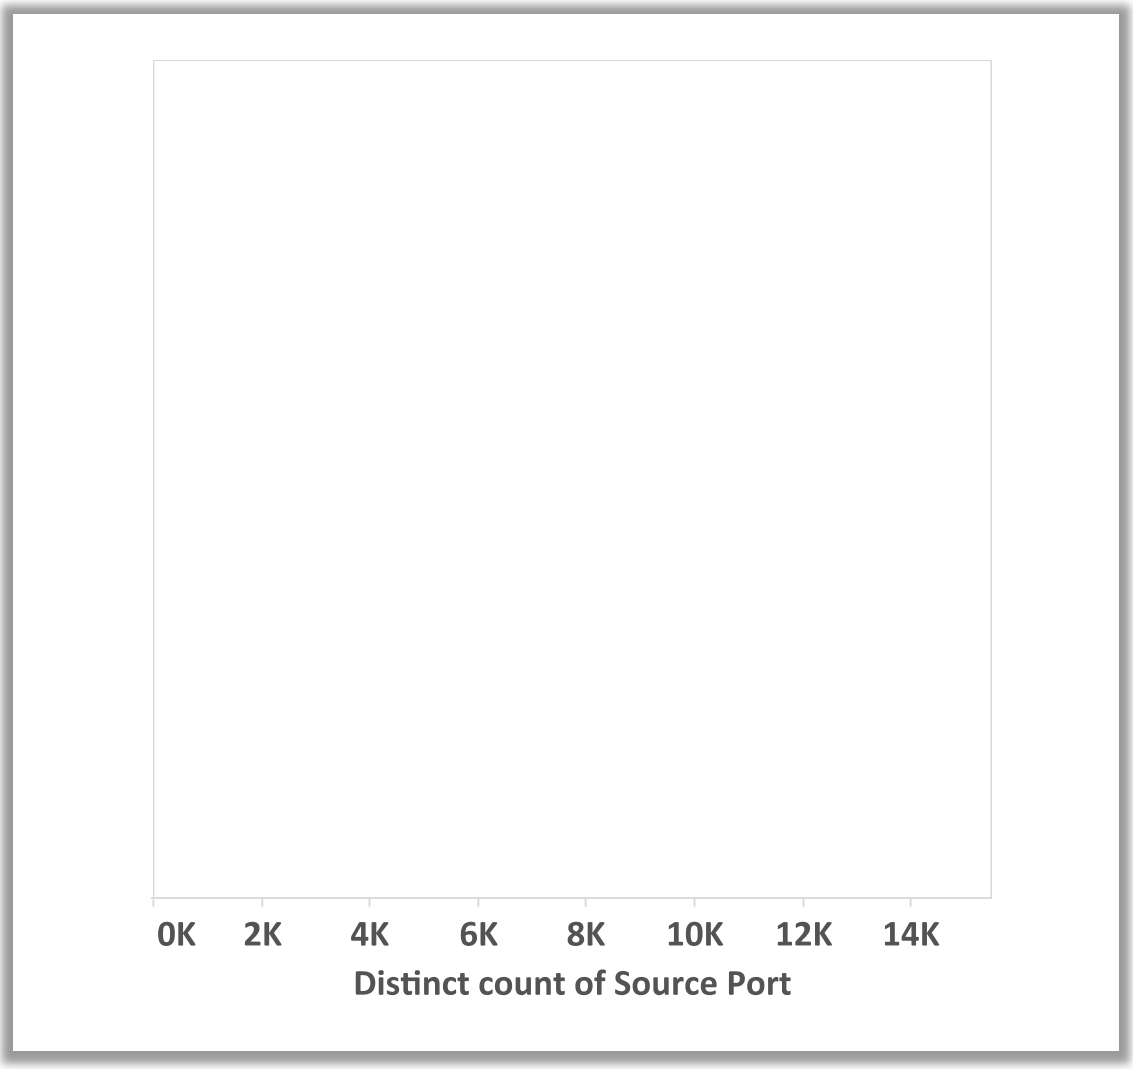

Figure 10: Counts of unique destination ports vs. source ports for each sample zone

The data for ports shares some broad trends with the IP address data. Zone 4 has the smallest number of distinct ports and zone 0 has a relatively large number of distinct ports compared to zones 1 – 3. Interestingly, zone 3 behaves notably differently in the case of IP addresses (where the number of distinct destination addresses was large compared to the number of distinct source addresses) than it does in the case of communication ports, where the number of distinct *source* ports is high relative to the number of destination ports.

The network traffic data we received also provides information about the communication protocols that were being used in each zone. The following figure (Figure 11) shows a breakdown of the kinds of communication protocols used in each zone as a percentage of the total traffic.

Figure 11: Communication protocol breakdown by type and zone

As described in our introduction, one of our purposes in this section is to describe the kinds of analyses that might be helpful to network analysts when developing firewall policy strategies. Of note in Figure 11 is that significant fractions of network activity is marked as 'N/A,' which are highlighted in Figure 12, perhaps limiting the ability of network analysts to understand network behaviors.

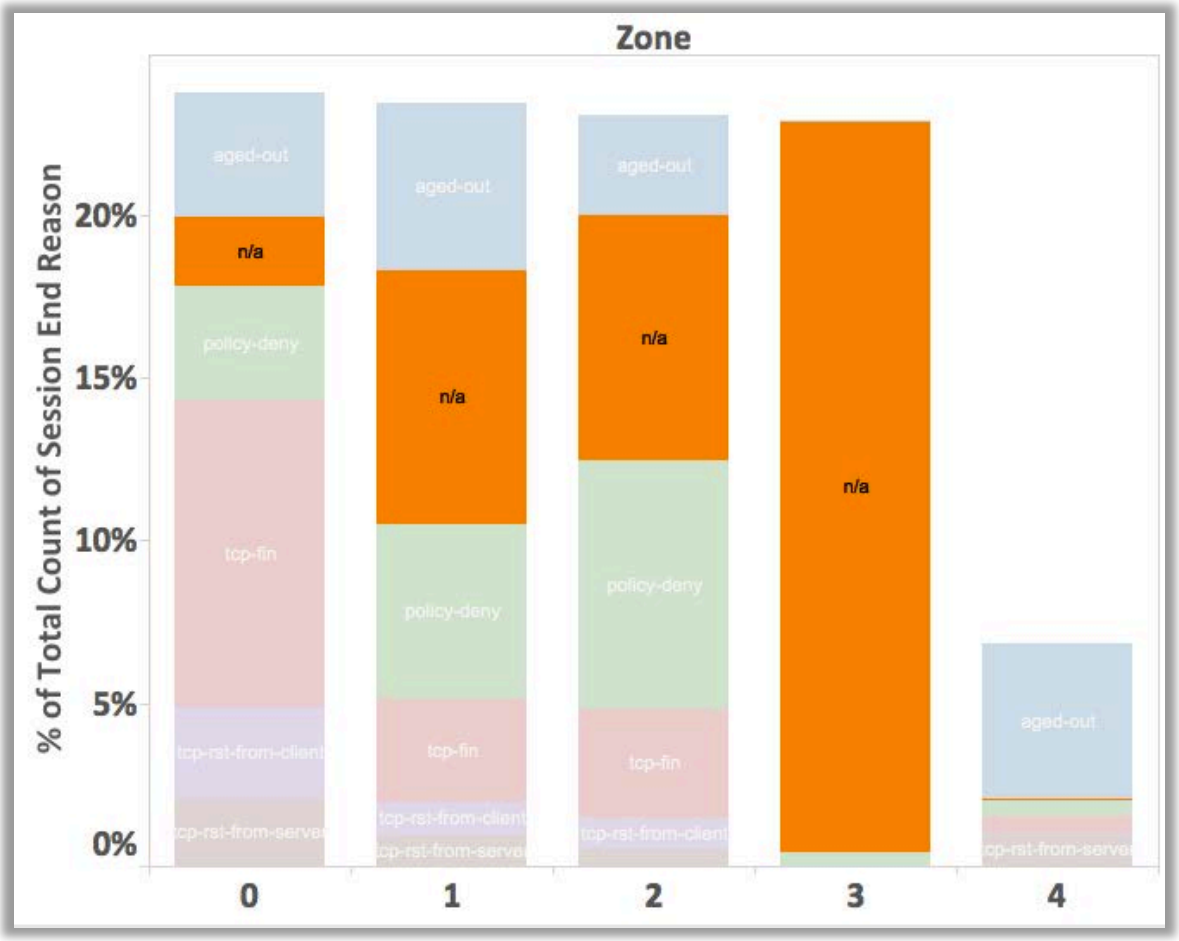

Figure 12: Percent of traffic labeled as 'n/a' in sample data

Also of potential interest are data that differ from zone-to-zone (as we saw above). In Figure 13 we see that the percent of traffic that terminates due to ‘tcp-fin’ in zone 0 is much larger than in any other zone.

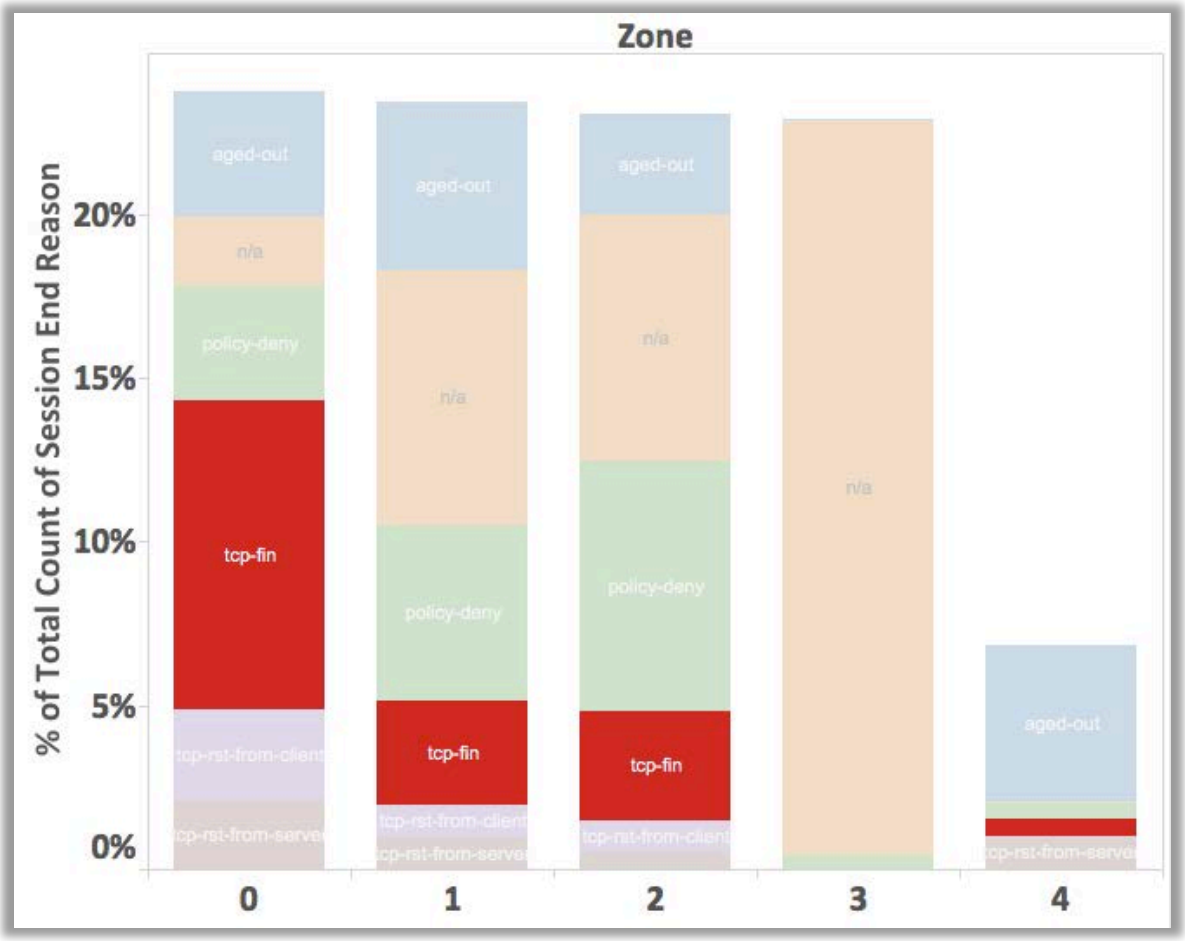

Figure 13: 'tcp-fin' traffic in each zone

The last approach in this line of thinking that we wanted to highlight that may be of use to network analysts is the percent of traffic that is dropped as a result of firewall rules. These kind of data may be helpful to analysts assessing traffic risk, where larger fractions of 'drop' traffic may indicate more relative risk or greater reliance on firewall policies to offer protection. Figure 14 highlights traffic dropped to a firewall policy.

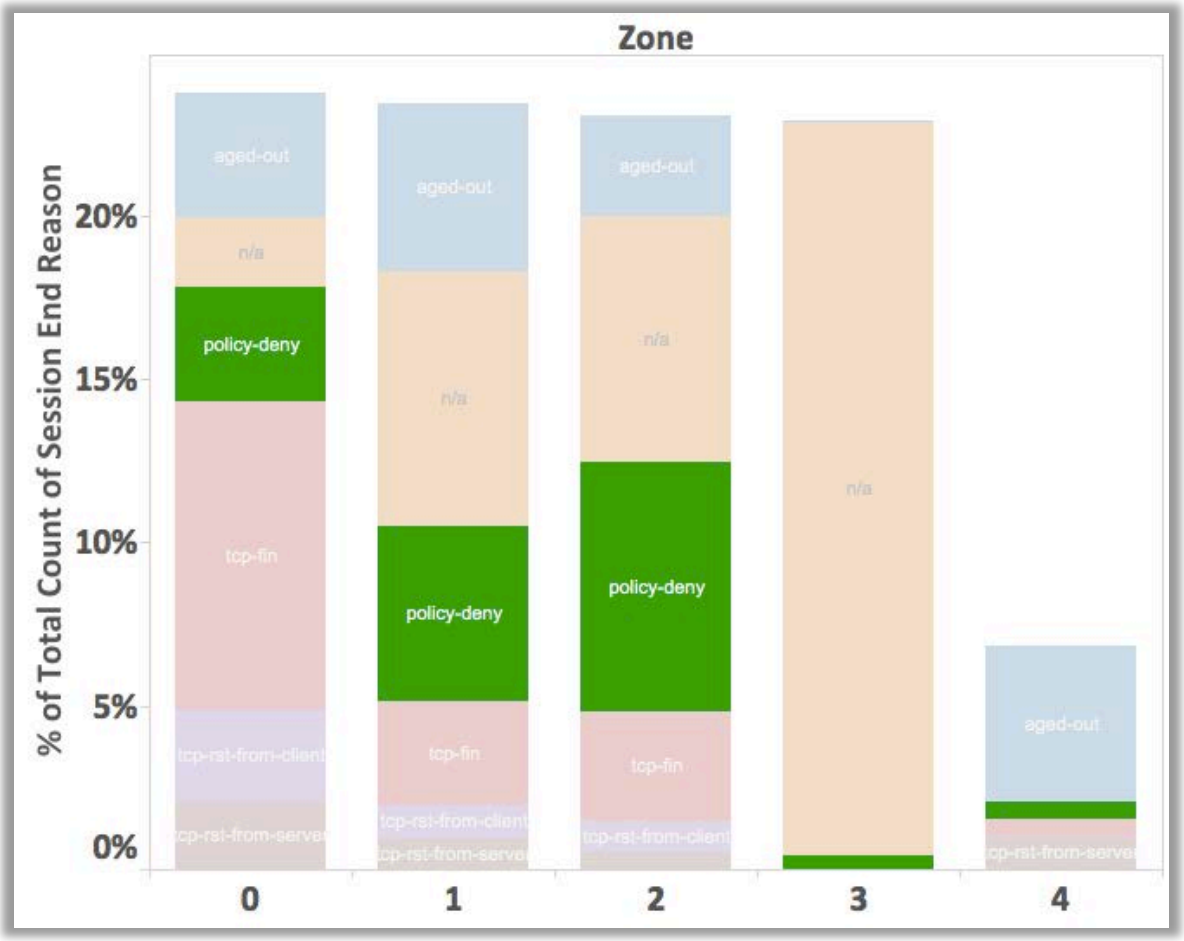

Figure 14: Percent of traffic dropped due to firewall policies

As can be seen in the previous figure, zone 2 has the largest fraction of 'policy-deny' traffic among all the zones of comparable sample size (i.e. zones 0 – 3). Using this information, analysts may have sufficient reason to investigate network traffic more thoroughly in that zone and perhaps develop insights that would contribute to more effective firewall policy implementation.

Up to this point, we have focused on analysis approaches that use all of the sample data, whether dropped or allowed by firewall policy. A natural next step is to attempt to extract features of dropped traffic beyond what the policy would typically specify. Since network administrators would know their firewall policies we will avoid discussing here approaches that attempt to infer those rules. Instead, the following charts and discussion are focused on number of packets and the size of the communication event.

5.2 Network Analysis

In addition to basic descriptive statistics of the firewall log traffic, we also wanted to demonstrate how simple network analyses could be used to help understand the traffic of individual zones in order to assess aspects of network traffic that may be affected by the implementation of new or changed firewall rules. We used Gephi software to construct all network graphs and used its built in analysis functions to calculate attributes such as node degree.

The following figure represents the source and destination IP addresses for the 5th network security zone at GMU. This graphic shows a few features that network engineers may be able to use when assessing their network traffic.[19] For example, Figure 15 shows an undirected graph of the source and destination IP addresses, the color of node represents the frequency of occurrence of the IP address in the records (where green means more occurrences and red less), and the width of the edges represents the frequency of the co-occurrence of the pair of IP addresses.

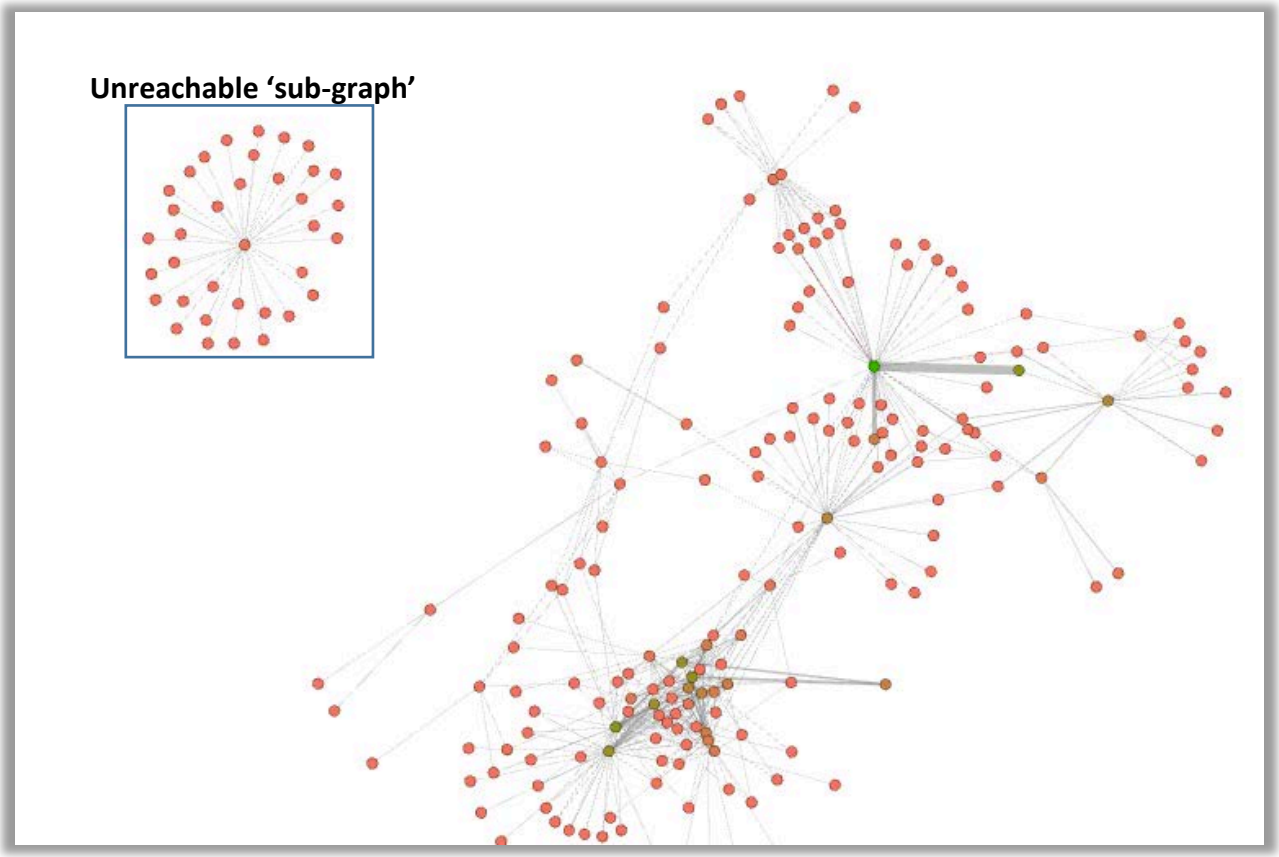

[19] As was mentioned previously, these results may not generalize to all network security zones. Due the sensitivity of GMU networks, we were only able to obtain roughly 10 minutes of network traffic logs for 5 out of approximately 250 network zones.

Figure 15: Zone 5 source and destination IP address graph

Looking at the structure of the graph, one notices that there are some nodes that are not reachable from the main graph: labeled as an 'unreachable sub-graph.' In this case, network engineers may consider whether distinct, unreachable graphs reflect the need for additional zones. Network managers may want to consider additional zones in those cases where the applications, traffic, or other factors are sufficiently distinct to warrant management as a zone. This graphic shows a method for discovering potential latent network zones.

In addition to developing insights based on the structures of the graphs, one may also be able to learn about network traffic looking at the frequency of the occurrence of pairs of addresses among all traffic. Figure 16 shows a node (highlighted with a bright green) with a high number of connections, which in this case are also highlighted in green. Additionally, one can see more clearly a case of the width of an edge representing the frequency of occurrences of pairs of IP addresses. The inset in Figure 16 shows a relatively wide connection between two nodes that dominate the frequency of the occurrences of pairs of IP addresses.

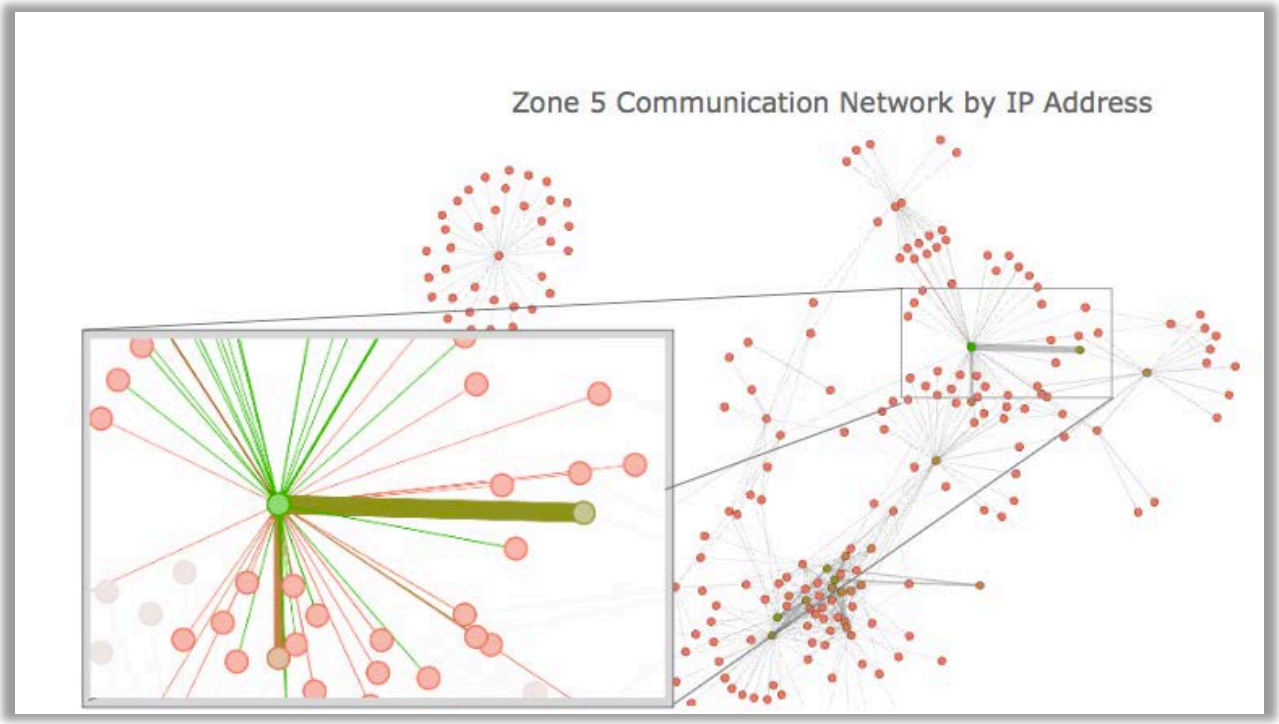

Figure 16: Zone 5 traffic graph with inset showing edges of a given node

As with Figure 15, the analysis results shown Figure 16 can help network traffic analysts assess where in the network frequent connections are likely to occur. Understanding the structure of and frequency of network communications are additional tools that could help GMU network analysts assess the impacts of new or changed firewall rules.

5.3 Anomaly Detection using Petri Nets

5.3.1 Petri Net basics

Petri nets are a method for graphically visualizing and executing discrete event systems such as in manufacturing, banking, software engineering, and command and control domains. Using a combination of places and transitions as shown in Figure 17, petri nets visualize the precedence relations and structural interactions. A circle node is called a place and a bar node is called is a transition. In this figure, the places are “Start” and “End,” and the trigger is labelled “Process.”

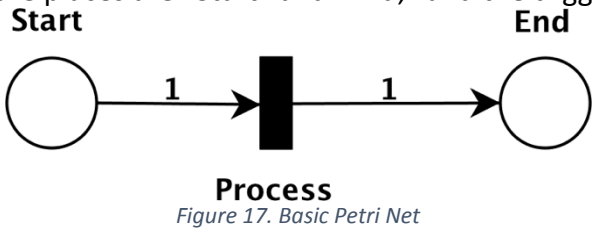

Figure 17. Basic Petri Net

Petri nets can model different states of a system by including tokens. When these tokens are allocated, the petri net is described as being marked. Using the structure of the basic petri net, Figure 18 shows the transformation from the basic petri net into a marked petri net. This is because a token now resides at the Start node.

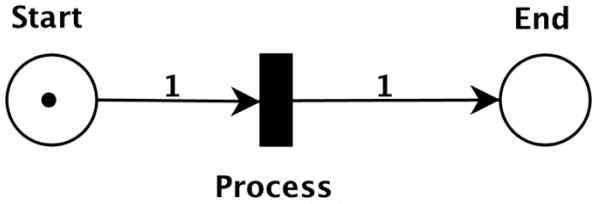

Figure 18. Marked Petri Net

Triggers are able to execute when all of it’s preceding places contain at least one token. When this precondition is satisfied, a trigger is able to execute – also called “fire” – by consuming at least one token from the preceding nodes, and then populating the output nodes. Firing the Process trigger in Figure 18 results in the marked petri net shown in Figure 19.

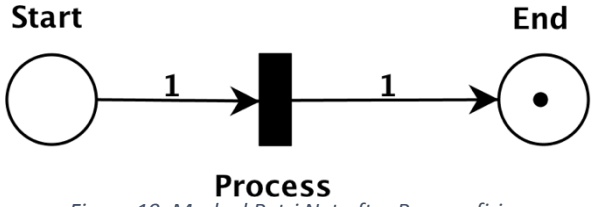

Figure 19. Marked Petri Net after Process firing

When multiple triggers can be fired such as the case in Figure 20, either one is allowed to be fired. Only one is able to be fired though, and then one token populating at the place End.

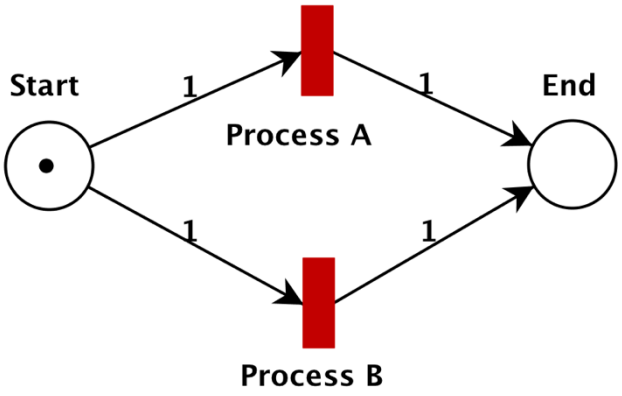

Figure 20. Petri net with two fire-able triggers

Another example of a scenario where either one of two triggers could be fired is shown in Figure 21. The difference in this scenario is that two different and distinct outcomes are possible. This means that the petri net execution will result in one of two possible distinct states – a successful end or an early end.

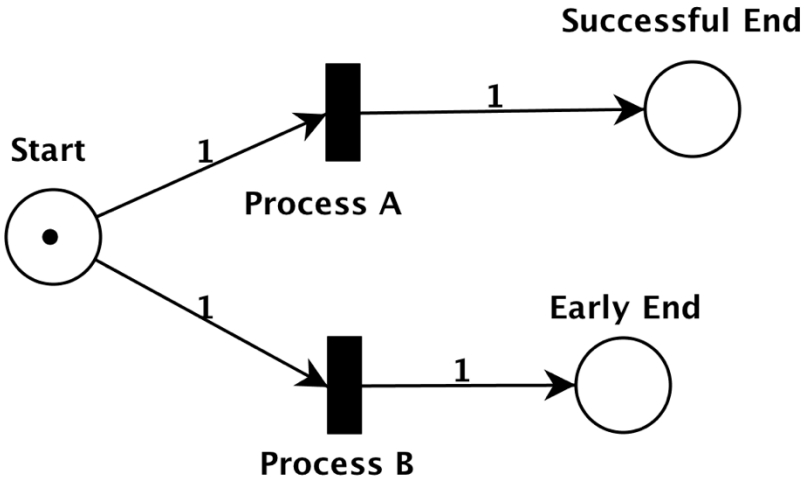

Figure 21. Petri net with different final outcomes

5.4 Why Petri Net analysis improves Firewall Rule Management

5.4.1 Anomaly Detection

Building on the Petri Net basics touched upon in 5.3.1, it is possible to design a discrete event system that would improve firewall rule management. This is accomplished by detecting anomalies in firewall rules, and having an end user address them appropriately. For example, consider the illustration in Figure 22. Network traffic is addressable by either rule 1 or rule 2, but they have distinctly different outcomes. This means there is a contradiction between rule 1 and rule 2, because they cannot both exist since they are mutually exclusive. Practically speaking, this scenario says that the network traffic is supposed to be allowed and also dropped at the same time, which is not acceptable. This is known as a contradiction anomaly because one outcome is contradictory to the other. In practice, this anomaly would be further identified

as being either a Shadowing or Generalization anomaly depending on the sequential ordering of the rules in the rule set.

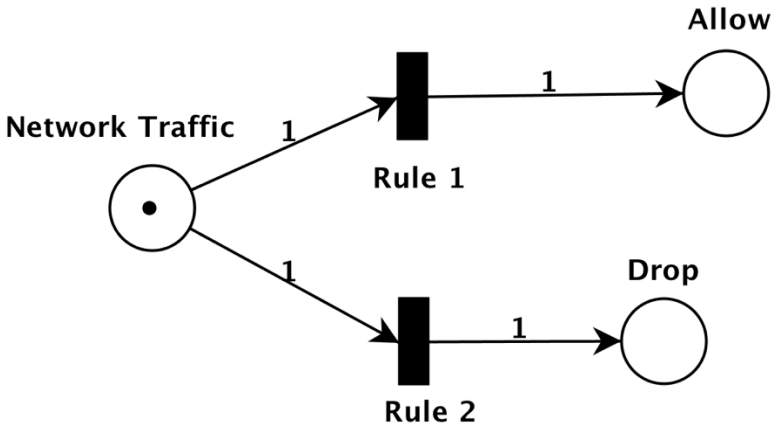

Figure 22. Example Firewall Rule Anomaly Detection

5.4.2 Anomaly Types

Below are descriptions of the five types of anomalies that are sought after when undergoing a review of the firewall rule set[20]:

Shadowing – a rule is shadowed by a preceding rule which matches all the packets that would match this shadowed rule. The actions are different.

Source Address	Destination Address	Action
192.168.1.1	*	Allow
192.168.1.1	192.168.1.2	Block

Correlation – two rules are correlated if they have different filtering actions, and the first rule matches some packets that match the second rule and the second rule matches some packets that match the first rule.

Source Address	Destination Address	Action
192.168.1.1, 192.168.1.2	*	Allow
192.168.1.2, 192.168.1.3	*	Block

Generalization – rule is a generalization of a preceding rule if they have different actions, and if the second rule can match all the packets as the first rule.

[20] Katic, T.; Pale, P., "Optimization of Firewall Rules," in Information Technology Interfaces, 2007. ITI 2007. 29th International Conference on , vol., no., pp.685-690, 25-28 June 2007

Source Address	Destination Address	Action
192.168.1.1	*	Allow
192.168.1.1	192.168.1.2	Block

Redundancy – rule is redundant if there is another rule that produces the same matching and action such that if the redundant rule is removed, the security policy will not be affected.

Source Address	Destination Address	Action
192.168.1.1	192.168.1.2	Allow
192.168.1.1	192.168.1.2	Allow

Irrelevance – filtering rule in a firewall is irrelevant if this rule does not match any traffic that may flow through this firewall. This exists when both the source address and the destination address fields of the rule do not match any domain reachable through this firewall.

Source Address	Destination Address	Action
172.14.0.1	10.0.0.1	Allow

5.5 Proposed application

The basic process of the petri net anomaly detection is illustrated in Figure 23. Given the network architecture and firewall rule set, a user would be able to execute the software to perform the anomaly detection. The analysis would take place, and the results would be processed. For this particular application, the targeted anomaly is for rules that are Shadowed or Generalized. These are anomalies where the actions between two rule are different, but one rule set is considered a subset of the other. To undergo the analysis, the user would start the application by importing the network architecture and firewall rule set. The software can be built to identify Irrelevance and Redundant anomalies. It is believed to identify Correlation anomalies also, but further research would be required. The results would be comprised of the following:

1. Firewall rule set containing edits with the reviewed anomalies.
2. Report of the removed anomalies

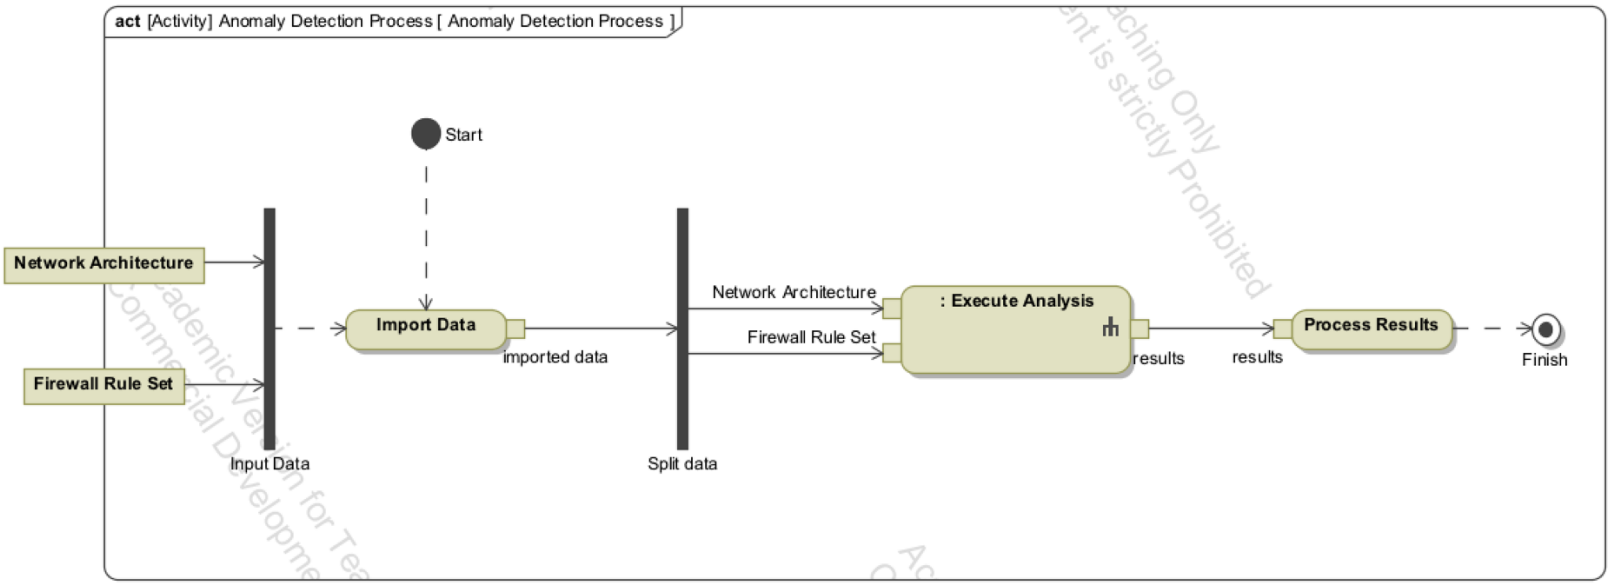

Figure 23. Activity Diagram for Petri Net Anomaly Detection

The analysis sub-process is broken down further in Figure 24. The network architecture is decomposed from a broad collection into hierarchical layers of mutually exclusive discrete nodes, similar to that of a tree structure. The hierarchical decomposition of the nodes is based on the attributes of the firewall rule. This means that if there are 6 attributes to a firewall rule, there would be 6 layers to the decomposition. The rules are then represented as triggers, which have each rule's dependencies mapped back to the mutually-exclusive nodes.

To determine shadowed or generalized rules, a reachability analysis is performed on each rule. Each rule is analyzed by placing tokens into the places that correspond to the network traffic. When a reachability analysis determines that two separate outcomes are available between two or more triggers, the rules involved in the anomalies are logged. After iterating through each rule, the anomalies log is then presented to the reviewer who deconflicts the anomalies. The resulting firewall rule set and log of the deconflictions are reported for processing and recording.

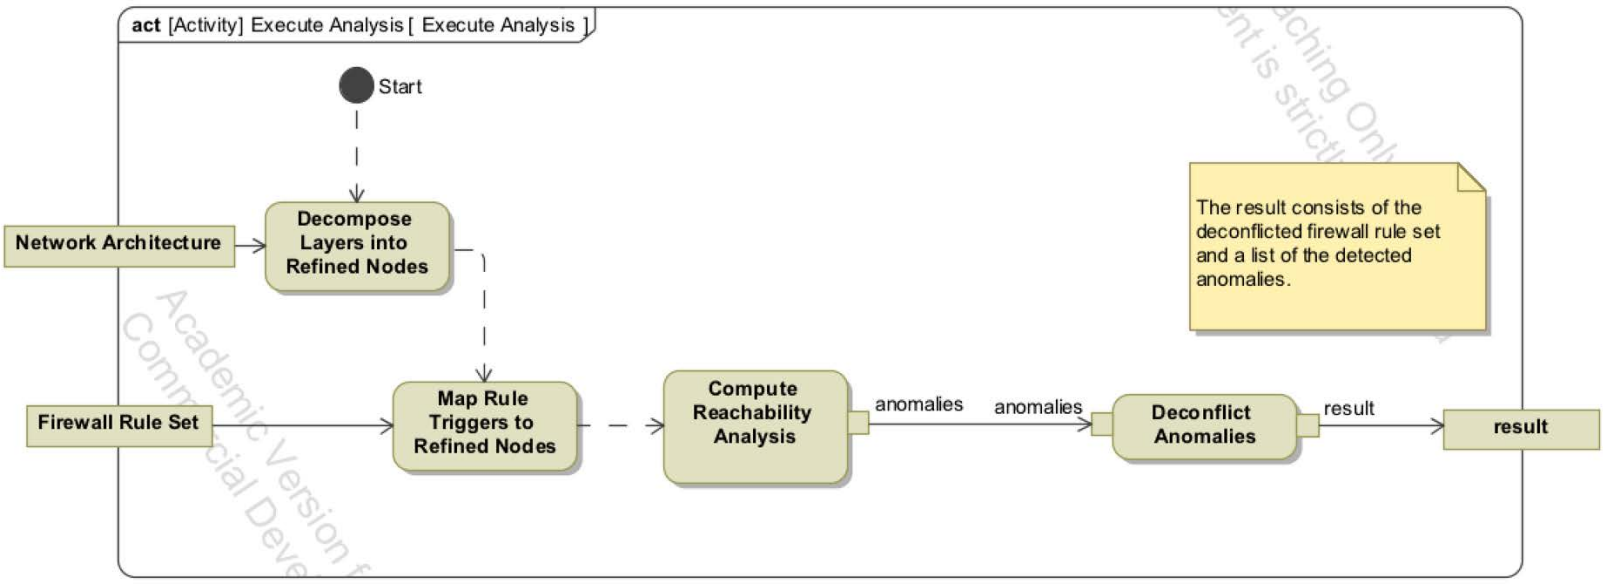

Figure 24. Activity Diagram for the Petri Net Analysis Sub-process

5.6 Manual Approach

For this project, a small-scale prototype was built to demonstrate the rule shadowing and generalization anomaly detection capability using dummy data. The goal of the prototype was to identify rules that were shadowed by preceding rules with different actions or generalized by following specific rules. An example list is provided below in table Table 1:

	Source	Destination	Action	Anomaly
Rule 1	192.168.1.1	192.168.1.2	Allow	
Rule 2	*	192.168.1.2	Drop	Generalization of Rule 1
Rule 3	192.168.1.2	*	Drop	Shadows Rule 4
Rule 4	192.168.1.2	192.168.1.1	Allow	

Table 1. Example of firewall rule anomalies

Another way to view this list is shows in Figure 25. Generalization is represented by the first two rules by the fact that rule 1 addresses a small subset of those source addresses in rule 2 and they have different filtering actions. Shadowing is represented by the 3rd and 4th rules by the fact that rule 3 addresses a broader set of destination addresses first, which makes rule 4 ineffective.

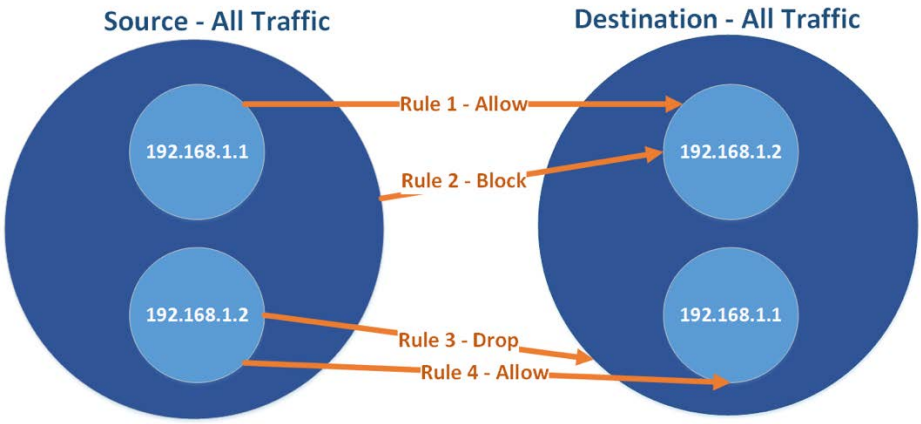

Figure 25. Example of firewall rule anomalies illustration

As described in 5.5, the firewall rule set as as well as the network architecture are required. The firewall rule set is described by Table 1, and the network architecture is described by Figure 25. The network architecture can be rendered as a petri net as shown in Figure 26. The head of the tree starts with the broadest classification, “All” traffic either from any source or any destination. The traffic can then be decomposed into the sub categories on the following layer.

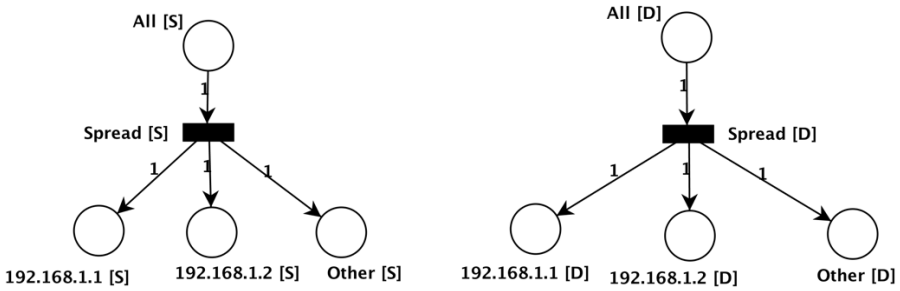

Figure 26. Petri Net of the network architecture

The next set of data to be imported is the firewall rule set. This can be modeled into our petri net by tracing the dependent places to the rule and the following action. Rule 1 was modeled this way as shown in Figure X.

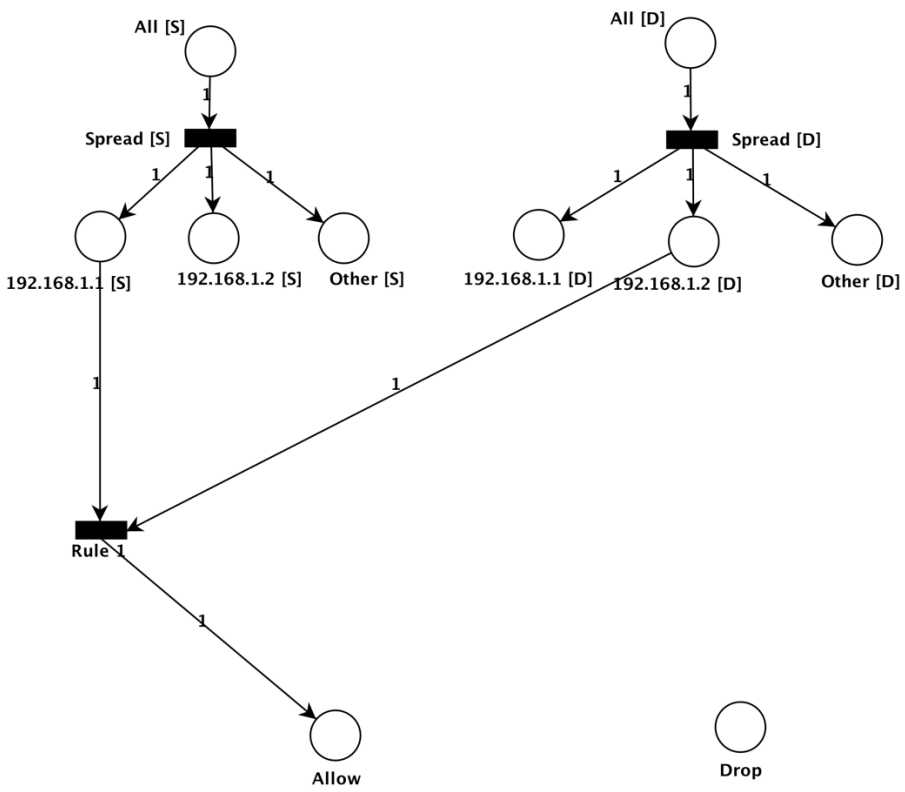

Figure 27. Adding Rule 1 to the Petri net

Mapping the remaining three rules results in the petri net below in Figure 28.

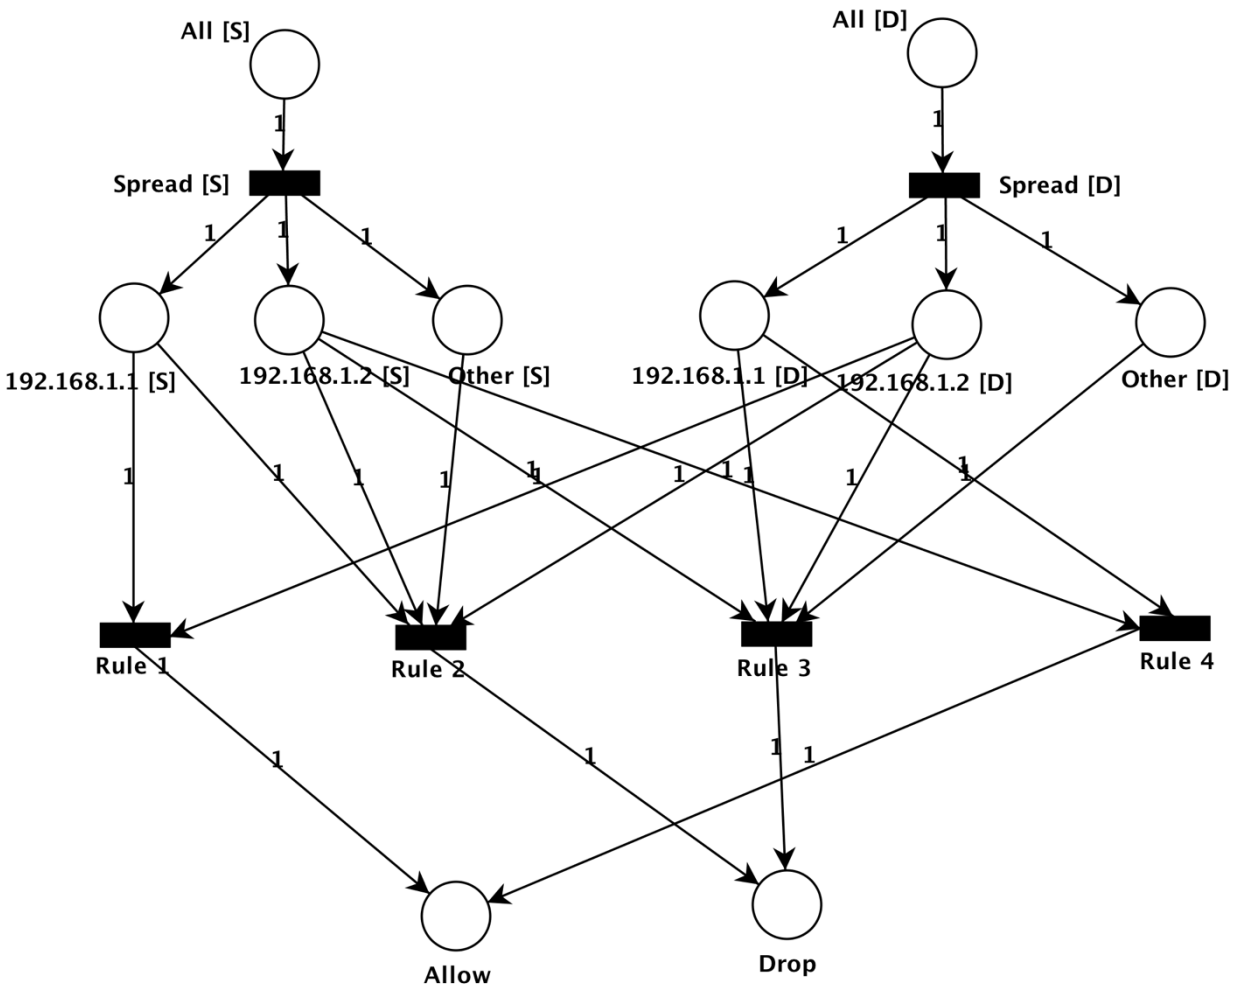

Figure 28. Petri Net of Network Architecture and Firewall Rule Set

The next step is to iterate through the type of traffic from each rule in the rule set, and see if anomalies show up. In this application, that means placing a token at 192.168.1.1 [S] and 192.168.1.2 [D]. The reachability analysis is completed and a red circle representing a final state is identified. Since only one final state is found, there is not a discernable anomaly between Rule 1 and the other rules. This makes sense because the type of traffic being addressed by Rule 1 is very specific, and other rules are not more granular than Rule 1.

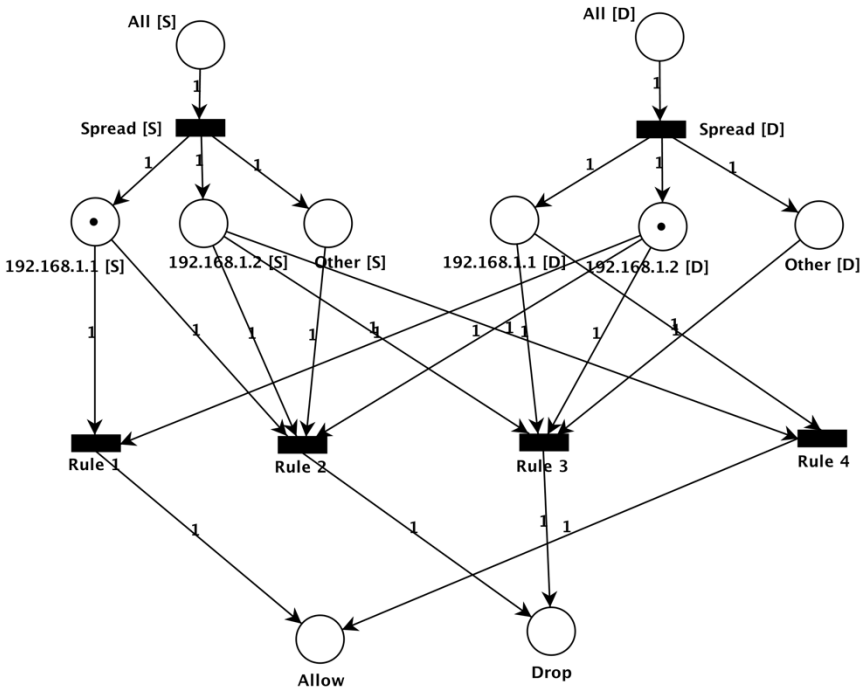

Figure 29. Analysis of Rule 1

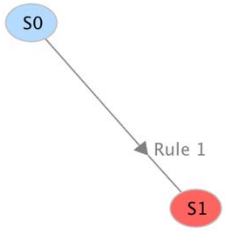

Figure 30. Reachability Graph of Rule 1

By analyzing Rule 2, a generalization anomaly is expected to be identified. The petri net is marked with the traffic that Rule 2 addresses. This can be shown by the petri net below. A token is placed at the All [S] place and at the place representing a destination address of 192.168.1.2.

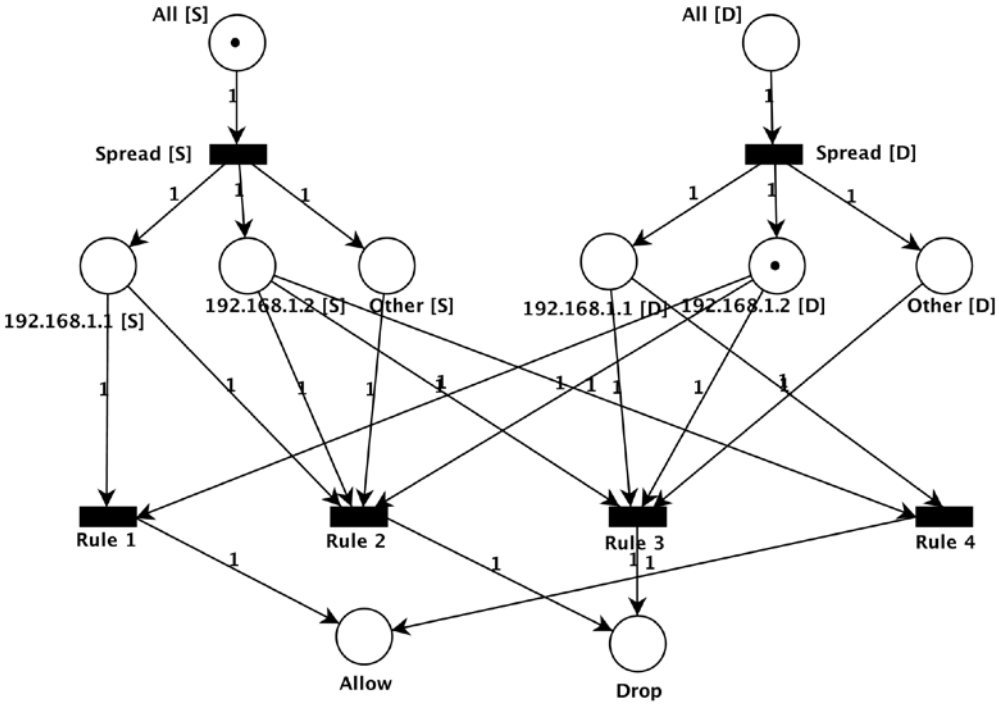

Figure 31. Analysis of Rule 2

The reachability analysis identified the anomaly as expected. Effectively, Figure 32 illustrates that two separate outcomes are possible by filtering with Rule 1 or Rule 2. This identification should be logged and presented to the firewall rule evaluator to determine if the anomaly is acceptable or not. Typically, generalization anomalies are acceptable because they act as more granular filters in a collection of addresses that might normally follow a particular action. Shadowing anomalies can be removed though, because they block the possibility of the preceding rule from triggering. In the reachability graph of Rule 2, it can be seen that Rule 2 and Rule 1 are the final triggers to the end states. Therefore, it can be classified as a Generalization because Rule 2 is in subsequent order to Rule 1.

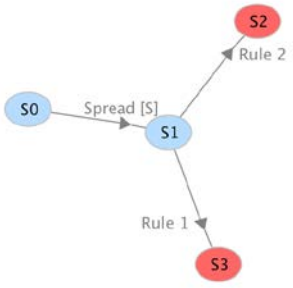

Figure 32. Reachability Analysis of Rule 2

By analyzing Rule 3, a shadowing anomaly is expected to be identified. The petri net is marked with the traffic that Rule 3 addresses. This can be shown by the petri net below. A token is placed at the 192.168.1.2 [S] place and at the place representing all destination addresses.

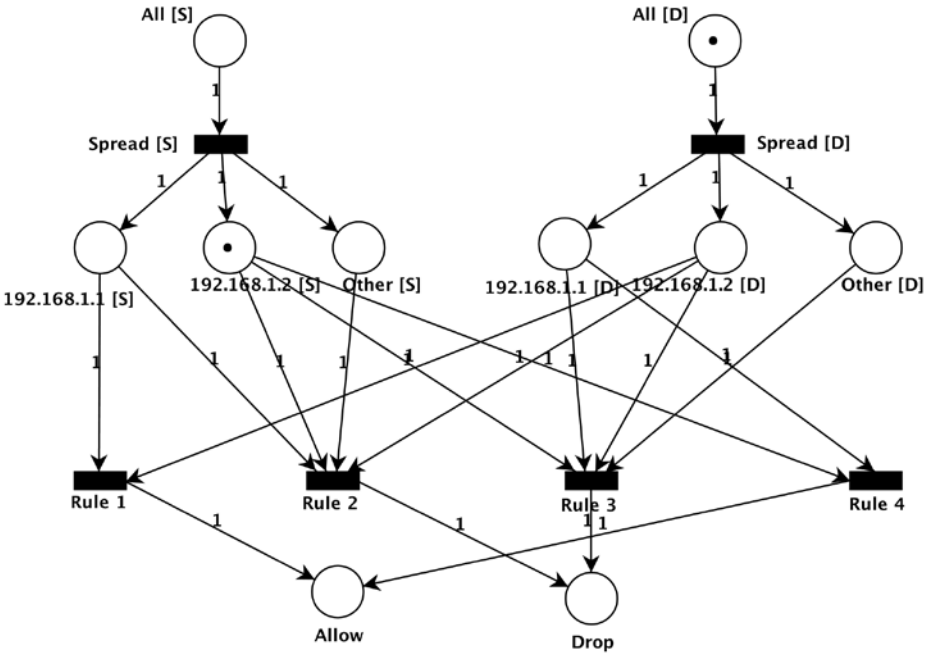

Figure 33. Analysis of Rule 3

The reachability analysis identified the shadowing anomaly as expected. Effectively, Figure 34 illustrates that two separate outcomes are possible by filtering with Rule 3 or Rule 4. This identification should be logged and presented to the firewall rule evaluator to determine if the anomaly is acceptable or not. Typically, shadowing anomalies are not acceptable because they block the preceding rule from triggering, which renders the subsequent rule ineffective. In the reachability graph of Rule 2, it can be seen that Rule 3 and Rule 4 are the final triggers to the end states. Therefore, it can be classified as a Shadow anomaly because Rule 3 is in preceding order to Rule 4.

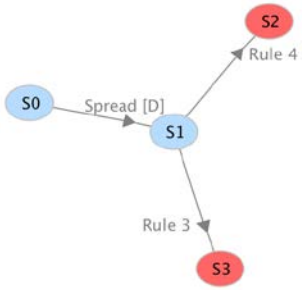

Figure 34. Reachability Analysis of Rule 3

During the development of the POC, it was determined that Irrelevance and Redundant anomalies can also be detected using the Petri Net approach. Irrelevance anomalies can be detected given that Other [S] and Other [D] map to a separate place named 'Irrelevant'. If the analysis shows a rule as having a state space mapping to only the 'Irrelevant' place, then an

Irrelevance anomaly has occurred. Additionally, further review can show that Redundant rules can be detected. If two separate rule triggers can fire and result with the same outcome and no tokens remain, then a Redundant Anomaly is present.

5.7 Scalability

In order to scale this application, the firewall rule set and network architecture need to be in a structured, machine-readable format. An approach would be to use a software library named Snakes that allows petri nets to be constructed using the Python programming language. This allows the program to perform the analysis without presentation to the user, and also execute analysis based on the detected anomalies to determine the type. This would conclude in a review by the end-user who would take the appropriate next actions.
Because of Python's flexibility, it could theoretically be incorporated into a behind-the-scenes process for change control. This automated analysis would reduce the burden of manual processes for impact testing and provide a formal analysis of the firewall rule set.

5.8 Review of tools used to create and analyze Petri Nets

The petri nets above were built using the Java software application named 'Pipe'. It is open-source software available for download. The software allows the manual creation, simulation, and analysis of petri nets.

6 Further Research

Given that GMU NET is at early stages of bringing their firewall rule set under configuration management and establishing auditing procedures, there are plenty of opportunities to extent our initial research. Systems engineers could be consulted to facilitate the development of requirements, policies, and procedures related to not only firewall rule set configuration management but also rule set auditing. This work can extend to process representation, requirements tracing, roles and responsibilities definition, among the many other aspects of system engineering.

First, although our examples show that Petri nets can be used to represent firewall rule anomalies, additional research is required before they can be applied to a real-world firewall rule set. First, software would be required to read a firewall rule set and translate those rules into a Petri net representation using standard Petri net representation schemes such as those found in Pipe. For a rule set of the approximate size of GMU's, this is a non-trivial task, that will likely require substantial software engineering experience.

Second, evaluation of a large-scale firewall rule set using Petri nets is likely a computation-intensive activity. However, because packets can be evaluated independently, the task of evaluating packets can be parallelized. We believe that joint research among system engineers

and computer scientists with experience in parallel computing could make a contribution to research in this area.

Next, given the central role that GMU NET's ticketing system plays in the management of its service work in general, and, in particular, work related to firewall rules, there is substantial work that could be done to improve not only its structure, but also operations based on its current structure. We recommend research in the area of discrete event modeling, simulation, and process optimization to explore methods for helping GMU NET use its engineering resources more efficiently to provide service.

Finally, additional statistical and network analysis using larger, unchanged firewall traffic data sets is likely to provide researchers the opportunity to develop insights into traffic patterns, fire rule effects, and other general trends in GMU network traffic. However, given the likely size of such data sets, this will likely require researchers with experience in constructing and using 'big data' architectures, such as those based on map-reduce computation or the Hadoop file system.

7 Appendix A: Firewall Rule-set Audit Checklist[21]

1. Review Change Process

1.1. Is the requestor documented?
1.2. Is the requestor authorized to make firewall change requests?
1.3. Is the business reason for the change documented?
1.5. Are there proper review and approval signatures?
1.6. Were the approvals recorded before the change was implemented?
1.7. Are the approvers all authorized to approve firewall changes?
1.8. Were the approvals recorded before the change was implemented?
1.9. Are the approvers all authorized to approve firewall changes?
1.10. You will need to ask for a list of authorized individuals.
1.11. Are the changes well-documented in the change ticket?
1.12. Is there documentation of risk analysis for each change?
1.13. Is there documentation of the change window and/or install date for each change?
1.14. Is there an expiration date for the change?

2. Review the Rule Base

2.1. Policy maintenance

2.1.1. How many rules does the policy have?
2.1.2. How many did it have at last audit? Last year?
2.1.3. Are there any uncommented rules?
2.1.4. Are there any redundant rules that should be removed?
2.1.5. Are there any rules that are no longer used?
2.1.6. Are there any services in the rules that are no longer used?
2.1.7. Are there any groups or networks in the rules that are no longer used?
2.1.8. Are there any firewall rules with 'ANY' in the 3 fields (source, destination, service/protocol) and a permissive action?
2.1.9. Are there any rules with 'ANY' in two fields and a permissive action?
2.1.10. Are there any rules with 'ANY' in one field and a permissive action?
2.1.11. Are there any overly permissive rules – rules with more than 1,000 IP addresses allowed in the source or destination?

2.2. Risk and Compliance

2.2.1. Are there any rules that violate our corporate security policy?

[21] http://www.sciencedirect.com.mutex.gmu.edu/science/article/pii/S1353485811700174

2.2.2. Are there any rules that allow risky services inbound from the Internet?

2.2.3. Are there any rules that allow risky services outbound to the Internet?

2.2.4. Are there any rules that allow direct traffic from the Internet to the internal network (not the DMZ)?

2.2.5. Are there any rules that allow traffic from the Internet to sensitive servers, networks, devices, or databases?

8 Appendix B: Links to Compliance Matrices for Higher Education

1. https://net.educause.edu/ir/library/pdf/CSD5876.pdf
2. http://www.higheredcompliance.org/matrix/