

A guide to firewall intervention for threat response

Using Athena FirePAC

Abstract

SIEM solutions detect real attacks from the thousands of events that are happening in the network. When an unexpected and potentially dangerous event is recognized, engineers take a quick (temporary) action to block the security leak, diagnose what really went wrong, determine what else is at risk, and decide on a fix. At the same time, the engineer might have to figure out what changed in the network configurations to cause the incident. Depending on the complexity and number of devices required to be investigated, executing a timely response can be unrealistic. This paper discusses how Athena's firewall analytics solution can be used to find the rule changes related to the incident and to verify if the remedy implemented actually works without creating unintended side effects to the network.

Introduction

When a SIEM solution identifies a potentially dangerous event that requires firewall intervention, engineers might plug the hole using a quick temporary action. Still, they will need to ensure that they have correctly addressed the actual problem, as well as any additional exposures, in order to prevent the attack from happening again.

In the face of a potential attack, engineers perform the following diagnostics to identify what needs to be fixed:

1. Find the source from where the threat is originating, the destination it is impacting, and the services or applications that are affected.
2. Find the path the particular threat took and determine if there are alternate paths that are possible.
3. Find the firewalls and routers that exist along the path that did not block the attack.
4. Isolate the firewall rules that allowed the attack to go through.
5. Modify the rule base to block the threat without affecting the service availability of other business services. This requires making sure rules are added in the correct place to not only block the threat, but in a way that does not disrupt other business services.
6. Identify what changed in the configurations to cause the problem, and understand if something else is exposed because of the same changes.
7. Investigate what else is accessible from the same sources or to the destinations used in the attack. Plug those holes as well before the other assets are actually attacked.
8. Find out if there are any alternate paths through which the same threat can take and repeat the same process for those paths as well.

The challenges for engineers

While the SIEM may have identified one origination point for the attack, it may not always detect all other origination points and paths that the attack can take. The firewalls might not log all packet data, and even if they do, the rules might not have been identified in the logs. This requires the engineer to perform a more thorough analysis of the firewall rule bases to isolate the rules that are allowing the traffic through. In cases where firewalls are highly complex, with hundreds/thousands of rules and object groups, manually analyzing the firewalls leads to an increased risk of disrupting critical business services. Another challenge is that the ultimate destination address might be a private internal RFC address. In order to isolate the ACL rules that are affected requires using the public addresses prior to translation to look at the ACL rules that allowed them.

Counter-acting the challenges with an automated firewall analysis solution

Firewalls and routers protecting enterprise networks with an already complex web of inter-connections will inevitably grow more complex because of the need to add rules in order to provide network access and protect against attacks. Ideally, rules would be added to the firewall in an organized manner. Furthermore, rules would be organized and enhanced to suit specific business purposes. Unfortunately, that is not reality.

Firewall administrators change; as new people transition into the role, rules are added in an ad hoc manner and the configurations across the network eventually becomes a disordered, chaotic mess.

Manually understanding the complete effect of a rule that refers to object groups having multiple levels of membership hierarchy is not only painfully tedious, it is error prone. As the rule base increases, the number of possible combinations explode.

For example, we have observed rule bases consisting of a total of 875 rules with 125 Deny rules using almost 4000 address objects/groups and 800 service objects/groups has hundreds of thousands of combinations.

If there are many overlaps between the rules and if the rule base is sprinkled with many rules blocking dangerous services (which tends to be the case in open network environments where the desire for open policy has to be reconciled with a policy to protect certain critical assets), then it becomes virtually impossible to determine the impact of each rule manually.

Also, in most networked environments, firewalls from multiple vendors exist to provide security defense-in-depth. However, there is no unified interface for accessing and managing these firewalls across vendors; they are often managed from separate consoles. Getting access to the configuration or pushing changes might often involve logging into the device using SSH or telnet. Without a unified view of what exists in these firewalls, one cannot easily compare rules. Even though firewalls from different vendors serve a similar purpose, their design and architecture are different.

Cisco firewalls have rulesets that can be enforced on an entering or exiting interface of the traffic. Cisco firewalls also have a "NAT control" feature that serves as an additional access control function. Juniper NetScreen firewalls enable users to apply rulesets based on the origination zone and the destination zone, where each zone contains networks that are partitioned using interfaces and routing tables. It is rare to have firewall administrators who have an understanding of all firewall types and this will introduce inconsistencies in policies deployed to the firewalls.

Athena FirePAC Firewall Analytics solution

As mentioned above, firewall configurations can easily grow complex. Analyzing configurations for firewalls from multiple vendors makes this an extreme burden. What is needed is a technical assistant, if you will, that understands the science of firewalls. This assistant is the firewall analytics tool. It completely understands all components of the firewall configuration for its *meaning* and *complete effect* on firewall behavior and can provide the following help to the firewall administrator:

- An ability to search across firewalls for address and service objects and object groups using object name or object content showing the complete hierarchy of the object groups involved, and the rules that refer to them. This aids the engineer in identifying which firewalls actually use the source, destination and service elements in the threat identified by SIEM. This will also help the user in identifying the proper object groups that he needs to use for making the change and the effect of that change.
- Address and service based advanced ACL rule search across firewalls (by names or content) will aid the user in figuring out the rules and their dependencies that are allowing the threat identified by SIEM. Understanding the dependency between the rules is very important to modifying the right rule or adding a new rule in the right place; otherwise the change might not have addressed the threat or caused an adverse effect on other business services. With this analytic function, administrators can easily find out the proper rules and rule sequence in the rule set for making efficient changes.
- Queries against the firewalls using private RFC addresses identified in the threat without worrying

about the address translations happening in the network and multiple steps to complete a report identifying all the ACL, NAT, VPN and routing rules involved. With FirePAC, the firewall administrator can use one query to identify the firewalls and rules that are involved and automatically account for address translations.

- An analysis of the effect of a change before a change is pushed to the device will help in better understanding risks to service availability as well as the exposure to any security holes. This also will result in few configuration changes and less rule “bug” fixing.

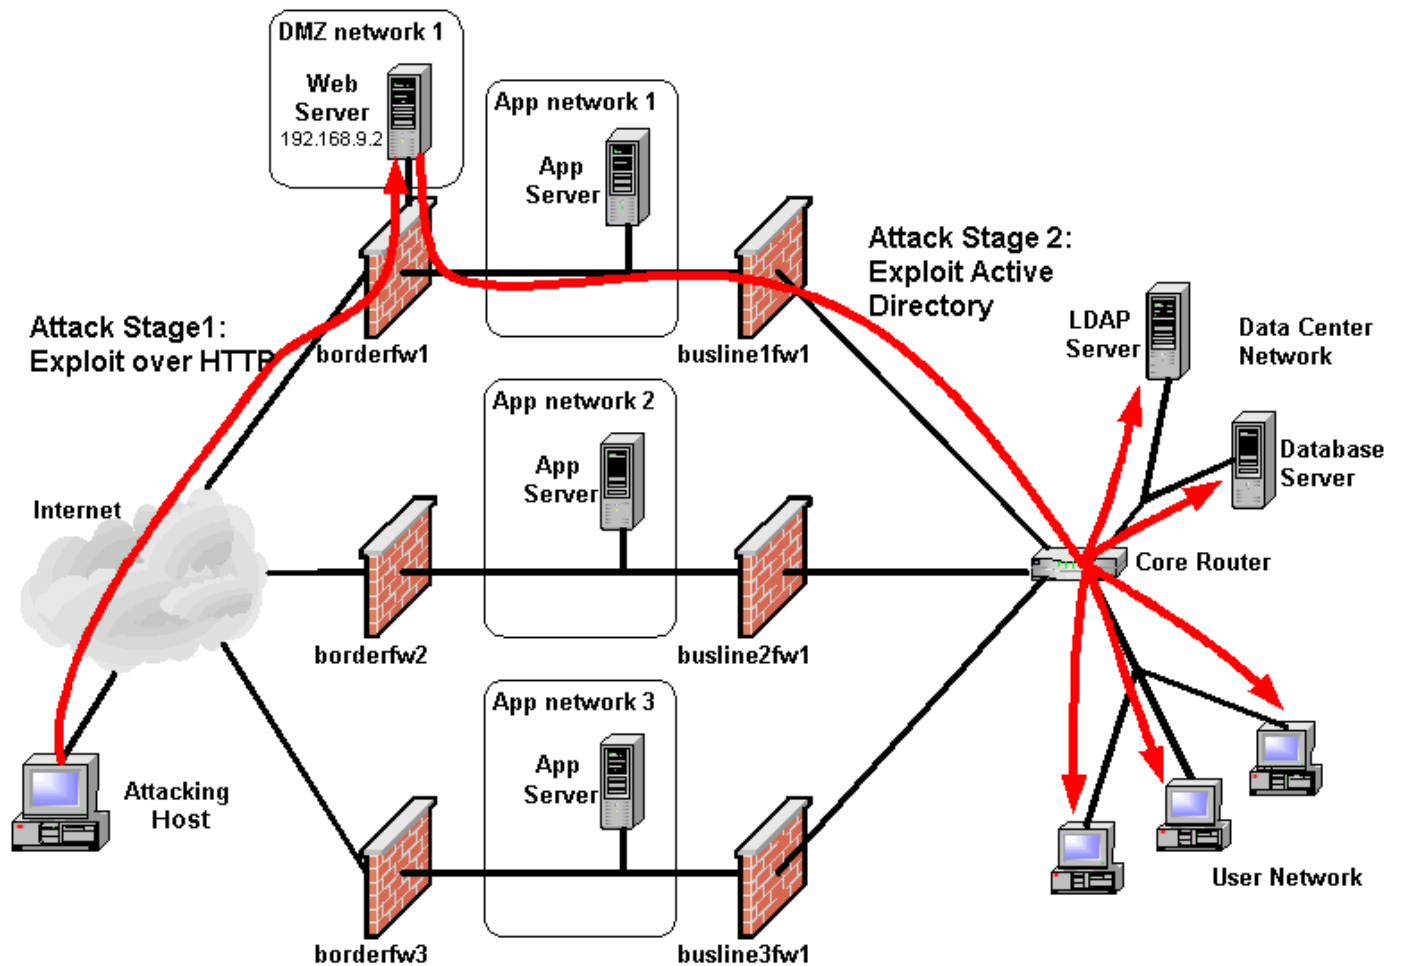
Step by step guide for isolating and validating changes to close the threat using FirePAC:

1. **Isolate the packet information from the threat signature:** Isolate the address of the source from which the threat was originating, the address of the destination being targeted, and the service ports that were present in the attack. The diagnostics help the user block the attack from the source(s) to the destinations for the ports isolated in this step.
2. **Import firewalls into FirePAC’s inventory:** Import the configurations of the firewalls that need to be analyzed into FirePAC’s firewall inventory. FirePAC can import firewall configurations from an NCM repository, configurations saved as files in the file system or by connecting to the devices using an SSH/Telnet channel.
3. **Run Object Query:** Using the packet information from step 1, run object query across the firewalls in the FirePAC inventory to quickly find objects that represent the source, destination addresses and service elements of the packet.
4. **Run ACL Rule Query:** Using the packet information from step 1, run ACL Rule query across the firewalls to quickly find ACL rules that allow or deny the packet.
5. **Run Policy Query:** If the destination represents an internal private RFC address, use the post-nat data flow query to find all firewalls in the inventory that allow the given packet. This will identify all ACL, NAT, VPN and Routing rules within each firewall that allows the given packet to go through the firewall.
6. **Run Rule Dependency Report:** For the firewalls that require firewall intervention, run the Firewall cleanup and optimization report to identify the rule order dependencies that exist between the rules in each rule set. This will help you in understanding the proper position if a new rule to block the packet is being added.
7. **Run Policy Diff Report on the new configuration:** Once the changes have been made, run the policy Difference report using the original and the modified configurations to make sure that only the packet(s) related to the threat are being blocked.

Example Scenario

This section explains how to use FirePAC to diagnose a SIEM event using an example scenario. The scenario involves a company which has a web presence with multiple web applications. Each web application has its own DMZ with a perimeter firewall separating it from the Internet and an application network with an internal firewall separating the application network from the internal and data center network. The perimeter firewalls allow HTTP from the Internet to a publicly accessible web server in the DMZ. The web server in turn requires access to resources on the application network and an LDAP server on the internal network. The internal firewalls are connected to a core router that provides access to a data center network containing the corporate LDAP server and the Data base server. The core router also connects to network segments containing user workstations. The hole in this architecture is the access to the LDAP server in the internal network from the web server.

Example Scenario: Access to the LDAP server in the internal network from the web server



SIEM detects a 2 stage attack on the internal networks as depicted below. In the first stage, vulnerability in the HTTP application is used to compromise the web server in the DMZ. The compromised host is then used to launch a stage 2 attack on multiple internal assets using a hole in the network policy that allows access on the Active Directory service to any internal host not just the LDAP servers. The access rules on the border firewall for the outside interface allow access to the DMZ server using a public IP 62.59.14.169. The access rules on the web dmz interface allow application access to the application servers on the application network and the LDAP access to all hosts in the internal network. This is a security hole in the configuration introduced by not restricting the access to LDAP servers only. The internal firewall makes the same mistake. The firewalls used to explain the scenario are Cisco.

borderfw1:

```
access-list acl_outside permit tcp any host 62.59.14.169 eq 80
static (webdmz,outside) 62.59.14.169 192.168.9.2 netmask 255.255.255.255 0 0
access-list acl_webdmz remark - TODO needs to restrict AD access to AD servers only
access-list acl_webdmz permit tcp host 192.168.9.2 any object-group AD_svcs
access-list acl_webdmz permit tcp host 192.168.9.2 host 172.16.10.101 eq https
access-list acl_webdmz permit tcp host 192.168.9.2 host 172.16.10.100 eq 7778
static (webdmz,inside) 192.168.9.2 192.168.9.2 netmask 255.255.255.255 0 0
```

Businessline1fw1:

```
access-list acl_outside remark - TODO needs to restrict AD access to AD servers only
access-list acl_outside permit tcp host 192.168.9.2 any object-group AD_svcs
static (outside,inside) 192.168.9.2 192.168.9.2 netmask 255.255.255.255 0 0
```

Threat Response using FirePAC

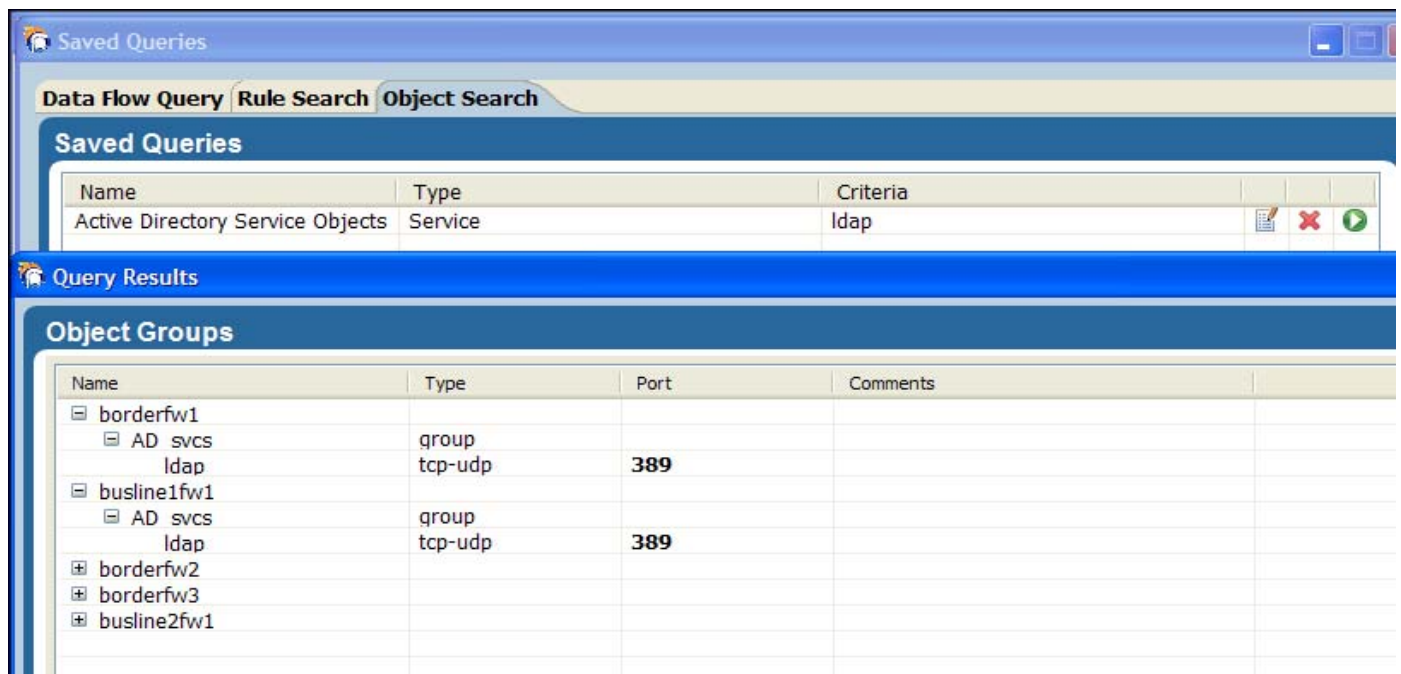
When the SIEM solution identifies the attack, it flags the external IP address where the attack originated, the HTTP service exploited, and the IP address 192.168.9.2 of the DMZ web server compromised by the attack in stage one. Then it identifies the Active Directory service and the internal host IP addresses targeted by stage two of the attack. The response to this incident should be to patch the vulnerability in the web server HTTP application and close the security policy holes in borderfw1 and the internal firewall for the Active Directory service that allowed the attacker access to sensitive assets on the internal network. Best would be to completely eliminate dependencies on the Active Directory service from the DMZ, but this may not be possible because of architectural requirements in the application.

Here is how you can use FirePAC to identify the firewalls and the rules that have the problem and see if the changes made actually fix the problem.

1. Using FirePAC Object and ACL rule queries, quickly find all the objects and rules in the border and internal firewalls that allow the Active Directory service from the DMZ network to the internal hosts. Use the AD service and 192.168.9.2 as the source address for these queries.
2. Use the Data flow query to find all the IP packets and the ACL, NAT and ROUTE rules that allow the above packet.
3. Modify the firewalls to close the security holes in the firewalls.
4. Compare the modified firewalls for policy differences and make sure that the IP packets that are supposed to be blocked are actually blocked.
5. Check if any similar security holes exist from the DMZ host or networks in the same firewalls.
6. Repeat the same process with other paths and firewalls with in the network.

Object Query in FirePAC

Using the object query, you can find all service or address object groups that are used across all firewalls in the firewall inventory. The image below shows a service object query searching for ldap service across all firewalls. The query result window shows all the object groups that refer to this service in a tree structure. The complete hierarchy can be browsed in the results window.

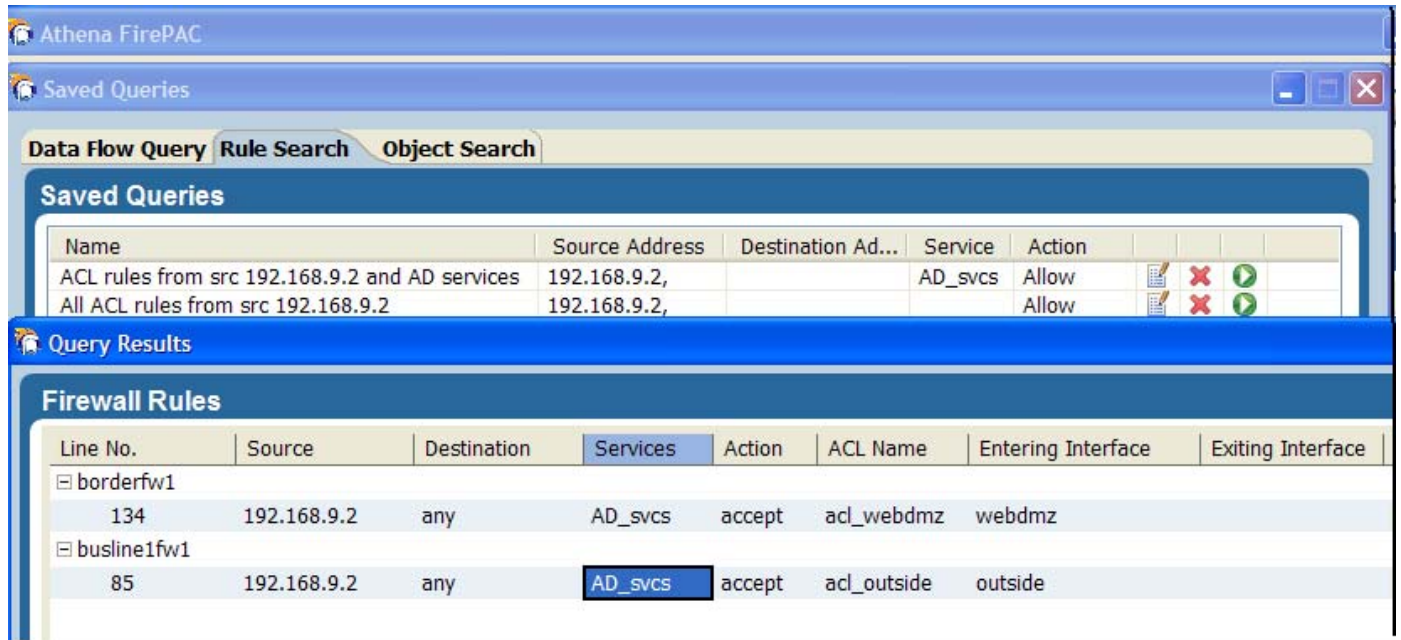


The screenshot shows the FirePAC interface with the 'Object Search' tab selected. The 'Saved Queries' section shows a query named 'Active Directory Service Objects' of type 'Service' with criteria 'ldap'. The 'Query Results' section displays a table of object groups:

Name	Type	Port	Comments
borderfw1			
AD_svcs	group		
ldap	tcp-udp	389	
busline1fw1			
AD_svcs	group		
ldap	tcp-udp	389	
borderfw2			
borderfw3			
busline2fw1			

ACL Rule query in FirePAC

Using the ACL Rule query, you can find all rules that refer to the given source, destination and service objects or ranges across all firewalls. The image below shows all the firewalls and the rules in them that have the 192.168. DMZ subnets or hosts as the source address and Active Directory as service in a tree structure. You can click on any object in the source, destination and service columns and see the complete definition of the object in a pop up.



Complete Policy Rule trails in FirePAC

Using the data flow query, you can see the actual IP packets and the ACL, NAT and ROUTE rules that allow them through the firewall entering and exiting specific zones/interfaces. The report below shows a data flow query results using 192.168.9.2 as source and Active directory service from DMZ to internal zone. The rule trails column shows as links the line numbers of acl, nat and route rules within the firewall configuration that allowed the IP packets.

Entering/Exiting Interface	Service	Src Address	Dst Address	Action	Rule Trails
DMZ : webdmz -> borderfw1 -> Internal : inside	tcp/ldap (389)	192.168.9.2	10.0.0.0/8	accept	acl:134 nat:197 route:8
			172.16.0.0/16	accept	acl:134 nat:197 route:3
			172.17.0.0/16	accept	acl:134 nat:214 route:9
			192.168.2.0/24	accept	acl:134 nat:214 route:10

Policy Rule Trails to check for Private RFC addresses

The firewall ACL rules use the PUBLIC ip addresses when defining access to the servers on the DMZ network. FirePAC data flow query allows you to specify the destination address directly as the internal RFC-1918 private IP without worrying about the translated public addresses for the Policy Report. The report below shows all policies from External zone to the internal address 192.168.9.2. This shows there is http access allowed using the public address 62.59.14.169 and the nat rule on line 213 indicates that the address 62.59.14.169 is being translated to 192.168.9.2.

```
Line 69: access-list acl_outside permit icmp any host 62.59.14.169
Line 70: access-list acl_outside permit tcp any host 62.59.14.169 eq 80
Line 213: static (webdmz,outside) 62.59.14.169 192.168.9.2 netmask 255.255.255.255 0 0
```

Services Passing Through Firewall					
<i>This section lists source and destination addresses for all IP services that are allowed to pass through the firewall device. The list of allowed services is grouped by each output interface from which the allowed IP traffic leaves the firewall.</i>					
Entering/Exiting Interface	Service	Src Address	Dst Address	Action	Rule Trails
External : outside -> > borderfw1 -> DMZ : webdmz	icmp/any	any	62.59.14.169	accept	acl:69
					nat:213
					route:6
	tcp/http (80)	any	62.59.14.169	accept	acl:70
					nat:213
					route:6

The following report shows all the services allowed from the source 192.168.9.2 into the application and internal networks. It indicates that application services to two application servers are being allowed besides the LDAP access to many internal networks.

```
Line 134: access-list acl_webdmz permit tcp host 192.168.9.2 any object-group AD_svcs
Line 135: access-list acl_webdmz permit tcp host 192.168.9.2 host 172.16.10.101 eq https
Line 136: access-list acl_webdmz permit tcp host 192.168.9.2 host 172.16.10.100 eq 7778
Line 197: static (webdmz,inside) 192.168.9.2 192.168.9.2 netmask 255.255.255.255 0 0
```

Entering/Exiting Interface	Service	Src Address	Dst Address	Action	Rule Trails
DMZ : webdmz -> borderfw1 -> Internal : inside	tcp/ldap (389)	192.168.9.2	172.17.0.0/16	accept	nat:214
					route:9
	tcp/https (443)	192.168.9.2	172.16.10.101	accept	acl:135
					nat:197
					route:3
	tcp/interwise (7778)	192.168.9.2	172.16.10.100	accept	acl:136
					nat:197
					route:3

Policy Difference Comparison in FirePAC to validate the fixes

If we look at the border firewall, the problem is the security hole created by allowing access on Active Directory service to all internal hosts not just to the required Active Directory servers that are running with the latest patches.

```
Line 133: access-list acl_webdmz remark - needs to restrict AD access to AD servers only
Line 134: access-list acl_webdmz permit tcp host 192.168.9.2 any object-group AD_svcs
```

So to close the threat, the rule on line 134 needs to be modified to restrict the access to the specific Active Directory servers. This problem can be fixed by defining an object group for AD Servers and using it in the rule on line 134 instead of “any”

```
object-group network AD_svcs
network-object host 172.16.1.50
network-object host 172.16.1.51
network-object host 172.16.1.52
Line 134: access-list acl_webdmz permit tcp host 192.168.9.2 object-group AD_svcs
object-group AD_svcs
```

After making this change to the configuration, the modified configuration can be compared to the original configuration to make sure that the access from 192.168.9.2 to the internal destinations other than 172.16.1.50-172.16.1.52 are now blocked.

In the reports below, you can see that all policies from 192.168.9.2 to anything other than the range 172.16.1.50 to 172.16.1.52 are marked as deleted policy. Clicking on the Deleted policy shows the rule trails for the original (Left) and the new configuration (Right). The rule trail for the new configuration shows implied-deny in the access-list once the ACL on line 134 is narrowed.

Entering/Exiting Interface	Service	Src Address	Dst Address	Comment
webdmz -> Left: borderfw1 -> inside	tcp/ldap (389)	192.168.9.2	10.0.0.0/8	Deleted Policy
			172.16.0.0 to 172.16.1.49	Deleted Policy
			172.16.1.53 to 172.17.255.255	Deleted Policy
			192.168.2.0 to 192.168.5.255	Deleted Policy

7	webdmz -> Left: borderfw1 -> inside	tcp/ldap (389)	192.168.9.2	10.0.0.0/8	acl:134
	webdmz -> Right: borderfw1 -> inside	tcp/ldap (389)	192.168.9.2	10.0.0.0/8	nat:197 route:8 acl:access-list-acl_webdmz-implied-deny
8	webdmz -> Left: borderfw1 -> inside	tcp/ldap (389)	192.168.9.2	172.16.0.0 to 172.16.1.49	acl:134 nat:197 route:3
	webdmz -> Right: borderfw1 -> inside	tcp/ldap (389)	192.168.9.2	172.16.0.0 to 172.16.1.49	acl:access-list-acl_webdmz-implied-deny
9	webdmz -> Left: borderfw1 -> inside	tcp/ldap (389)	192.168.9.2	172.16.1.53 to 172.16.255.255	acl:134 nat:197 route:3
				172.17.0.0/16	acl:134 nat:214 route:9
	webdmz -> Right: borderfw1 -> inside	tcp/ldap (389)	192.168.9.2	172.16.1.53 to 172.17.255.255	acl:access-list-acl_webdmz-implied-deny

Conclusion

SIEM solutions are powerful tools for identifying attacks against your network. They can provide you with the path of an attack through your network and help determine the firewall rules that allowed it to happen. This intelligence is just the starting point for remediating the problem.

The proper response to an incident requires understanding the complexities of the firewall configurations in the context of your network, identifying which specific changes will prevent it from happening again, and verifying that there were no unexpected side-effects from the change. These capabilities provided by Athena FirePAC in combination with a SIEM tool are a powerful solution for responding to security events.

About Athena

Athena offers infrastructure analysis tools that identify the precise relationship between firewall rules and network services in a single device or across a complex network. With a comprehensive focus on configuration data, Athena helps network and security engineers perform what-if analysis that reduces the reliance on diagnostics and validation by testing. Over 300 companies turn to Athena products, Athena FirePAC and Athena Verify, for standardized and consistent intelligence to reduce the time and effort required for policy management on network security devices. For more information see <http://www.athenasecurity.net/>.