Modeling Flow Information and Other Control System Behavior to Detect Anomalies

Brian Moran, Rick Belisle
IBM Internet Security Systems
6303 Barfield Road
Atlanta, GA 30328
bmoran@us.ibm.com, rbelisle@us.ibm.com

Abstract: The stereotypical behavior of most control systems lends itself to modeling the system and detecting anomalies that could be cyber attacks or other dangerous network activity. In this paper we will discuss and demonstrate how the flow information available from routers and other network devices can be used in an anomaly detection system (ADS) to detect attacks from the communication flows from a source / destination / service perspective. We will also use this flow information to identify anomalies in the volume of communication on the network as a whole or between any two hosts. Finally, the paper will discuss the potential to model control system protocol usage to detect anomalies caused by cyber attacks.

Keywords: Netflow, network behavior analysis, network anomaly detection, SCADA

1 Introduction

Most SCADA and Process Control Systems in use today were developed years ago, long before public and private networks or desktop computing were a common part of business operations. As a result, the need for network security measures within these systems was not anticipated. At the time, good security for SCADA systems meant limiting and securing the physical access to the network and the consoles that controlled the systems. Planners rationalized that if they were suitably isolated from any physical entryways, and if access was limited to authorized personnel only, the systems were fully secure and unlikely to be compromised.

The increasingly networked and linked infrastructure of modern SCADA systems has rendered those early security plans obsolete. As companies have added new applications, remote access points and links to other control systems, they have introduced serious network risks and vulnerabilities that cannot be addressed by their physical control policies. Often, these risks are underestimated due to the complexity of the network architecture, the lack of formal network security guidelines and assumptions about the privacy of the network. Organizations are now realizing the security of these systems means more than physically separating the system and the components they control and monitor.

By analyzing network traffic flow and modeling behavior of the control system, commercially available security systems can offer the much needed protection for

control networks by identifying unusual traffic patterns. This paper defines the existing capabilities of network behavior analysis to identify network anomalies and maps this functionality to specific needs of control networks and SCADA systems.

2 Evolving Control Networks & Legacy Security

While legacy control networks face known IT threats, a bigger risk comes in the evolution of control networks. Many enterprises are now investing – or plan to invest – to upgrade their existing proprietary control networks with standard IP networks and operating systems. This investment paves the way for deeper integration with enterprise networks as well as reducing the costs associated with maintaining proprietary legacy networks.

2.1 Standardization

Network upgrades based on accepted IP and operating standards greatly reduces the heavy cost of maintaining legacy control networks. By adopting TCP/IP standards for their next generation control networks, enterprises can deploy a proven, reliable network infrastructure just like a typical enterprise network. Standardized operating systems eliminate the cost and reliance upon specialized knowledge of non-standard, proprietary legacy systems built 20-30 year ago.

However, the move toward these standard network infrastructure and operating systems also opens the door for IT threats that control system operators are not accustomed to seeing in their proprietary legacy networks. With standard TCP/IP networks and operating systems, control networks now face global risks that are common to every enterprise network. Sophisticated hackers are always searching for new vulnerabilities to exploit in these standard platforms. And "script kiddies" or moderately technical hackers can use widely available hacking tools to exploit known vulnerabilities that enterprises have not patched.

Network behavior analysis and anomaly detection have a role to play in maintaining security over TCP/IP networks by identifying any unusual traffic patterns that might indicate an ongoing attack or result from a successful attack against the control network.

2.2 Integration with Enterprise Networks

This evolution of control networks opens the door for deeper integration with the enterprise network. Business benefits of this integration provide compelling motivation to make the investment necessary to upgrade legacy control networks. However, integration introduces new dynamics to a control network, which have typically relied upon strict network segregation as a means of physical security.

Legacy control systems remained secure as long as they were closed off from other networks. However, control network operators must rethink their legacy security practices as these physical barriers erode. Network behavior analysis can be applied to monitor all integration points and identify threats that enter the control network through these points of integration.

2.3 Legacy Security

Because physical access to the network was the primary means of security for legacy control networks, the designers of these networks saw little need for the security that enterprise networks now employ. For this reason, many control systems were designed without access controls, encryption, authentication, user access rights and other security measures that are now required of any network that utilizes network and operating system standards and integrates with other networks. Unfortunately, many upgraded control networks still lack these basic layers of security because they still rely on proprietary protocols that lack this basic security.

Network behavior analysis and anomaly detection can be applied to compensate for the lack of basic security by applying advanced analysis of network traffic to identify network threats.

3 Security Challenges for Control Networks

While legacy control networks relied upon limited physical access to the network, security for next generation control networks demands the best practices of layered security for these new standardized networks. Enterprise networks have proven the value of firewalls, intrusion prevention, anti-virus and other standard security practices. However, SCADA systems demand 100 percent uptime and have no tolerance for latency introduced by in-line security systems. Certainly, any security measure that fails in a "closed" position to block traffic threatens network resiliency and performance. For this reason, most operators of control networks are reluctant to deploy these basic security measures.

Control networks demand security that identifies known and unknown threats without effecting network performance or risk shutting down the network or its critical applications.

3.1 Default Operating System Configurations

The critical applications that run across control networks often complicate security plans. In many cases, the vendors of these critical applications demand default configurations of the operating systems on which they run. Most of these applications run on the Windows platform, and thus control system operators cannot patch these Windows servers from the literally thousands of known vulnerabilities. As a result, any 10-year-old virus can bring down the control network once it reaches the systems. And any moderately sophisticated hacker can use widely available hacking tools to exploit security holes in the operating system if they can simply gain access to the network.

Network behavior analysis and anomaly detection can be applied to identify unusual behavior from profiled systems that could indicate an attack against an operating system's known and unknown vulnerabilities.

3.2 Critical Network Traffic

Despite the investment to upgrade control networks, little has changed in their actual operations – with the exception of the critical integration points. Control networks still employ a network design where a central command center conducts routine polling cycles to remote terminal units, which respond with the critical information needed at the control center. Of course the points of integration with other networks present the biggest difference in the upgraded networks. With this basic network design, security of control networks should focus on the relatively limited number of critical network points that include:

- Traffic from field sites and remote terminal units to the command center
- Traffic to and from the enterprise network
- Traffic to and from third parties, such as vendor support through VPN's and dial-up connections

By simply focusing on this critical network traffic, commercially available network behavior analysis solutions can identify threats by identifying anomalies and behavioral patterns of known threats.

4 Anomaly-Based Threat Detection

So far, this paper has focused on trends with control networks and areas where network behavior analysis can provide value by identifying anomalous traffic. The remainder of this section of the paper will focus on the functionality of commercially available Network Based Anamoly (NBA) solutions. All screen shots featured are of the IBM Proventia Network Anomaly Detection System. However, the functionality mentioned in this paper is also available in other NBA solutions, such as Peakflow X by Arbor Networks, StealthWatch by Lancope and Mazu Profiler by Mazu Networks. All these systems operate by passively collecting flow data from routers and switches. Technical requirements of all NBA solutions – including Netflow – are addressed in Section 5.

Because network behavior analysis is not widely deployed in control networks, many of the screen shots and examples given in this section draw upon experience with enterprise networks. This paper attempts to apply the lessons learned from enterprise networks to control networks as much as possible.

4.1 Modeling of Relationships between Network Assets

In passively monitoring network traffic, network behavior analysis solutions build a model of relationship of all network assets and entities. In short, NBA solutions maintain a log of all client-server connections and the services between these network assets. For control networks that lack user access controls and authentication, these relationships provide a baseline of accepted use – a default list of authorized connections. With this model of network asset relationships, network behavior analysis can alert to anomalous conditions, such as new clients or servers identified on the network or unusual connections made between two network assets that have not previously connected with each other.

For example, a utility company has three different field sites for three different neighborhoods. Each of these field sites operates with a remote terminal unit (RTU) that reports data back to the control center. Connections between field sites or RTU's would be identified as unusual and the network behavior analysis solution would alert to this anomalous activity.

To demonstrate how network behavior analysis solutions build their model of relationships, the screen shot in Figure 1 shows the host relationships of all client and servers in an enterprise network.

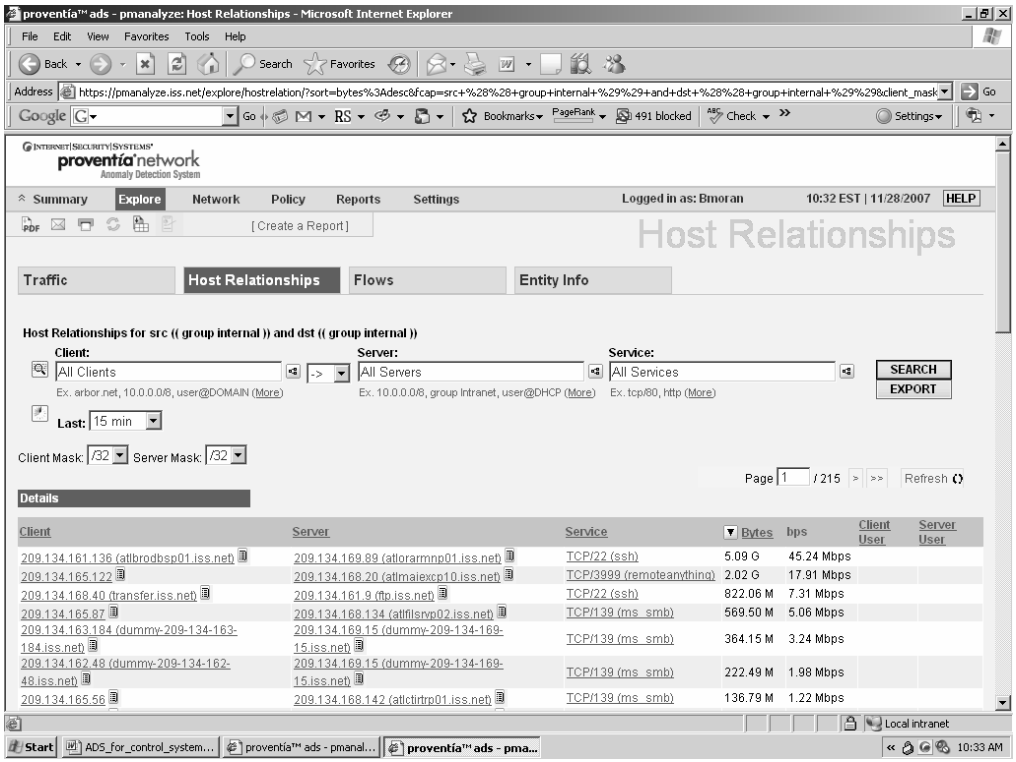

Figure 1 - Screen Shot of IBM Proventia Network Anomaly Detection System. List of all client-server relationships observed on this network with details into client and server IP addresses, DNS identifying date where available, service and amount traffic observed over the designated time.

4.2 Baseline of Traffic Patterns: Bandwidth & Protocol

While client-server relationships provide a foundation for analysis, most network anomalies are identified by establishing a baseline of traffic patterns of bandwidth and protocols. Network behavioral analysis solutions build a baseline model by passively monitoring network traffic and profiling each client and server based on its client-server relationships, bandwidth consumed and protocols or services used. Figures 2 to 4 provide screen shots of an NBA solution's profile of a network entity.

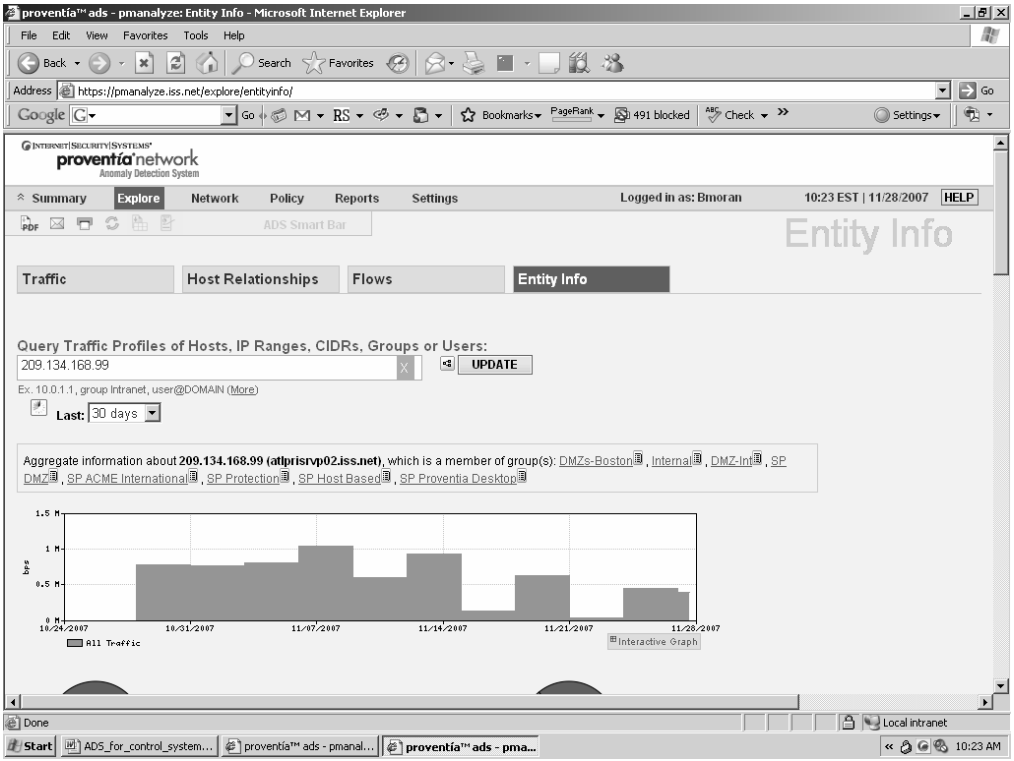

Figure 2 – Screen Shot of IBM Proventia Network Anomaly Detection System. Here the system focuses on an entity profile with a simple bar chart that shows all network activity over the defined period of time (30 days).

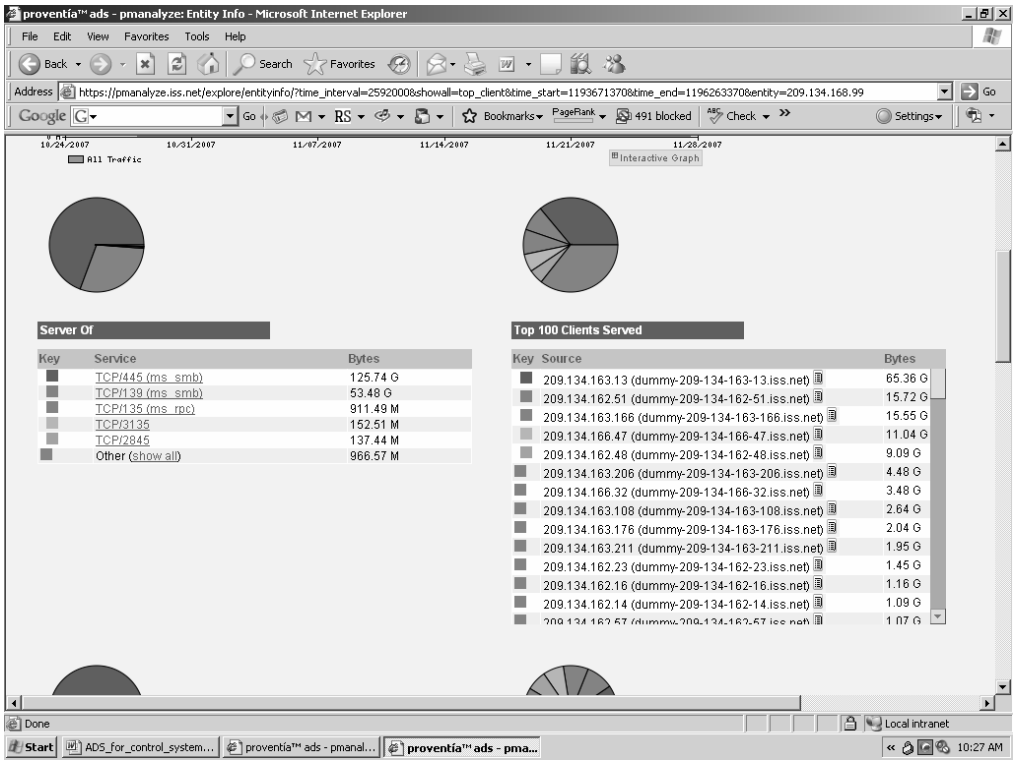

Figure 3 – Screen Shot of IBM Proventia Network Anomaly Detection System. Under the bar chart seen in Figure 2, the entity profile includes a list of clients served by this entity and protocols or services for its client relationships.

With this baseline and profile for every network entity, network behavior analysis solutions can identify anomalous events, such as:

- Spikes in network activity
- Unexpected drops in network traffic
- Atypical direction of network traffic

For example, an operator is likely interested in knowing about sudden drops in traffic from remote terminal units to the control center. The Modbus protocol used by many control networks uses function codes for the control center to monitor and control the remote terminal units. If a hacker gained access to the network, he could generate a Modbus packet to insert a function code for the remote terminal unit to delay or turn off its reporting. In this scenario, the control center could then miss critical information about pressure building in a valve that requires immediate attention.

While network behavior analysis could identify the sudden drop in activity from the remote terminal units to the control center, this scenario could have also been identified by the hacker's activity on the network. Network behavior analysis would have identified a new asset on the network that generated the function code.

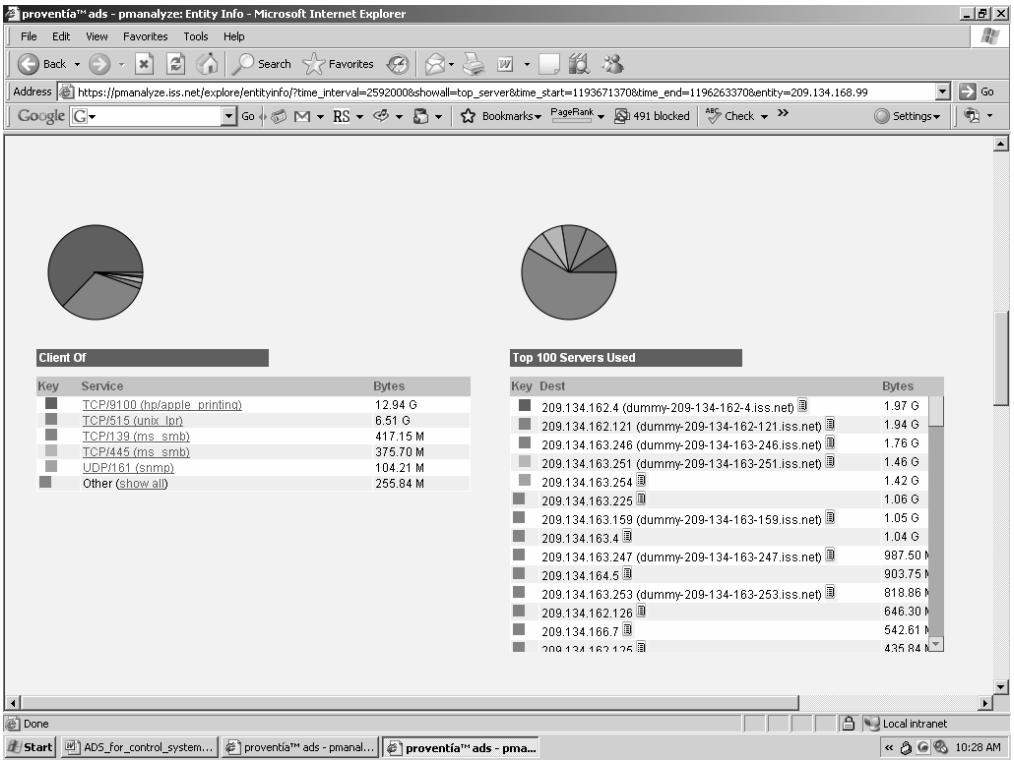

Figure 4 – Screen Shot of IBM Proventia Network Anomaly Detection System. The entity profile also includes a list of servers this entity has been a client of and the services used in its server relationships.

4.3 Known Threatening Behaviors

Network behavior analysis can also identify threats to a control system network with analysis that matches network traffic to the traffic patterns of known threats. This analysis does not necessarily rely upon a specific control system network's traffic baseline, but instead draws upon common anomalies to all networks. Network worms are the best example of matching network behaviors to known threats. Even if a worm has not been previously identified, network behavior analysis can identify the propagating behavior of a worm as it spreads from an infected host to other network entities. Figures 5 and 6 below show the identification of a worm spreading on an enterprise network.

Other known threats than can be identified through their anomalous behavior on a control network include port scans, host scans, connections to known command and control botnets, connections to known phishing servers, or traffic routed through proxy servers.

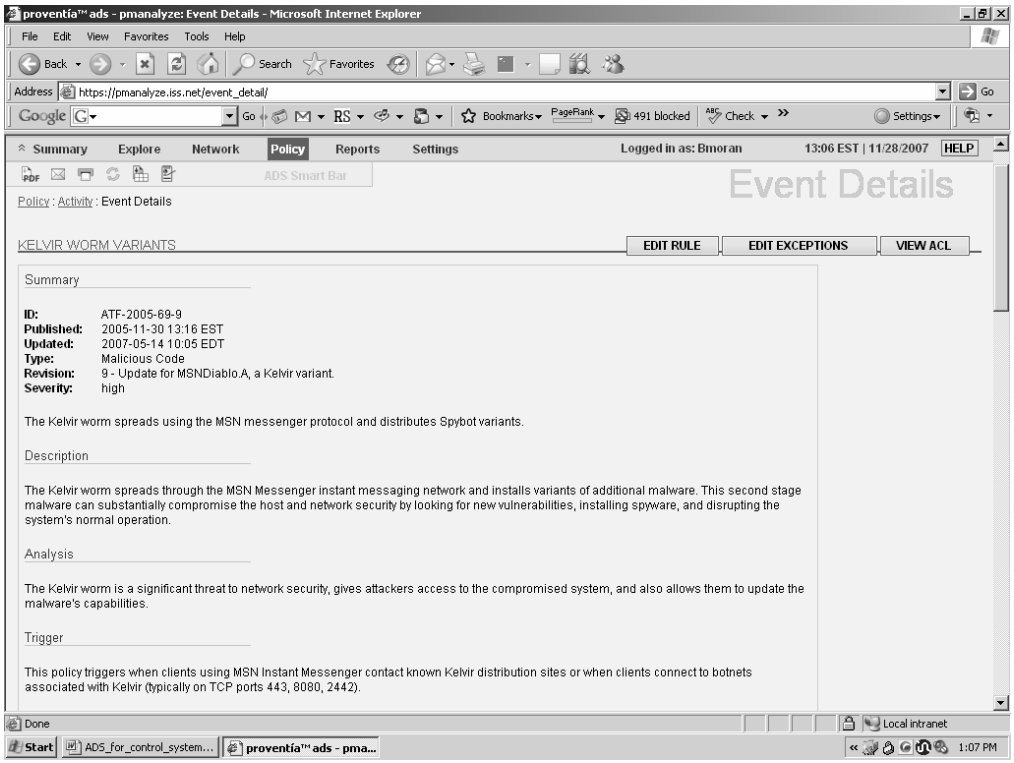

Figure 5 – Screen Shot of IBM Proventia Network Anomaly Detection System. This rule or policy identifies the propagating behavior of a worm that appears to be a variant of the known Kelvir Worm. This is an example of a known network threat that does not require a traffic baseline. This screen shot gives a description of the threat, a brief analysis of the risk it poses to the network and trigger of how this worm was identified.

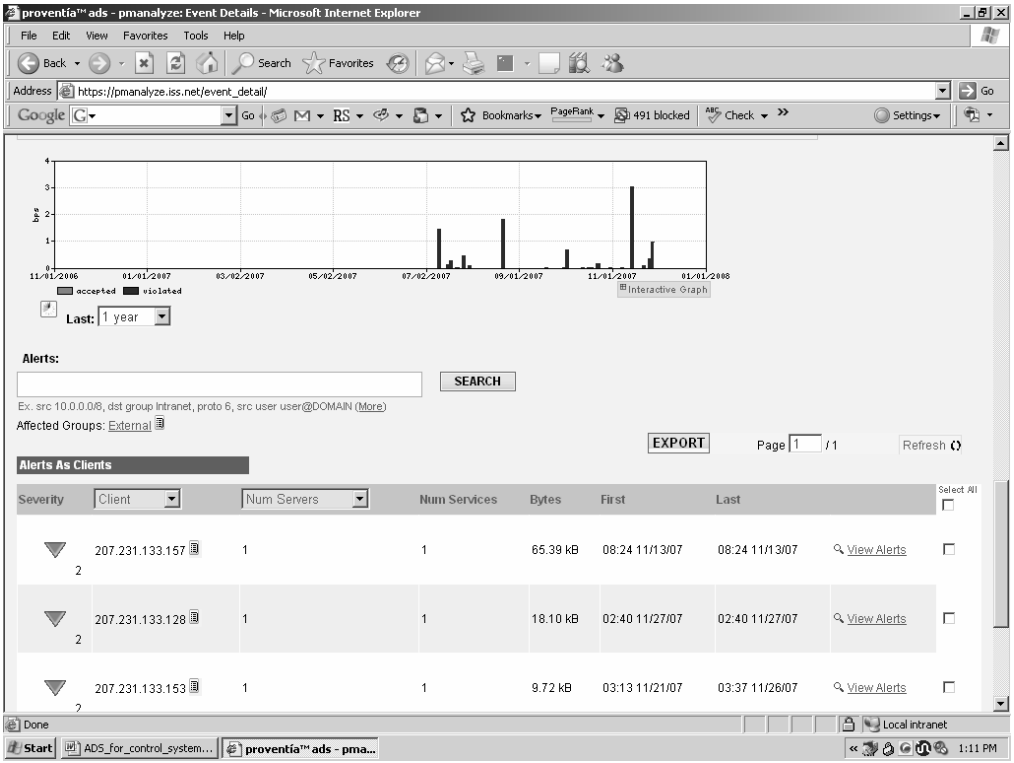

Figure 6 – Screen Shot of IBM Proventia Network Anomaly Detection System. This screen shows further detail of the Kelvir Worm variant as it appears on this enterprise network with traffic profiles and a list of clients who appear to be infected with this worm.

4.4 User-defined Norms and Acceptable Use

Network behavior analysis solutions can also identify anomalies and threats based on network traffic that violates what the user defines as normal and acceptable for the control network. By building policies for appropriate use, the user can clearly define the threats that are specific to the control network. In many ways, network anomaly detection provides an ideal solution for control system networks because these networks rely on such a simple architecture. A user can easily define the protocols and services of control network traffic and the authorized entities on the network. By default, anything else is anomalous.

Because all traffic from a remote terminal unit to the control center should be Modbus encapsulated over IP (TCP port 502), a good example of a user-defined policy would be to alert any activity from the remote terminal unit that is not over TCP port 502.

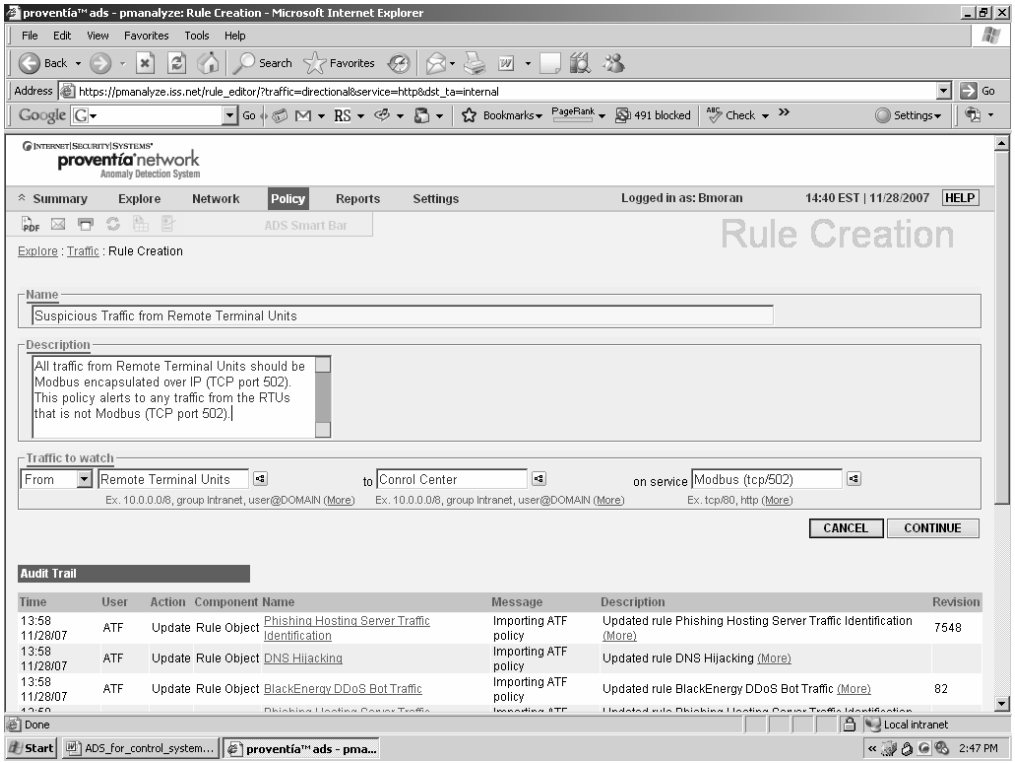

Figure 7 – Screen Shot of IBM Proventia Network Anomaly Detection System. This screen demonstrates the creation of a policy to alert for any traffic from remote terminal units that is not Modbus encapsulated over IP.

5 Technical Requirements

5.1 Flow Data

The anomaly detection capabilities profiled in this paper are based on commercially available network behavior analysis solutions that analyze flow data from routers and switches. As network traffic passes through a router's interface, it is possible to maintain traffic profile monitoring statistics based on the packets flowing through that router. These statistics are also referred to as "flow." Specifically, "flow" is defined as a unidirectional sequence of packets between two endpoints.[1] Every "flow" record is made up of several fields, some of which include:

- Source and destination IP address
- Source and destination port

[1] Detecting Worms and Abnormal Activity with NetFlow, Part 1, Yiming Gong (www.securityfocus.com/print/infocus/1796)

- Protocol
- Type of Service (ToS)
- Input and output interface number
- Next hop address
- Number of packets in the flow
- Total bytes in the flow
- TCP flags

Cisco's version of flow is called Netflow with version 5 the most widely deployed. Juniper Networks and Huawei Technology provide a similar feature for its routers called Jflow and NetStream, respectively.[2] Networks deployed before 2000 are not likely to support any form of flow. When upgrading control networks, flow generation should be a procurement requirement when designing the network infrastructure.

Flow generation introduces some overhead that can tax routers that operate at near capacity. However, most deployments will generate less than 10,000 active flows per second, which consumes 7.14 percent of the router's CPU utilization.[3] Because of this extra load on the routers, some enterprises have chosen to only sample flow data. However, sampling greatly reduces the ability to identify anomalies and enforce policies for appropriate use.

5.2 Deployment

Network behavior analysis solutions rely upon a deployment that provides visibility to all critical areas. As previously mentioned, this deployment should focus on the critical areas of integration with enterprise networks, traffic from field sites to the control center, and traffic to and from third parties, such as vendor support through VPN's and dial-up connections. This requires that the network behavior analysis solution to collect flow data from the routers and switches that carry the traffic for these three main areas. In short, visibility is limited to the collection of flow data. For example, a centralized deployment that only monitors flow data at the control center as traffic comes from remote sites would not be able to identify direct traffic between two remote sites if that traffic is not routed through the routers in the control center.

6 Summary of Benefits & Limitations of Anomaly-based Detection

Control networks can benefit greatly from anomaly-based detection of network behavior analysis by identifying network threats and compensating for a lack of security that is often not available on process control networks. However, NBA solutions are not a

[2] Wikipedia: "Netflow" http://en.wikipedia.org/wiki/Netflow
[3] Cisco: "Netflow Performance Analysis", page 19.

panacea for security. Other security measures, such as network intrusion prevention systems, should also be evaluated for control networks.

6.1 Protection that Does Not Threaten Performance

By passively collecting flow data from routers and switches, network behavior analysis solutions avoid the common concerns about other security measures, such as anti-virus and firewalls, that they could introduce network latency and threaten network uptime if they failed in a "closed" position that blocks all traffic. NBA solutions are "off-line", so there's no threat of network latency or blocking traffic if the solution failed.

6.2 Network Visibility

Network behavior analysis solutions are one of the best technologies available for delivering visibility to network assets, services and behaviors. This consolidated view of the network is often cited as one of the greatest benefits of NBA solutions. With this visibility, control network operators gain valuable insights into network performance and potential bottlenecks. This information is commonly used for capacity planning and ensuring network availability for critical applications. However, network visibility is most valued for delivering a higher level of overall control of the network by providing all relevant information needed to properly manage and secure the network.

6.3 Flow Limitations

While flow delivers much information about what's going on across a network based on traffic profiles, it does have some protection limitations. Inline IPS is able to decode a packet stream based on the protocol, detecting and blocking attacks at the packet level. Most IPS solutions are able to decode and investigate a packet stream for layer 7 attacks. Today, flow data is more general in that it focuses on layer 3 information, but future flow versions promise the ability to leverage layer 2 and packet header information as well.

7 Conclusion

As control networks and SCADA systems evolve to integrate with enterprise networks and adopt network and operating system standards, control system operators face security threats that are common to enterprise networks. Many control systems choose not to implement standard security practices because of their fears that security could impact network uptime and guaranteed network performance causes. For this reason, commercially available network behavior analysis solutions provide an attractive alternative for protecting control networks. With off-line analysis, NBA solutions identify network threats by recognizing network anomalies.

About the Authors – Rick Belisle is the Manager for the X-Force Professional Security Services Southeast Region. He has been with IBM Internet Security Systems for over seven years, and as Services Manager he is responsible for all professional service delivery in the region. Mr. Belisle has 15+ years of information technology experience, and has specialized in information security for over 10 years. He has conducted numerous information risk assessments and penetration tests for a variety of government and civilian clients. Most recently Mr. Belisle has been focusing his efforts on Process Control Network (PCN) security in support of large scale SCADA assessments.

Brian Moran is a Product Marketing Manager for IBM Internet Security Services.