360-FAAR Scenarios For Use

Evaluation of Methodologies

Suggested Solutions using 360-FAAR

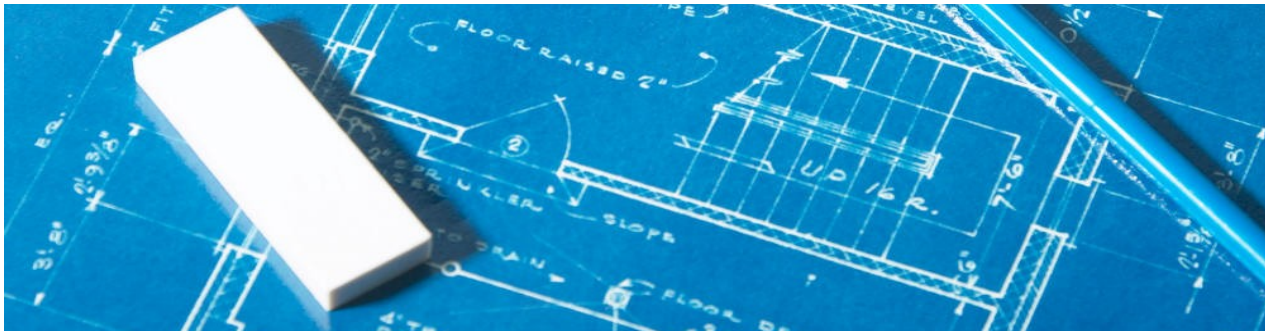

Purpose of This Document

This documents purpose is to pose every day scenarios that are found commonly in networks. These scenarios have, historically, not been well resolved using standard tools or manual methodologies.

The reasons for this lack of resolution are then examined and SOLUTIONS are proposed using the 360-FAAR (360° Firewall Analysis Audit Repair) program.

360-FAAR is capable of standard rule evaluation techniques but they are not covered here as the intention is to cover scenarios that require new tools to resolve.

Intended Audience

This documents intended audience are technical professionals, technical managers, operations teams, and company directors interested in technologies that offer a distinct advantage over competitors solutions.

About the Author

Dan Martin is the Director of 360 Analytics Ltd. Prior to this, he worked as a Network and Security Analyst/Engineer for fourteen years, employed by companies such as Intel, Nokia Internet Communications, Sun Microsystems, Qualcomm, Verison, Diageo Ltd, Party Gaming Ltd. and 'The Cloud' wifi network.

Table of Contents

End of Contents

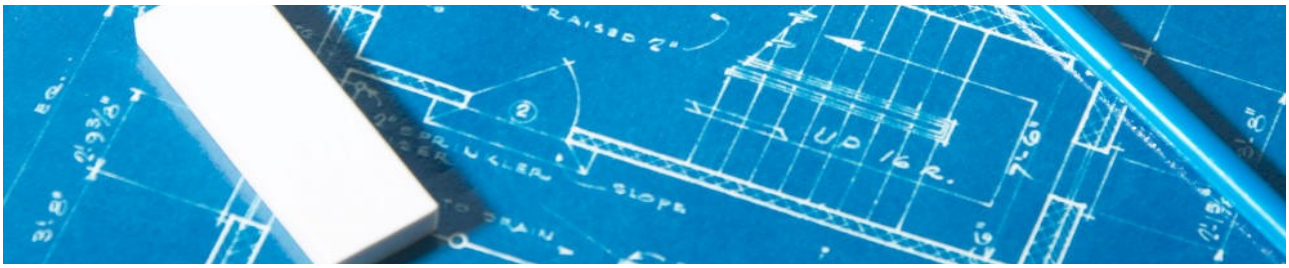

About The Scenarios Described In This Document

The following twenty sections describe simplified or amalgamated, but none the less real world, scenarios in which the 360-FAAR program has been used successfully, in interesting and effective new ways, to solve common network and firewall issues.

Many of the scenarios described below have been encountered during customer projects and references from the technical managers who requested the work and saw it completed can be provided.

It is hoped that by reading these examples you will be able gain a feel for the capabilities and potential uses of 360-FAAR.

Many of the later sections build on the examples given during the first five sections so it is recommended that you read these examples first... many will be familiar problems, resolved in new ways.

1. The Firewall Rulebase Cleanup

1.1 Reasons for Cleanup

Firewall rulebases are some of the most persistent structures within any network.

In large environments the firewalls have often been managed by several different service providers, each with their own object naming conventions, rule management conventions (rule grouping / organisation / section types) and various security policies that were agreed with the customer at different points of the firewalls life.

This inevitably leads to inconsistencies within the firewall policy/rulebase.

Objects are duplicated with different names, rules (or subsets of rules) are duplicated over time, especially in large rulebases where the problem of 'rule usage visibility' restricts the usefulness of managing the rulebase with traditional methods (manually with a GUI or terminal).

Old rules may violate current security policies, but cannot be removed because of the general nature of the rule or because of a lack of 'rule usage visibility' of the traffic the rule is permitting.

1.2 Restrictions of Current Methodologies and Tools

The manual method for performing a cleanup on large scale firewalls is very time consuming and is a user/engineer driven manual process.

This means that an engineer will need to evaluate each rule in the rulebase to assess its current usage profile and decide weather to keep or to drop it after asking the user/customer/service provider for confirmation of its use or redundancy.

The process usually entails calling the services provider to request information on the services listed in the rules concerned, and asking weather the rule is needed.

This is a very inefficient process, fraught with human error, that requires 'Local Knowledge' on the part of the specific engineering departments involved (sometimes years after the event) which is often spread amongst many project groups that are difficult to identify from historical documentation.

It also requires that duplicate rules and other such anomalies are found and evaluated MANUALLY by the engineer to ensure that unused rules or rule duplicates (or duplicate objects within rules) are removed.

This makes it very difficult to strip unused objects or even rules from large rulebases manually (in the order of a billion cross references for a usual midrange - 500 rule, 1000 object, 100 group-rulebase) with any certainty that the objects removed will not affect current services.

It relies heavily on the individual engineers skill and tenacity!

Once you have completed this process, the chances are, that in a large organisation the rulebase will have changed sufficiently (within the time of the project) so that matching the rulebase used for the analysis to the current rulebase becomes difficult in its self, causing yet more work for a dedicated engineer or team, adding unneeded engineering time to the project.

1.3 Solution Provided by 360-FAAR

360-FAAR solves this problem by automating the firewall cleanup operation.

It identifies ALL connectivity used in the log files, applies this as a filter across the rulebase and outputs a clean version of the current rulebase and and a rebuilt optimised version built by the 360-FAAR policy engine.

A years worth of logs will generally take 2 to 5 DAYS to process on a reasonable size server, after which you are presented with a WIZARD that can be run many times to build as many new configs as you choose from your existing firewall policy.

Once the log files and configurations are loaded into 360-FAAR, you simply request that the logs

are used to filter the full configuration and that the new configs are output as spreadsheets showing the newly proposed rules and the commands to update the firewalls AUTOMATICALLY in their NATIVE languages.

This requires no extra hardware on the network and no additional connectivity to or from the firewalls management station, it is completely off-line!

The firewall commands are generally text (so they can be read and cross checked by in-house engineers easily) and can be cut and pasted in to the firewall or uploaded in file format (dbedit, netscreen, cisco supported).

This procedure can reduce the time a firewall cleanup project takes from 6-12 MONTHS to 1-2 WEEKS, including reprocessing and updating the rules that the customer requested removed, changed or updated from the spreadsheets created by 360-FAAR.

Even rules or objects from other firewalls can be cut and pasted into the new rulebase spreadsheet so long as all the firewall configs used are available for processing.

360-FAAR will identify the sources of the objects and rules automatically and translate them (between manufacturers or firewalls) before rebuilding the firewall commands needed to match the updated spreadsheet.

This allows all parties with a stake(!!) in the firewall policy cleanup to examine the proposed rules and make changes that can be automatically included when the final commands are generated from the updated spreadsheet.

Also, any rules can be tagged for retention -such as DR rules- to ensure that they are not removed, and the tags can be firewall specific.

Simplified UNUSED rules will be output in a separate spreadsheet ready for easily analysis of any rules that maybe needed for the new policy, that were not used and were set to log traffic (rules set not to log are required to be included by the policy engine as they cannot be analysed for their usage although they can be filtered for addresses, services, text strings, object names, comments, names or section headers).

2. Moving Networks or Routes Within Network Infrastructures

2.1 Reasons for Route or Network Moves Between Firewalls

During a firewalls lifetime many of the networks it protects will be de-commissioned, new networks, on which new services are hosted will be brought into commission.

Old services may need to be relocated to new more suitable network locations and servers will be reassigned for new services.

Companies may choose to change service providers and readdress their networks or possibly move an existing provider independent IP space to another place in their own network or to Ipv6 address space.

Each of these tasks requires in-depth firewall analysis and a lot of manual work for the engineering department.

2.2 Restrictions of Current Methodologies and Tools

A network is a dynamic entity in almost every facet except for its firewall policies.

The policies are specific to manufacturers methodologies. Each manufacturer has their own style of configuration and language for rule creation.

Despite this a companies rulebases are often treated as malleable structures by service provider companies.

The standard tools to move or merge rulebases are manufacturer specific, are considered to be ineffective across manufacturer boundaries, often do not have the functionality to deal with moving specific subsets of a firewalls policy rather than the whole thing and are not production grade for large environments... that is until 360-FAAR.

As a result the engineers employed for the project are generally expected to write the new rules using only manual (or automated yet unrelated to the firewall config) log analysis and intuition.

Often this results in a route move project that translates only the 'known good rules' and plans to fix 'on the fly' (known more generally as 'by the seat of your pants') the rules that have not been captured by the manual or policy-less analysis.

The trouble with this standard approach, is that you only find out a rule is missing when a service fails because you only get visibility of its use at this point.

2.3 Solution Provided by 360-FAAR

With 360-FAAR the result of a firewall analysis project is no longer defined by the subjective opinion of the engineer tasked with completing it.

The CIDR ranges of the networks that are intended to be moved can be entered into 360-FAAR and the connectivity from all rules will be evaluated for connectivity to this CIDR range.

Rules that entail connectivity to the chosen network will be stripped of only the connectivity requested to be moved and 360-FAAR will build new rules consisting of exactly the connectivity specified that can be applied to a new firewall, while the commands to remove the original connectivity can be applied to the original firewall.

The new rules built can be configured to consist of only connectivity that is not already specified in the destination firewall, and if requested will use all object definitions from the destination firewall that match the objects being moved, as well as creating all new objects needed with their original properties preserved (even an objects colour can be retained across firewall and manufacturer boundaries).

The process can additionally include filtering of the newly created rules using the log files loaded into 360-FAAR, which provides as output only the translated filtered rules that are in use.

The Newly proposed rules and objects are output in spreadsheet format so that they can be modified easily and signed off by the customer before generating the signed off rules, and objects for application to the destination firewalls and to remove the connectivity from the rules of the source firewalls.

Many firewalls can be used at once as the source of rules to be combined and filtered to create many firewall configs for many differently configured firewalls across many different firewall manufacturers... many many possibilities indeed!

3. Network Object Translation to Newly Assigned IPv4 or IPv6 Addresses

3.1 Reasons for Assigning New IP Addresses to Objects in Firewall Rulebases

Many firewall projects require that a group of servers be assigned new IP's as a consequence of network changes or the need to readdressed networks during network restructuring.

These projects are often costly to the customer simply because of the engineering time required to isolate the servers or networks to be translated within the multitude of rules.

Soon ALL IPv4 address will have been assigned for usage. Unfortunately at the point in time when this happens it will necessitate a ramp up of IPv6 networking to 100% of all new addresses assigned almost overnight.

This leads to many ugly consequences (such as transit provider NAT gateways etc. some of which are unavoidable) and theses in turn lead to large network restructuring projects in order to incorporate the new addressing scheme.

3.2 Restrictions of Current Methodologies and Tools

Many of the issues discussed in the other sections of this document can play a part in adding to the complexity of translating a firewall policy and its objects to new addresses.

The process involves removing old connectivity and adding new connectivity in its place, often in very large environments.

This process entails a huge amount of work if undertaken manually and requires many man hours of expensive expertise, currently tools to automatically do this work are restrictively expensive and largely unused.

3.3 Are Other companies Offering CIDR Translation to IPv4 or IPv6 in this way?

360-FAAR is the only current tool that is capable of re-writing rulebases using CIDR translated addresses for ALL HOSTS, NETWORKS and RANGES specified within a rulebase.

No other tools offer this methodology for readdressing objects across the spectrum of firewalls we support (360°!)

3.4 Solution Provided by 360-FAAR

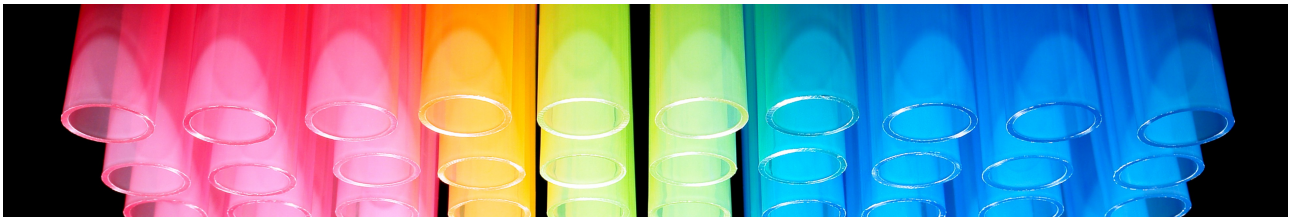

360-FAAR can translate ALL objects found within a rulebase to new IP addresses using:

- Simple translation tables (two column spreadsheets, from CIDR and to CIDR) or
- Inclusion/Exclusion tables
  (a four column CSV – In/Out, In/Out First/Last Priority, From, To,)
- Masked Bit Switching (Swapping of masked portions of the IP addresses)
- Other custom methodologies (such as higher order octet subnetting etc)

It is then possible to build rulebases using these new objects.

Networks, servers and ranges can be readdressed by replacing an address's masked bits or using translation tables. This methodology can be used for translation to IPv4 or IPv6 addresses.

360-FAAR is also capable of matching these new IP addresses, networks or ranges with current addresses, networks or ranges that exist in the destination firewall already, removing any possible duplicates that would have been created by a manual process of readdressing each of the objects.

The Newly proposed rules and objects are output in spreadsheet format so that they can be modified easily and signed off by the customer before generating the signed off rules, and objects for application to the destination firewalls and to remove the connectivity from the rules of the source firewalls.

New rulebases can be ordered by service, alphanumerically or by port, by source or destination column names, or by IP address numbers.

If log files are loaded into 360-FAAR's policy engine the rules can be ordered by usage, either by object hits or connectivity hits.

4. Change Object Naming Conventions

4.1 Reasons for Renaming Firewall Objects

A firewalls objects are its understanding of 'places'.

In the real world a places name does not often change, however, in the digital world resources and locations are much more likely to be used for more than one service during the lifetime of a firewall and may be called by many names during this time.

The names assigned to addresses tend to reflect an objects first usage profile. This name becomes less relevant the more often the address is assigned new services.

Often to counter the reduced readability that incorrect object names create with a policy duplicate objects are created to signify the new services an address is providing. This in turn leads to rulebases that are difficult to manage with many definitions for a single address, range or network, reflecting the networks historical usage.

4.2 Restrictions of Current Methodologies and Tools

A network is a dynamic entity in terms of address allocation, dynamic routing and other such services. The tools used to manage objects names have not kept pace with the development of automated IP services and DDNS (Dynamic DNS).

To a large extent this is due to quantifiable security concerns with the potential dynamic objects present as an attack vector through any networks security policy (e.g. to inject a dynamic address into the firewall permitting connectivity to an IP of your choice).

Object definitions are specific to manufacturers methodologies. Each manufacturer has their own style of configuration and language for object creation.

The standard tools to analyse object names for adherence to naming conventions are considered to be ineffective in most situations, often do not have the functionality to deal with renaming specific subsets of a firewalls policy rather than the whole thing and are not production grade for large environments.

As a result the engineers employed for the project are generally expected to run manual searches using a firewalls policy editor or by examining the config with a text editor.

Names are notoriously difficult to predict and when an engineer encounters many duplicate objects used in several different firewall rules or policies this increases the complexity of his/her job significantly.

Often this results in a naming convention consolidation project that translates only the 'well known' object definitions and assigns more complicated clean-up operations to a current risk list.

4.3 Solution Provided by 360-FAAR

360-FAAR resolves the problems associated with manual name configuration within firewall policies vs DDNS style name to address mappings in networks, by cross referencing an approved text based zone file (possibly output from a DDNS server in your AD or network infrastructure at a specific and recent time).

360-FAAR locates all name to IP mappings.

These mappings are used to locate all objects with known IP addresses from the firewalls policies and 360-FAAR's policy engine rebuilds a firewalls objects and rules using the newly assigned names from the current zone file as well as removing all object duplicates and using the newly named object in the place of all duplicates removed.

New rulebases output, can be ordered by service, alphanumerically or by port, by source or destination column names, or by IP address numbers.

If log files are loaded into 360-FAAR's policy engine the rules can be filtered or ordered by usage, either by object hits or connectivity hits.

5. Firewall Policy / Group Reassignment and Restructuring

5.1 Reasons for Reassigning a Firewalls Objects to New Groups

Grouping objects within firewall rules allows administrators to quickly identify functional groups of servers, networks or ranges.

These groups can be used to write very descriptive rules that are easy to manage and reflect the policies purpose in a way that simply adding objects to firewall rules cannot.

However, a firewalls policy and function changes over time.

Server and object usage change over the life of the firewall as do the relevance of the groups configured into the firewalls policy, rendering many object and group names incorrect after only a few months or years.

5.2 Restrictions of Current Methodologies and Tools

A firewalls groups are currently one of the most difficult elements of the configuration to remove or change using conventional tools or manual analysis

Complications, when dealing with moving or removing or updating groups, arise mostly when dealing with nested grouping structures. These structures are relatively easy to unwrap but no so easy to wrap up again

Identifying large numbers of objects group memberships and locating those groups within the nested structures manually quickly becomes a difficult task on even medium size firewall installations.

Once you have located the objects and groups you then need to determine their usage in rules and also their inclusion in larger groups used in sometimes separately configured rules.

5.3 Other Companies Offering a Group Restructuring Function for Firewall Policies

- At present we are not aware of any other companies offering a group restructuring methodologies for firewall policies.

5.4 Solution Provided by 360-FAAR

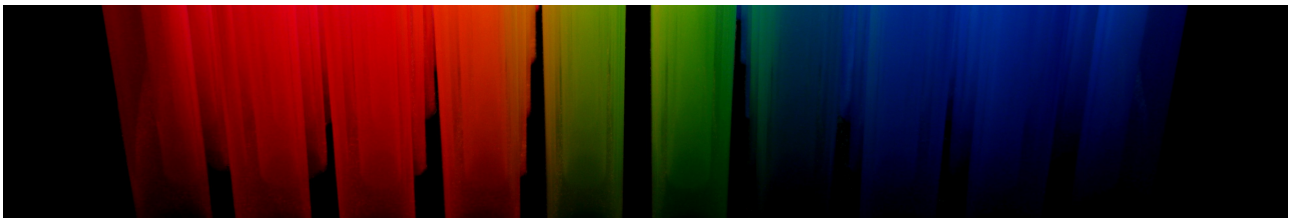

360-FAAR is, not only capable of expanding out any grouping structures within a firewalls policy, but is also capable of reconstructing new structures from the existing connectivity to replace inconsistent policies in production systems.

Infinitely deep grouping structures are handled easily and a groups subgroup/supergroup memberships are tracked against all connectivity existing in the rulebase.

New groups can be interjected at any point and old groups can be removed while leaving the objects they contained (including or excluding subgroups) in place in the rulebase.

This means that a firewall's groups and rules can be re-written by the 360-FAAR policy group engine using a newly defined set of groups and also restructured and reordered while maintaining the same connectivity.

This process uses a system of 'Sequential-Rule-Masking', and set of priority based choices to un-group all connectivity specified within a policy, and replaces groups of objects with their associated group when all objects from a group have been found to have a similar connection profile.

No new connectivity is ever created by this process, it should be considered a methodology that is useful for examining existing connectivity and reorganising it in new ways!

It can dramatically reorganise a firewalls rulebase in minutes, reducing rule numbers (sometimes by as much as 60%) while maintaining all connectivity and increasing its readability many times (all original rule details are retained and can be added to comments, section headers or other properties of the new rules).

Once the rules have been processed 360-FAAR outputs the newly proposed rules and commands to automatically update the firewalls in spreadsheet format that can again be easily updated or changed before reprocessing to generate the final commands to be applied.

New rulebases can be ordered by service, alphanumerically or by port, by source or destination column names, or by IP address numbers.

If log files are loaded into 360-FAAR's policy engine the rules can be ordered by usage, either by object hits or connectivity hits.

6. Close Open Rules

6.1 Why Firewalls Have Policies With Less Than Optimal Rulebases

It is not uncommon to find at least one firewall within a organisation that has one or more problematic rules in use in its policy.

These rules can be historical, but allow a sufficiently large proportion of a firewalls traffic that removing them is not an option.

This does not help to resolve the security engineers problem, how to secure this firewall rule.

Other rules may have been specified during at test phase of a services life and were unable to be removed because of connectivity issues at go live day.

Business needs in situations like this are such that an engineer may have had to add a more general rule than a policy requires in order to resolve unforeseen connectivity issues.

6.2 Restrictions of Current Methodologies and Tools

The most common methodology used in operations departments to close open rule definitions is manual (or automated but unrelated to firewall config) log analysis and to match rule numbers, services and objects manually to rules profile.

If, however the log file was archived at a previous time, the changes to its current policy will most likely render the numbers out of date adding complexity to an engineers job.

Some firewall analysis tools do offer rule usage details, for very general rules however, consisting of a few simple but large objects, these 'in use / not in use' flags are not sufficient to build the new connectivity needed to remove general rules and replace these with targeted rules defining only the connectivity in use.

6.3 Solution Provided by 360-FAAR

360-FAAR's policy engine permits individual rules to be compared to all connectivity found in the log file.

A report consisting all the objects and services found to be communicating through a rule, all connectivity a rule permitted, top talkers using the rule, and the most specifiic rule definitions required to permit the traffic seen that are possible to be built from your firewalls current config.

For rules that are yet more specific, 360-FAAR can build new objects to configure the tightest rules possible, if requested.

7. Security Policy Enforcement

7.1 Uses for Policy Enforcement

Many companies, in fact most, have security policies to protect their networks data. These security policies are enforced in practice by your firewall policy and other network (such as IDP) devices.

These policies protect vital assets your company possesses, however they are restrictive for day to day operations of the network and are often violated by high priority business needs, before the services make it to more stable locations.

7.2 Policy Enforcement Provided by 360-FAAR for Firewalls WITHOUT Zone Methodologies

CIDR to CIDR + Inclusive / Exclusive filters of objects and warnings of traffic or rules that permit traffic violating the policy prescribed, e.g. an input policy file could contain a three field CSV (In/Out,IP,NM) listing permitted or restricted network connectivity or possible the four field format described above could be chosen as an input method instead. Almost any format that is consistent can be read by 360-FAAR's custom reader sections.

The statements listed in the input file are used as a filter to search the rulebase for connectivity in violation of the statements.

Once a hit has been identified it is reported and potentially tighter security rules are suggested.

7.3 Policy Enforcement Provided by 360-FAAR for Firewalls WITH Zone Methodologies

Netscreen and Fortigate firewalls employ zone based methodologies.

Policies using this style of configuration are readily mapped to security policy documentation because of the similarity between conventions for writing security policy documents and the concept of zones or areas such as DMZ's.

360-FAAR's policy engine will report any violations of security policy specifications. These are listed in the form of inclusive and exclusive statements regarding zone to zone communication.

Any rules found to be in violations of these policies will be reported and corrections suggested.

8. Split Large Policies Into Smaller Policies for Virtualisation

8.1 Reasons for Virtualisation

Since the advent of firewall virtualization environments (such as MDS or Crossbeam) service providers have been able to offer customers their own virtual firewalls separate from other organisations infrastructure yet running on the same hardware.

This is cost effective and can increase manageability and security - if done well.

A large firewall policy can be split to make smaller individual policies from each individual customers rules.

Service providers are also aided by virtualization technologies because it allows them to reorganise their own network infrastructure easily with out the need to change their customers environments in the way they would if they were to relocate a firewall with a single large policy.

In well managed rulebases this can often be achieved by simply moving a customers section of the rulebase to the new virtual firewall, however, matters are rarely this simple in practice.

8.2 Restrictions of Current Methodologies and Tools

Often a companies management connectivity will be specified in large firewall rulebases, some of which will be managed by outside entities.

As a result it is far more difficult to isolate and to ensure 'no service is left behind' during network operations tasks.

A companies response is often to virtualise their environment. However, large numbers of man hours are required to analyse large firewall policies and each simplified policy created (for the new smaller virtualised firewalls) adds to this total.

A quick fix that is used from time to time in virtualised environments is to copy the complete original rulebase to each of the new firewall and attempt to remove each new firewalls unneeded rules.

We regard this as a false economy because the new virtualised rulebases diverge from each other quickly after their installation and this increases the complexity of a firewall rulebase cleanup project that now needs to be performed for each of the new firewalls.

8.3 Solution Provided by 360-FAAR

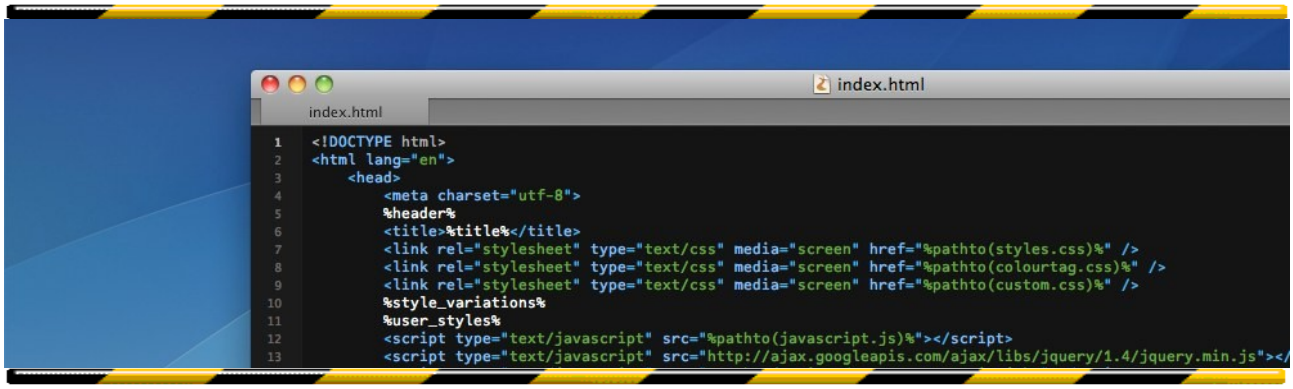

360-FAAR permits you to filter, translate and subdivide existing rulebases within its core processing.

After having loaded your required time periods log files, you are presented with a WIZARD that permits you to text filter, CIDR filter, zone filter, group filter your existing policy to create new policies that are output as spreadsheets and firewall commands to automatically make the required changes... all completely off line.

The spreadsheets can then be modified and re-uploaded to 360-FAAR which generated the new command sets needed to make the changes.

See points 1 – 5 for more information.

9. Merge Firewall Configurations Together Seamlessly

9.1 Reasons for Consolidating Firewall Policies

Firewall policies are usually specific to a firewall location within a network.

Changing a firewalls location or consolidating firewall hardware often requires merging firewall policies together.

As network hardware grows in processing and throughput capacity less physical platforms are required to carry the networks traffic.

Firewall policies also need to be consolidated during management take over (MTO) tasks and during network restructing avtivities.

In many cases, several firewalls policies will need to be considered to ensure that a policy for a specific firewall within an infrastructure will be removed of connectivity that is routed via other firewall in the network.

9.2 Restrictions of Current Methodologies and Tools

The tools currently available for merging firewall rulebases are manufacturer specific and in many cases cannot filter or translate subsections of a firewall configuration.

Merging complete firewall policies consisting of rules (Security and NAT), objects and groups requires large cleanup operations on the merged config to ensure that the duplicated objects and rules are removed before the policy goes into production.

The more firewalls that need to be considered during the merges process the larger the firewall cleanup operation and associated cross referencing will take.

9.3 Solution Provided by 360-FAAR

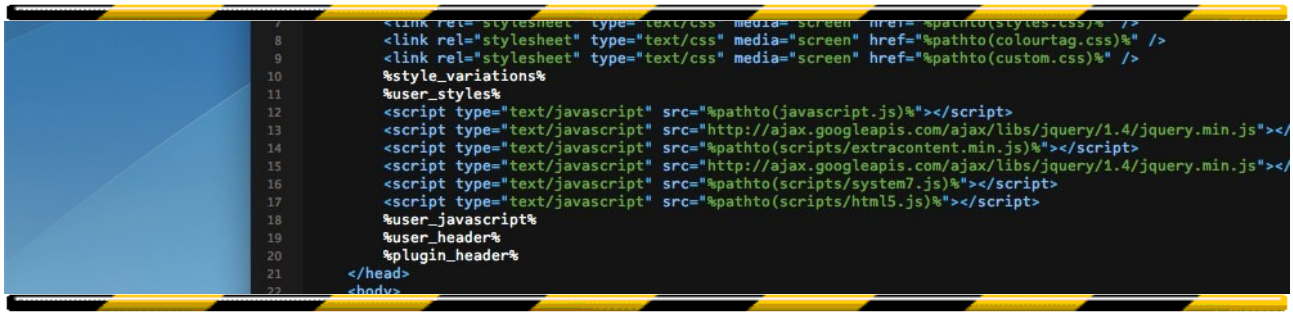

360-FAAR's policy engine is capable of loading as many complete firewall configurations as the server you are running it on has memory to store!

From these configurations you can select to pull rules objects and groups from as many firewalls as you choose as input for the merge process.

You can also use any of the other functionality of the 360-FAAR policy engine to filter and automatically select rules to merge. See sections 1-5 for more information.

You can also select to pull rules, objects and groups from many firewalls to use as a 'merge to' rulebase.

The 'merge to' rulebase first merges together the connectivity selected and then uses this to eliminate any matching rules that were built when objects and rules were selected as the source of the merge.

This is especially useful if an intermediate firewall between the two (or more) firewalls you are moving configuration information between is encrypting some of the traffic that was originally in the clear.

The VPN rules specifying the encrypted tunnel can be added to the 'merge to' rulebase and connectivity matching their profile will be removed from the final rules to be applied after the merge.

Rulebases can be translated between manufactures and firewalls reprocessed to be relevant to a different firewall and all rule, object and group definitions are output in the format described previously.

All connectivity matched or found to be duplicated as the rules are read and built into the various rulebases is reported and removed, as is any connectivity that is matched to the log files or that matched filters.

Very large amounts of debug are available for the merge process and can produce 10 or more GB of data if fully enabled.

10. Translating Between Firewalls and Manufacturers

10.1 Reasons for Translating Rules and Objects Between Firewalls or Rulebases

Modern company networks rarely have a single firewall infrastructure and often different departments will use firewalls from different manufacturers or with differing policy types.

Services will, on occasions, need to be moved to new locations within a network that require several firewalls to be updated along the new route.

10.2 Restrictions of Current Methodologies and Tools

Network changes of the type described above are almost always done manually in present day operations teams or by specialist contractors.

These resources are expensive and the time taken to do through analysis in this way is often restrictive and can mean that budget constraints win out over engineering tenacity.

One of the problems with this approach is the number of firewalls the new connectivity will need to be added to.

Cross checking (for duplicates and group mismatches) is time consuming but is an essential process if an engineer is to be sure the rules they are applying mean the same thing to all firewalls they are applied to.

10.3 Solution Provided by 360-FAAR

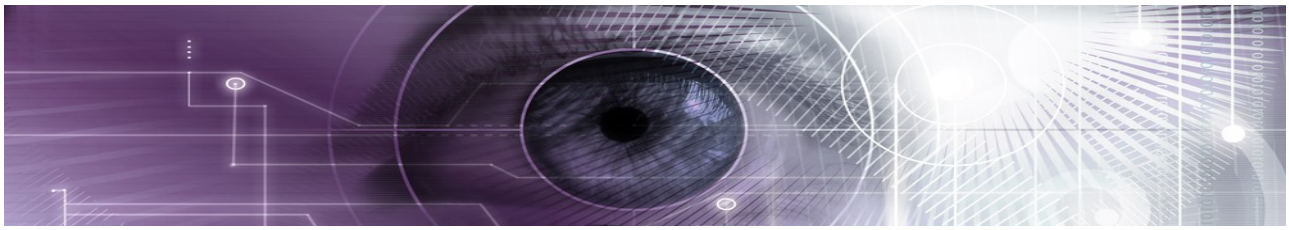

360-FAAR gives you incredible visibility, permits you to filter, from a current rulebase, only the connectivity you require and to then translate this to another firewall, removing all the duplicate connectivity or object definitions that would be created if rules and objects were not translated and checked for existence already before being applied.

See points 1 to 5 for more information.

In Netscreen firewalls it is impossible to create duplicate objects so rules specifying duplicated objects that were dropped by the ScreenOS will not be correctly applied if the objects they contain are not correctly translated.

Lists of the previously existing connectivity that matched the connectivity you were adding are provided, as are translated objects and details about partial matches that were not in the end used but are reported to increase visibility during the process.

Once the 360-FAAR policy engine has processed the requested rules, their definitions and the commands to create them are output in spreadsheet format.

These spreadsheets can then be re-read by 360-FAAR in order to translate these rules across many firewalls using the existing objects and groups from each destination firewalls configurations in turn and creating new objects only were needed.

Detailed log analysis can be performed on each of the source and destination rulebases using the results of the above process as the input rules to be checked for use.

This provides a very powerful tool for moving rule and policies between firewall and across manufacturer boundaries as though they are dynamic structures.

11. NAT Rule optimisation

11.1 Reasons for Firewall Policies to Have Less than Optimal NAT Rules

Firewall NAT Policies have many uses.

From hiding the identity of the hosts communicating, to readdressing server destination addresses or for managing routing concerns in multi-homed systems.

These layers of network addressing sit one on top of the other and form part of the underlying routing infrastructure despite describing addresses that do not physically exist in many situations.

This means that removing these rules poses network connectivity issues as well as security concerns, as a result few are removed.

11.2 Restrictions of Current Methodologies and Tools

Current methodologies to analyse NAT rule usage are far and few between.

Log Analysis provides only half of the picture and does not necessarily associate NAT translations logged with the individual rule used to translate the traffic

Also, very few policy editors have the needed functionality to provide usage stats for NAT rules.

Add to this the complexity of looking for differently addressed traffic on either side of the NATing firewall and you very quickly have a situation where NAT rules enter a 'stack' and the stack grows exponentially higher as visibility of the NAT rule usage drops.

11.3 Solution Provided by 360-FAAR

360-FAAR uses an unusual but highly effective strategy for analysing NAT rule usage.

All original NAT rule definitions are exported in spreadsheet format.

360-FAAR performs log analysis on all the traffic seen and NAT translations logged. It also matches traffic that matched NAT rules but was not translated (i.e. because it hit a policy that was defined at a higher location in the rulebase).

360-FAAR then builds a new spreadsheet listing each of the objects or IP's found to be using the NAT rule in the format:

NAT Rule Num, Source Object or IP, Destination Object or IP, Service Used, Translated IP

This format permits an engineer to quickly establish which networks or hosts are utilising a NAT rule currently, which hosts or networks can potentially use the NAT if the policy were to be altered and which NATs (by their absence from the usage spreadsheet) are not in use.

If you choose to automatically restructure your NAT statements you have 3 options available to you:

- CIDR supernet and subnet matching is available for all route or policy based NATs.
- Traffic Analysis from log files can be used to establish which NAT rules passed traffic and which rules can be safely removed
- Rule Order analysis can identify policy statements that mask later rules permitting subnet/supernet NAT connectivity, and while sequentially checking this connectivity against the rulebase all subnet supernet connectivity masked is removed.

Spreadsheets are output containing the updated rule suggestions and commands to create them automatically.

These spreadsheets can be edited in the ways described previously in this documentation

12. VPN Rule Optimisation and Simplification

12.1 Reasons for Firewall Rulebases to Contain Less than Optimal VPN Rules

A companies VPN rules have most usually been built over time.

As a result the VPN rules in use are often fragmented, include overlaps between VPN encryption domains, and in some situations can be asymmetric.

Large VPN Rulebases reduce the efficiency of a firewalls cryptographic capability as well as its traffic forwarding capabilities, introducing network latency and reducing firewall throughput.

Large VPN concentrators often provide more than one methodology to implement VPNs, such as Policy Based (initiated because of a policy match) or Route Based (initiated because traffic is routed to a tunnel).

This can, in turn lead to large numbers of unneeded IKE sessions increasing the size of memory the firewall needs to hold its state tables.

12.2 Restrictions of Current Methodologies and Tools

Large VPN Rulebases make visibility of a rules usage very difficult via manual log analysis.

Analysing these rules requires large numbers of man hours and highly technical engineers to simplify or rationalise the rulebases.

Policy and Route based VPN can be configured on either end of a VPN connection. Depending on the policy's specific configuration, this can lead to different tunnels (different SA pairs) providing from and to connectivity.

In cases where policy VPNs specify both ends of the VPN traffic can be routed asymmetrically simply because of a policy entry (or rules) location within the rulebase.

12.3 Solution Provided by 360-FAAR

360-FAAR is capable of analysing VPN rules using several methodologies

- CIDR supernet and subnet matching is available for all route or policy based VPN's.
- Traffic Analysis from log files can be used to establish which VPN tunnels passed traffic
- Rule Order analysis can identify policy statements that mask later rules permitting subnet/supernet connectivity.

The process presents the user with lists of options and walks you through the process of choosing the options you need to isolate the VPN traffic of interest.

Once 360-FAAR has analysed a VPN policy it will output spreadsheets containing the newly proposed VPN rules or Tunnels and the commands to create them. These can be used as template commands to configure the VPN's but will not include any secret keys or certificates. This is because both ends of a VPN tunnel are assumed to require analysis and when analysing large numbers of VPN rules it is inevitable that some of the firewalls at the remote end of the VPN will not be within a companies control.

However, once you have visibility of the VPN's in use and templates to update the connectivity VPN cleanup projects are far less problematic.

Client Based VPN's can also be considered. Methodologies for analysing Client Based VPN's are specific to a manufacturers firewalls and are beyond the scope of this document.

13. Security Policy Optimisation

13.1 Reasons Why Firewall Rulebases Need Optimisation

Firewall rulebases are structures that have traditionally been built over time or by merging rulebases together.

Firewalls have often been relocated or are handed between operations departments or companies support departments, rulebases are merged with existing rules or are moved to sub sections of more complicated rulebases.

In some cases policies have been merged wholesale, duplicating both objects and rules in single policy.

Some rulebases contain rules for more than one firewall and many of these firewalls rules may not be needed or are duplicated across all firewalls.

Simplifying these rulebases or organising them in more sensible ways increases security, throughput and allows greater visibility for the engineers managing them.

13.2 Restrictions of Current Methodologies and Tools

Firewall policy cleanup projects to restore order to rulebases are highly problematic and require skilled engineers to be able to assess the exact connectivity that needs to be moved and how to organise the new rulebase.

These skills are expensive!

Tools to optimise rulebases are, for the most part, based solely on the policy it self and do not consider the log files in the context of a current rulebase.

Log files are considered but are analysed separately from the policy providing network usage details and traffic analysis information, but not policy usage information.

Engineers are required to combine the various sources of information manually despite the tools having carried out the individual analysis components automatically.

13.3 Solution Provided by 360-FAAR

360-FAARs policy engine allows engineers to combine a firewalls logs and its configuration to create a detailed map of all connectivity used, hit counts specific to individual types of traffic and the locations of connections source and destination points.

This map can be related to the firewalls policy and can be used to provide detailed information regarding each rules usage in their current format.

The policy engine can also build new rules from the existing policy and apply the connection map to these new rules instead.

The rules output can be ordered via usage statistics either based on connectivity or how often an object spoke or was spoken to etc. They can also be organised alphanumerically or by their netmask or IP address.

The newly generated rules are output in spreadsheet format containing the rule definitions in readable format and also the commands to create the changes needed to update the firewalls automatically.

This process is completely off-line.

For more information regarding the optimisation capabilities of 360-FAAR see sections 1-5.

All of these processes can be used during a single pass of the many configurations it is possible to load into 360-FAAR.

14. Removing or Decommissioning Networks

14.1 Reasons for Removing Firewall Objects from Policies

Firewall objects are representations of networks, addresses, services and groups that are known to the firewall.

As the surrounding networks change, so to must the objects and rules that relate to them.

Some networks are inevitably decommissioned and need to be removed from the firewall policies of the surrounding firewalls.

14.2 Restrictions of Current Methodologies and Tools

Removing the connectivity of well used networks from large numbers of firewalls when such networks are no longer needed or services are consolidated elsewhere is a laborious procedure.

With very little payoff (removing unused connectivity from firewalls is risky and is usually invisible to the business) engineers are often reluctant to accept very large decommissioning projects.

This reluctance is due to the difficulty of locating all objects within the multitude of a companies rules spread across many rulebases and firewall types.

Add to this the issues with naming conventions, duplicates and group structures discussed earlier and simply finding all the connectivity you intend to remove becomes more problematic than is usually anticipated.

14.3 Solution Provided by 360-FAAR

360-FAAR is capable of finding subnet or supernet connectivity in rules permitting traffic through many different firewalls and generate spreadsheets and commands proposing changes that can be automatically applied to firewalls.

An objects group memberships, and a groups subgroup and supergroup memberships are tracked for infinitely deep groups.

Using a range of filters you can easily locate objects and rules from many firewalls, assess their use and specify to remove their connectivity from the policies of your choice.

This is a hugely powerful tool and is handles intrinsically by 360-FAAR's policy engine. This engine is equally capable of renaming networks objects and rules, removing them or translating subsets of policies.

Removing objects is one aspect of this functionality, there is plenty more!

See sections 1-5 for more information.

15. Build New Policies From Objects, Groups and Logs

15.1 Reasons to Build Firewall Rulebases from Logs and Existing Object and Group Definitions

A common scenario, when the ability to build a policy from log files and known objects can reduce a projects run time significantly, is during the procedure of replacing a router with a firewall.

In this situation a firewall can be swapped into the place of the router with a initial policy that permits and logs all traffic passed.

Another, less common but more serious scenario, is when a policy has been miss-configured and a rule permits too much connectivity, or in worst case scenarios everything.

Sometimes these rules can go unnoticed for some time, requiring detailed log analysis to ensure that removing them will not break services that have been setup since the rules introduction.

These services are likely to have happened to 'just work' at the time of their implementation, but are in fact, using the open rule.

15.2 Other Tools to Build Rulebases from Logs

Other tools to build rulebases from log files do exist. They rely heavily on manual intervention to choose between rule choices and often build host heavy rules that are difficult to manage.

No other production grade tools permit objects to be selectively copied from many firewalls and used to create new policies with simplified rules.

15.3 Solution Provided by 360-FAAR

360-FAAR is capable of filtering objects from selected configuration files as well as from selected rules.

It is possible to use all the filtering and translation methodologies listed in sections 1 to 5 to select objects and groups from source firewalls.

These objects, groups and a special 'any' object can then be used to match all traffic listed in the log files and rules can be built specifying the most specific connection profile permitted by the objects loaded into 360-FAAR.

360-FAAR first matches all connectivity (most specific to least specific for both address and service definitions) found in the log files to the imported object definitions. It then uses the imported group definitions to wrap the connectivity using the group definitions in place of objects, when all objects in the group made similar connections.

The grouping procedure reduces the complexity of individual rules immensely, it also helps to ensure that similar objects are associated together and has a direct affect on the rule building algorithm. Changing the priorities of this algorithm can provide different perspectives on the same traffic.

Two pre-setup versions of this algorithm are available for use in parallel, with each other, providing two perspectives on the rules that can be built to permit the traffic the firewall has seen ,using the objects and groups that were imported.

The results are output as spreadsheets that contain the new rules and the commands to create the needed rulebases. These spreadsheets can be modified and re-uploaded to generate new rule definitions for one or more firewalls.

Detailed analysis of traffic assigned to the special 'any' object is provided and this information is listed in a format that can be uploaded after updating and used to build further rules.

16. Antispoofing Group / Routing Table Cross Referencing

16.1 Reasons to Cross Reference Antispoofing Groups with Routing Tables

Anti-spoofing functionality is an essential part of any firewall policy.

Despite this, it is all to common to find that anti-spoofing has been miss-configured, or disabled on many interfaces of production firewalls.

Re-enabling a firewalls anti-spoofing functionality, without proper visibility of where the objects that you are expecting to see, need to be routed to and from, is a hazardous task.

Without this information it is all to easy to block large sections of a customers network from communicating with their most vital services

16.2 Restrictions of Current Methodologies and Tools

Most modern policy editors have a network map of some sort, however these maps are usually built from the information detailed in a firewalls anti-spoofing configuration.

The map is of little use to you if you are trying to rebuild the anti-spoofing groups or zones that it is relies on to build the networks structure (catch 22).

Manually locating routes for each of the network and firewall objects configured is a laborious and error prone task.

It often leads to smaller networks and groups being missed and assigned the same anti-spoofing or zoning properties as one of their super-nets or super-groups.

Log analysis (if it is unrelated to a firewalls configuration) is of little use when resolving anti-spoofing or zoning issues. This is because the anti-spoofing function of a firewall will only log in situations when a packet has been dropped so is only visible via policy analysis.

Systems outside of the firewall should never see the firewalls anti-spoofing functionality unless they are spoofing their IP's or the firewall is badly configured.

16.3 Solution Provided By 360-FAAR

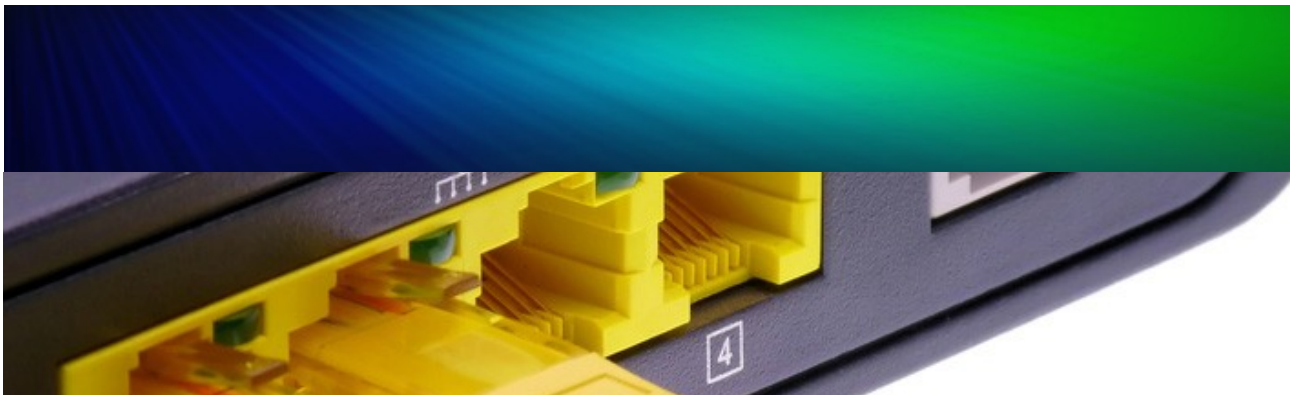

360-FAAR's policy engine has a route analysis subsystem built in.

It relies on the same functionality that the zone and area assignment algorithms use.

These algorithms are able to construct routing or zoning tables from a firewalls existing objects, routes, zones and interface configurations.

Routing tables of almost any style can be loaded into 360-FAAR for analysis, such as:

- Unix/BSD/Linux style netstats,
- Cisco -show route,
- MS -route print,
- Netscreen and Fortigate -get route),
- Object definitions (with zones) can be used to assign extra routes to zone information
- Custom routing tables formats can also be read!

On firewalls running dynamic routing algorithms, it is possible to load an exported routing table to catch dynamic routes that are configured buy the routing algorighm in use on the firewall and not the firewalls config its self

.

The routing / zoning algorithms create a CIDR maps of the routing table, the surrounding routers and the local interfaces. All object definitions are then matched to the route map in the same way a router would match network traffic during forwarding (most specific match first).

The updated object definitions and the commands to create and modify the firewalls automatically are output in modifiable spreadsheets than can be reprocessed after updating.

17. Custom Analysis and Rebuild Projects Using 360-FAAR

17.1 Situations Requiring Custom Analysis Procedures

In very large networks, many of the tools described in this document will need to be deployed in concert with each other in order to provide truly useful information regarding rule usage throughout a complete infrastructure.

The order of processing and the priorities and algorithms used to process rules, logs, objects and groups can produce differing rulebase so 360-FAAR can be standardised between implementations and service runs to ensure the consistency of the rules output.

To ensure users of 360-FAAR within an organisation adhere to the conventions agreed, the user interface can be updated to remove the algorithm and priority options to prevent them being changed.

Larger, highly structured procedures can be created by sequentially linking various parts of the 360-FAAR policy engine together in a process analogous to the way symbols are used in algebra.

17.2 Situations Requiring Custom Analysis for Specific Areas of a Firewalls Implementation

The 360-FAAR policy engine is a highly versatile tool. The code for it is well defined, very general and highly malleable. These features mean that it lends it self easily to many other uses, most of which we probably haven't even thought of.

If you have an analysis problem that you need solved please call 360 Analytics Ltd. and talk your problem through with us. It may be that we can solve your problem with a well defined procedural approach but if not, our data structures and sub routines are defined in such a way as it is easy to integrate new algorithms into almost any point of the existing process.

In terms of network, firewall, object and rule manipulation we really can find a way do almost anything you need!

18. Automating 360-FAAR to Implement Secure Dynamic Security Policies Within Your Infrastructure

18.1 Full Architectural Functionality!! (A Future Project)

360-FAAR can be setup to safely automate the removal of unused connectivity from rulebases periodically from within your network.

It can also provide optimisation services or any of the other types of analysis documented here and will output spreadsheets for review before automatically applying the policy changes to the relevant firewalls

18.2 Secure Dynamic Security (SDS) Policies Within Your Infrastructure

360-FAAR can implement SDS Policies automatically within a network infrastructure.

SDS policies are held by 360-FAAR and are not installed on the firewalls. SDS policies apply last usage timers to subsets of connectivity specified on the firewall. When the time-out is reached the connectivity is removed from the firewall.

18.3 Creating SDS Policies

Any of the filtering methodologies listed in this document can be used to create subsets of a policy.

Alternately a single SDS policy can be specified for all rules in all firewalls. Such an SDS policy could be used to time out rules after two years and provide documentation to be signed off (and possibly modified) before the changes are automatically applied.

18.4 Infrastructure Possibilities

360-FAAR can be implemented in a multitude of different ways within a networks infrastructure.

Hopefully some of the functionality described in this document will have interested you and it is highly likely that this could be automated safely. Please contact info@360analytics.co.uk for more information.

19. Has a Network Audit Identified Rules that Violate Your Security Policy?

Has a network audit identified many rules in your rulebase, that you do not know if you need, that violate your security policy?

19.1 The Fastest Methodology to Resolve Security Policy Violations

The fastest methodology to remove these problems without disrupting the network services you are providing will be to remove all connectivity you do not need – see point 1-5, after this confirm your security policy is met – see point 6 and 7 - and repeat the process.

20. Firewall Documentation

20.1 Automatically Generate Firewall Documentation

360-FAAR currently prints its rule definitions in spreadsheet format. These can be configured to be e-mailed to operations e-mail addresses or consolidated within 360-FAAR its self.

360-FAAR can also update Tiki systems automatically with safe subsets of a firewalls configuration. The safe subsets can be created using any of the filtering possibilities available within 360-FAAR. This ensures that only the configuration information you want to store automatically is exported.

Your firewall documentation is kept up to date and various subsections can be recorded at different levels of access within the tiki.

End of 360-FAAR Scenarios For Use

Contact and Company Details

360 Analytics Ltd.

LUTIDINE HOUSE
NEWARK LANE
RIPLEY, SURREY
UNITED KINGDOM
GU23 6BS
TEL: +447960 028 070
Company No. 07533060

- For General Information in the UK
  Please Visit: www.360analytics.co.uk
- For Further info please visit the blog: 36zeroanalytics.wordpress.com

- For General Queries Please Contact: info@360analytics.co.uk
- For Sales Requests Please Contact: sales@360analytics.co.uk

- For International Information Visit: www.360-analysis.com
- International Contact: info@360-analysis.com

- For Operations or Project Work Requiring Onsite Support
  Visit Our Sister Company Site: www.36ZeroNetworkAnalytics.com
- Contact 36ZeroNetworkAnalytic Ltd: info@36ZeroNetworkAnalytics.com