An Active Defense Decision Model

Sergio Caltagirone
University of Idaho
scaltagi@acm.org

Abstract

Active defense poses many problems; among those is the issue of what active defense actions to take given a threat against a particular asset. The model presented here allows any organization to develop an active defense policy and accompanying escalation ladder, which are used in conjunction to describe the best plan of action in a threatening situation. By investing in this model, an organization frees itself to pursue active defense as a legitimate security (and protection) tool while limiting the associated risks.

1. Introduction

Security against defeat implies defensive tactics; ability to defeat the enemy means taking the offensive. [4:5]
- Sun Tzu, The Art of War (c. 510 B.C.E.)

There is no question that our systems are vulnerable in too many ways to mention. Almost every day, new security vulnerabilities are discovered and exposed sending users scrambling to find additional methods to protect their systems. However, what does a user or administrator do when their systems come under such a remarkable attack that their only choice is to disrupt service (hoping the attacker does not return) or attempt to stop the attacker permanently? What should they do? These are the questions posed by active defense.

Active defense is, as Sun Tzu so eloquently phrased it, the "ability to defeat the enemy [by] taking the offensive." In other words, the only way to permanently disable the enemy is to attack them. More specifically, active defense is any action performed, in an automated or non-automated fashion, to acceptably mitigate a threat against a particular asset. But how does this strategy apply to a company whose business is rendered lame when a cyber criminal removes their ability to effectively communicate with their customers and supply chain? What of a university whose enrollment and student records databases are constantly being probed for vulnerabilities?

Both of these situations, and in any situation where the risks to an organization are great enough to contemplate an active role to stop the attackers, requires a proper model for organizations to fully understand the implications of their actions. Without a model, organizations that assume a position of active defense do so with incredible risk. A decision model for active defense should allow any organization to create an active defense policy and escalation ladder that allows them to protect their most valuable resources in a timely and, hopefully, satisfactory manner. Additionally, such a model should provide organizations with a sense of confidence that they are free to actively protect themselves while not assuming unacceptable risk.

This paper will present such a model. The paper will first define the stages necessary in a proper analysis of active defense, and identify to which stages the model applies. Secondly, the paper will define the goals and assumptions of the model. Next, the model itself will be defined and described. Lastly, the model will be analyzed with respect to the stated goals and assumptions.

2. A Quick Note on Terms

The semantic difference between the terms active defense and active response will probably be argued for many years; and although they are practically the same, it is still important to note why this paper chooses active defense over active response. The term active defense was chosen because the model includes many ethical aspects and active defense reinforces the idea that any organization contemplating action is only doing so to mitigate risk and not for retribution or vigilante justice. Secondly, active defense strengthens the fact that organizations should only do as much as is necessary to defend their systems.

3. Active Defense

Active defense is, as defined earlier, any action performed, in an automated or non-automated fashion, to acceptably mitigate a threat against a particular asset. Of particular note in the definition is the use of the phrase 'acceptably mitigate'. The means that a threat does not need to be eliminated, but only diminished until a specific protection goal is met. This goal can be anything from 'remove the threat to the asset until a more permanent solution can be found' (such as changing the port that a service is running on, or blocking traffic to the port at the firewall), to 'permanently remove the threat' (by possibly sending out a virus against the attacker).

However, in security, we must be very clear what the term 'permanently remove' means. Although this might be the goal of most active defense policies, it is unrealistic because the attacker could move to another computer, or another attacker could download a script that exploits the same vulnerability. Therefore, the goals of an active defense action must be understood in the context of a very dynamic environment and should be very strictly defined as such.

Additionally, numerous agents can perform an active defense action. These agents could be part of an intrusion detection system [1], or the agents could be system administrators. The decision of which agents perform the actions in which situations is for each organization to define based on their own security policy and organizational structure. In either the case where autonomous or non-autonomous agents perform the action, the same stages for a successful active defense are necessary.

The eight stages of active defense are: planning, detection, evaluation, decision, action, analysis, escalation, and maintenance. These stages can be formally or informally defined by an organization. The model presented later in this work applies to the planning stage of active defense. It will allow an organization to create both the active defense policy and an escalation ladder. Both of these devices will then be referenced in the evaluation, decision, analysis and escalation stages.

3.1. Planning

The most important stage, planning, is done well before any attempt at active defense is made by an organization. An active defense plan includes two parts: an active defense policy and an escalation ladder; both parts are required for the purpose of mitigating the risks acquired by assuming a position of active defense. It should be noted that planning is not a 'necessary' stage in active defense because an angry operator may not have previously planned to launch a virus against an attacker during an angry fit. But, planning should be a necessary step in any formal analysis of active defense.

3.1.1. Active Defense Policy

The active defense policy describes what assets are to be considered for active defense, a complete evaluation of the threats (or classes of threats) that exist against the assets, and the value of the asset with respect to the consequences of a successful attack. Additionally, the policy describes the potential actions that can be taken to mitigate the risk of the threat, as well as the risks assumed by conducting an action.

Not all assets should be considered for an active defense policy. For example a person who downloads a brochure from your website may not be considered a threat requiring an active defense action. An active defense policy should be consistent with an organization's formal security policy that describes valuable assets and the risks associated with those assets.

3.1.2. Escalation Ladder

The escalation ladder describes the series of actions that an organization can consider in the decision stage to mitigate the threat against an asset. Each step up the ladder will also contain more risk than the previous step. For example, if an attacker was discovered to be launching a denial of service attack from computers in the organization, the first action may be to notify the chief information security officer, then the organization may block the port of the firewall through which the attacker is gaining access. The organization may then attempt to track the attacker and ascertain their location, which then will be used by the organization to seize the offending machines and neutralize the attacker.

The escalation ladder for a given threat and asset will contain one or more steps, where each step will be associated with a risk level as defined by the active defense policy and developed through the use of the model.

3.2. Detection

Detection is the automated or non-automated discovery of a past, ongoing, or future threat against an asset. It is beyond the scope of this model, but is well researched in areas such as intrusion detection.

3.3. Evaluation

Evaluation is the process that places a detected threat in the context of the active defense policy. Evaluation does not in any way describe the quality of the policy. The evaluation of a threat will inform the organization of the risks of a successful attack versus the risks of an active defense action.

3.4. Decision

Decision is the determination of which active defense action to perform to mitigate the threat, but which will not expose an organization to more risk than a successful compromise of an asset.

3.5. Action

An active defense action is any automated or non-automated activity performed for the specific purpose of mitigating the threat against an asset. Actions can range from notifying the chief information security officer of a detected threat, to shutting down a port, to the use of a denial of service (DoS) attack against the attacker, to the initiation of a virus against the attacker. A successful action does not imply a decreased risk from the threat.

3.6. Analysis

After an action is performed, an analysis must be made of whether the action has successfully mitigated the threat to the satisfaction of the goals stated in the active defense policy. If the action has not satisfied the goal(s), then escalation is necessary. If the action has satisfied the goal(s), then determination must be made if the attack is ongoing – and whether actions taken need to be kept in place or whether an organization can revert to a state of less risk further down the escalation ladder.

3.7. Escalation

Escalation is the change in state by performing the next action described in the escalation ladder. Escalation also implies an increased risk assumed by the organization. Escalation is usually necessary when the action performed is not successful in mitigating the risk to the satisfaction of the goal(s) described in the active defense policy.

3.8. Maintenance

Maintenance is important to the security of any organization. It is of paramount importance that an organization keep their active defense policy as current as possible to minimize potential unknown risks. This includes adding or removing assets, threats and risks. Additionally, after the analysis and escalation stages of an active defense, the policy should be reviewed to reflect any lessens learned during the post-mortem of the active defense action. Clearly, it is also necessary to update the escalation ladder if the active defense policy changes.

4. Goals and Assumptions

For any model to be successful, its goals and assumptions must be explicit and verifiable. The reason is that the model may not be as powerful as marketed, or may not be relevant to certain organizations. Also, the model may contain assumptions that greatly weaken the model; more importantly is the fact that some assumptions may make implementation of the model unrealistic.

4.1. Goals

Because this proposed model is not purely an academic exercise, but a practical guide for any organization to create an active defense policy and escalation ladder, the goals of the model should mirror that practicality. The goals presented here are in no specific order.

1) **Generalizable**: The model should allow any organization or individual the ability to create an active defense policy and escalation ladder.

2) **Useful**: The model should be practical and useful to any organization contemplating active defense.

3) **Expandable**: The model should allow organizations to include elements that are not included in the model with only limited changes to the model in general.

4) **Mitigates legal risk**: Allows an organization to prove that they have practiced due diligence with respect to active defense in the face of any legal challenge.

5) **Consistent**: Every element in the model should be consistent with every other element in the model.

6) **Thorough**: The model should allow any organization the ability, with the proper time investment, to create a complete assessment of risk and benefit for each potential active defense action.

7) **Automated**: The model should allow explicit analysis and action by automated methods.

4.2. Assumptions

The assumptions in this model are by far its largest weakness. The problem is that the model rests on the principle that the assets and responses are properly and thoroughly evaluated in the active defense policy. Even in the smallest organization, the problem of accurately enumerating and evaluating all of its assets and responses is immense, if not impossible. Additionally, the problem of evaluating the ethical considerations of an active defense action is dependent on an organization's ability to expertly evaluate the ethical issues involved with the action.

The question also arises about whether all of the ethical consequences of an action can be enumerated. It may be that the organization chooses to denial of service the attacker, which actually causes a router far upstream of the attacker to fail, which causes some power plant controller to fail which further causes critical systems in a hospital to fail, whose redundant systems fail and people are harmed. This is a far-fetched example, but illustrates that all of the consequences of an action cannot be

enumerated because of the complexity and interdependence of our information systems.

On the other hand, how can we evaluate legal risks accurately? What of untested laws? Can we trust that the laws will be interpreted the same in our case as it was in others?

This model assumes that all of these questions have been answered by the organization that is developing an active defense policy. It should be noted that these questions do not need to be answered completely, but only to the organization's satisfaction. Therefore, the model will only predict outcomes and best actions based on the relative accuracy of the information provided and risk ranks.

Additionally, the unknowns faced by an organization grow rapidly as they progress up an escalation ladder. Therefore, very few unknowns will exist in the lower levels, such as shutting down a port, but as the action begins to include denial of service attacks, the number of unknowns grows and may become impossible to enumerate and evaluate.

However, the problems implied by these assumptions are lessened with the inclusion of contingency plans in the model. These contingency plans allow an organization the flexibility to dynamically assess the situation to relieve the stress of unknown quantities in the model thereby reducing risk.

1) **Assets can be estimated**: The model assumes that the assets and risks of an organization can be accurately estimated with respect to the given categories.

2) **Responses can be evaluated**: The model assumes that all of the active defense actions to a given threat have been included, and that the model will not be used to evaluate actions that have not been included.

3) **Consequences are enumerable**: The model assumes that all the consequences of an action are known and have been included in the active defense policy.

4) **Ethical considerations can be evaluated**: The model assumes that all ethical considerations have been evaluated correctly to provide their accurate weight.

5) **Legal consequences are known**: The model assumes that all legal consequences are known, and that the laws have been tested and interpretations will be static.

5. Escalation Stages

Of course not all active defense actions are created equal. Tracking down an attacker with common tools such as ping and finger is not the same as sending them a virus. Therefore there are implicit stages of active defense where each stage, naturally, contains more risk than the previous stage(s). It is important to identify the stages of active defense actions because as the model is concerned with an organization assuming liability, taking action should begin at the lowest stages and progress upward until the protection goal is met.

Additionally, a logical and measured progression through the stages can protect an organization legally by showing due diligence was practiced and the attack was not ad hoc.

The stages of active defense are [2]:

1) Internal Notification
   a) Using the an organizational structure to notify the appropriate persons of an active defense situation

2) Internal Response
   a) Applying active defense actions within an organization's boundaries (shutting down the port on a firewall)

3) External Cooperative Response
   a) Employing the assistance of other entities outside of an organization to mitigate a threat

4) Non-cooperative Intelligence Gathering
   a) Using external services (finger, nmap, netstat) to gather intelligence on the attacker

5) Non-cooperative 'Cease and Desist'

6) Retribution or Counter-strike

7) Preemptive Defense

[FINISH DESCRIPTION OF STAGES]

[A LOT MORE NEEDS TO BE SAID HERE TO DEFEND THE CLASSIFICATIONS AS BEING NECESSARY AND SUFFICIENT]

6. Model

The model's purpose is to provide an organization with an algorithmic method of developing an active defense

policy and an escalation ladder. The model is separated into three stages, asset evaluation, action evaluation, and the state machine creation. The asset and action evaluation stages represent the active defense policy, while the state machine represents the escalation ladder.

6.1. About Risk Assessment

This model requires that the organization invest heavily into risk analysis so that each asset can be accurately weighted. However, this model is not to serve as a business model for risk analysis or assessment and does not claim to be a risk assessment tool. The model simply uses an organization's existing risk assessment tools and documents and places that information into a form that is useful for understanding and comparing the complexities of active defense. It is suggested that an organization adopting or analyzing this model first adjust the model so that their existing risk analysis is compatible.

There has been some very good research done in the field of risk assessment, particularly in applying models to information assurance, such as [3].

[INCLUDE A FURTHER DISCUSSION OF CURRENT RISK ANALYSIS RESEARCH]

6.2. Asset Evaluation

The first stage in the creation of an active defense policy is asset evaluation. In this stage, an organization identifies which assets, if threatened, are candidates for an active defense action. Additionally, the threats against each identified asset are enumerated. More importantly in this stage, is that all of the risks to an organization are listed for each threat, and each risk is valuated. This helps to quantify an organization's exposure to risk if the identified threat materializes and is successful. Later this will be used to decide if the risks of an active defense action outweigh those of the loss of an asset to the threat.

6.2.1. Asset Identification

The first step in asset evaluation is the identification of the valuable assets of an organization (each asset should be given its own unique identifier). Assets can be as specific as a particular object in a system (such as a process or file), or as general as a service or area (such as internet connectivity or network behind xxx firewall). What is important is that the asset be valuable enough to be under consideration for an active defense action if threatened. The granularity of the assets is left to the organization based on their needs and investment in the development of an active defense policy. Most of the information about valuable assets should come from the organization's previously developed general security policy.

For example, a university may identify its valuable assets as: Internet connectivity, student records database, administration servers, certain sensitive files (or a file storage area where such files are maintained), etc.

6.2.2. Threat Identification

After identification of all valuable assets, the threats to each asset are enumerated under the classical categories of confidentiality, integrity, and availability. Just as with the asset identification, each threat needs to be given its own unique identifier. The threats identified can be as general or as specific as necessary to satisfy the organization. The threats can be as specific as 'an attacker probes port 25 and 26 in order during non-operational hours', or it can be as general as 'a probe of network ports is detected.'

Therefore, threats can be listed as very specific, or as categorizations of threats. Again, the granularity of the threat identification will depend on the needs of the organization and their investment in the active defense policy. But it is recommended that the threats be specific enough to detect and analyze easily, but general enough that new attacks could be placed into a categorization of threat; the more specific the threats, the more difficult it is to categorize new attacks as previously identified threats.

6.2.3. Scoring Chart

Because of the nature of this model, the risk of a threat materializing must be comparable to the risk of an active defense action. Therefore, an organization must have a reasonable method of scoring the risks. There are four threat-risk categories (although additional categories can be easily added), they are: legal, national security, financial, and ethical consequences.

The chart ranges in scores from 10 to –10, where the greater the number represents a higher cost of risk. For example, a financial risk score of 3 may be a loss to the organization of $100,000, while a score of 5 may be a loss of $400,000, and a score of –3 will be a gain of $100,000. The reason for a scale of negative as well as positive is that some active defense actions actually have a positive impact on an organization – such as saving lives. Since we wish to model both minimizing risk, and maximizing benefit, this structure seems to work best.

It must be stressed that each organization will develop the chart according to their own standards of how they

calculate loss. And accordingly, the chart must be maintained to be consistent with the realities of the corporation. For example, a Fortune 500 corporation may have a financial risk score of 5 represent one billion dollars, while a start-up's 5 may represent $10,000. Additionally, each category must be translated into the context of the organization. The CIA would calculate national security risk differently from a national laboratory. An organization may also not have any entries in a category.

The scoring is not limited to integers; an organization is free to use the entire real number domain between 10 and -10. By doing this, a company may use 4.5 to represent $4.5 million lost. The precision that an organization chooses does not have any ramifications on any other part of the model, therefore each category could have its own precision.

The method of developing a scoring chart is simple. All that is required is that at least the integers from 10 to –10 be defined as representing a real loss or gain to the organization, the scores must remain within 10 and –10, and 10 represents the greatest lost, –10 represents the greatest gain, and 0 represents no loss or gain. For example, given a university, the chart may look something like this:

Legal	
10	Criminal charges levied, long jail sentences for employees
9	
8	
7	
6	
5	
4	
3	
2	
1	
0	
-1	
-2	
National Security	
10	
9	
8	
7	
6	
5	
4	
3	
2	
1	
0	
-1	
-2	
Financial	
10	Loss of $1 billion
9	Loss of $500 million
8	Loss of $100 million
7	Loss of $50 million
6	Loss of $25 million
5	Loss of $1 million
4	Loss of $500,000
3	Loss of $100,000
2	Loss of $50,000
1	Loss of $10,000
0	No loss or gain
-1	Gain of $10,000
…	…
-10	Gain of $1 billion
Ethical Consequences	
10	

Table 1. Example Scoring Chart for a University

[COMPLETE THE CHART]

6.2.4. Risk Identification

After each threat has been identified, then it is necessary to calculate the risk associated with the threat. For each threat, the organization should list all possible risks (in the legal, national security, and financial categories) that an organization. Each risk must then be scored.

To calculate the score of each risk requires two steps. The first step is to assign a probability, between 0 and 1, that the risk will manifest itself. The second is to locate a score on the scoring chart that represents the total cost to the organization (over all time).

For example, a university may determine that if the threat is successful, it will result in the loss of some enrollment (financial risk). The probability of this is risk manifesting itself is 0.3. It then determines that the total cost (of lost enrollment dollars) would be approximately $100,000 – which corresponds to a score of 3 (in the financial category). The score of that risk is then 0.3 * 3 = 0.9.

The total score of the threat is simply the sum of the scores of all of the risks associated with that threat. Once this is completed, the organization should have a thorough

chart of assets, the threats against that asset, and the risk of the threat manifesting itself. This chart will later be used to determine the actions that are reasonable to levy against a threat. The next step is to determine the actions themselves, along with their risks.

6.2.5. Example Asset Evaluation

Our example will concern the fictitious Foo Bar University (FBU), and its database of student records. The records include student health information, academic records, employment records (with the university), financial aid data, and other sensitive data such as social security numbers. We will assume that the database is connected to the larger university network, as well as the Internet through the university network but is partitioned by the admin firewall as well as the campus firewall (see Figure 1. Network topology of Foo Bar University). The student records database contains information that is required to be kept confidential by law, and the information within the database is not kept encrypted but access is regulated through the use of usernames and passwords. An intrusion detection system and extensive logging also monitor activity on the database.

The information in the database is also very sensitive to integrity threats because it is the University's only record of student information that is necessary for correct reporting of diploma status, student enrollment verification, student health insurance, etc.

Here, FBU declares the student records database as an asset valuable enough for active defense consideration. It then lists the specific threats or categories of threats (for confidentiality, integrity, and availability) to the asset and the associated consequences of a successful attack. These risks are then associated with risk levels in the previously defined scoring chart (Table 1. Example Scoring Chart for a University). The resulting report is shown.

Asset (A1): Student Records Database			
Confidentiality Threats			
Threat (TC-1): Outsider gains access and copies sensitive data			**FINAL SCORE:**
Legal Risks	*Score*	*Probability*	*Score * Prob*
L1:			
National Security Risks			
NS1: Students' social security number are released			
Financial Risks			
F1: Loss of tuition	5	.8	4
F2: Loss of financial donations	7	.4	2.8
Ethical Consequences			
EC1:			
Integrity Threats			
Threat (TI-1): Outsider changes sensitive data within the database			**FINAL SCORE:**
Legal Risks	*Score*	*Probability*	*Score * Prob*
L1:			
National Security Risks			
NS1: Students' social security number are released			
Financial Risks			
F1: Loss of tuition	5	.8	4
F2: Loss of financial donations	7	.4	2.8
Ethical Consequences			
EC1:			
Availability Threats			
Threat (TA-1): DoS attack on database			**FINAL SCORE:**
Legal Risks	*Score*	*Probability*	*Score * Prob*
National Security Risks			
NS1: Federal agencies not able to track foreign students			
Financial Risks			
F1: Loss of employee productivity	4	1	4
F2: Loss of financial donations	4	.2	.8
F3: Increases security staff requirements	2	1	2
Ethical Consequences			
EC1:			

Table 2. Example Asset Evaluation

[COMPLETE THE CHART]

6.3. Action Evaluation

Action evaluation is the next, and final, step in the development of an active defense policy. In this step, an organization identifies all of the potential actions it can perform to mitigate threats and the risks associated with those actions. At the end of this step, an organization

should have created an active defense action chart, which will be used to develop the escalation ladder.

6.3.1. Goal Identification

Because we are not using active defense as a retributive or vigilante tool, but to defend our systems and valuable resources, we must understand when to stop attacking. This is important for a number of reasons, (1) it prevents an organization from accidentally assuming more risk than necessary, (2) allows the organization to prove in court that they did only what was necessary to achieve an appropriate protection goal, and (3) helps guide the development of a response to a threat by providing a threshold.

The goals for each threat are going to be different depending on the organization and their needs. For example, a national security organization may have a goal to prevent any future threat from that particular assailant, while a business may only be concerned with halting the current threat. The level of goal will be dependent on an organization's available resources and their protection needs.

Therefore, for each threat, a clear and unambiguous goal must be declared which will guide the responses to the threat. These goals must also be approved by the management in the organization responsible for assuming the risk if anything goes wrong while executing the active defense actions.

6.3.2. Action Identification

After the goals have been properly defined, then an organization must identify every possible action that can be performed to mitigate a threat against a particular asset to obtain the goal. This is easily done by listing, for each threat identified in the asset evaluation, all of the actions that can be performed. Additionally, actions must include organizational requirements, such as notifying the proper higher-ups, filing a report, etc.

For example, given threat TA-1 (in Table 2. Example Asset Evaluation), possible actions are: notify CISO (chief information security officer) of the threat, notify upstream provider of the attack, notify FBI of the attack, using a diagnostics tool (such as nmap or netstat) obtain the location of the attacker, DoS the attacker themselves, send a legal 'cease and desist' letter to the attacker, send a virus to the attacker's machine, notify the attacker's ISP, divert all traffic to a dummy database on a honeypot, etc. Of course, an organization is only required to list the actions that are available given their resources.

6.3.3. Action Classification

The next step is to classify the actions identified in the previous step according to which stage of active defense they are associated with as described in section 5. For example, notifying the FBI would belong to stage 3 and using nmap would belong to stage 4.

6.3.4. Additions to the Scoring Chart

To be able to successfully compare the risks of an attack being successful to the risks involved with each active defense action, the same scoring method must be used. However, a stronger ethical component must be present determining the score of an action because if the organization actively defending themselves accidentally causes more damage than the attacker or even acts more unethically, then the question is: who is worse?

Performing an active defense action clearly places ethical risks on an organization. And although the ethical risks may be out-shadowed by the other categories in the minds of most, the attempt to maintain an ethical organization cannot be dismissed.

Also, choosing between a consequentialist (only the consequences of an action are deemed necessary for ethical consideration) and a deontological (only the act in and of itself is considered for ethical consideration) ethical theory is almost impossible for an acceptable application of ethics to active defense because there are serious questions that abound on both sides. Therefore, it is necessary to represent both, where the consequentialist is represented by the Ethical Consequences category which defines the ethicalness of the potential consequences of an active defense action; and the deontological is represented by the Ethical Action category which describes the ethicalness of the action an organization takes in and of itself.

Therefore, using the previously defined scoring method the two categories will also be scored from 10 to –10, where 10 is the most unethical, and –10 is the most ethical. For example, killing a person would be considered 10, while saving a life would be considered –10 in the consequences category; while sending out a very destructive and uncontrollable virus would could be a 10 and releasing a patch to a virus could be a –10.

It is easily said that this scoring is the most subjective and difficult for an organization. However, to make things easier, in the Ethical Actions category, the only actions that need to be considered are those that are potential active defense actions. While in the Ethical Consequences category, all potential consequences need to be considered. An example of a scoring chart for these two categories is given below.

Ethical Consequences	
10	Killing a Person
9	
8	
7	
6	
…	…
0	Not an ethical consideration
-1	
-2	
…	…
-10	Saving a Life
Ethical Actions	
10	Uncontrollable Virus or Worm
9	
8	
7	
…	…
0	Not an ethical consideration
-1	
-2	
…	…
-10	Release Patch for Virus or Worm

Table 3. Example of Ethical Scoring Chart

[COMPLETE CHART]

6.3.5. Risk Identification

The method of identifying the potential risks of an active defense action is almost identical to identifying risks of threats as previously defined. For each action, all of the risks must be identified in the categories of Legal, National Security, Financial, Ethical Consequences, and Ethical Actions. Additional categories may be added by an organization if necessary.

For each risk, the probability of it occurring must be defined, as well as the risk's score (as defined in the scoring charts). These are then multiplied to calculate the total score of the risk. All of the risk scores are then added together to obtain the score of the action.

6.3.6. Utility Modifiers

Because each organization has its own unique goals, a modifier can be placed on a specific category to provide more weight based on the utility of that goal to the organization. This comes from the idea of a utility function developed by Keeney and Raiffa in [4].

If, for example, a national security organization was concerned with the national security implications of an action rather than any others, then it could place a modifier on the national security category to give it more weight in the escalation ladder.

Also, a utility modifier can be placed on each category if an organization wishes, or place it on none if they wish the same weight for all categories. If an organization wishes to act particularly ethically with regard to active defense, then it may place the same modifier on both the Ethical Consequences and Ethical Actions category, but place no modifiers on any other category.

To use the modifier, an organization simply multiplies each risk's score in that category with the corresponding modifier. For another example, we may multiply every National Security risk score by 1.2 while we multiply every Ethical Action score by 1.3. This would place a 10% greater weight on Ethical Action than on National Security, and a 20% greater weight on National Security over all other categories.

6.3.7. Success Ordering

One of the most important steps in the action evaluation is the action ordering. It is also one of the simplest. All that is required is that for each threat, and within each stage, partially order the nodes such that that if $n1 < n2$, then n1 has a lower probability of successfully mitigating the threat based on the defined goal. In other words, n1 has the least chance of meeting our goal, and n2 has the greatest chance. Nodes can also be numbered the same, but it is best if all the nodes are numbered as integers.

For example, if in threat TA-1, stage 3 had three actions, and those actions were: ask the FBI to find the attacker, ask your ISP to drop all offending packets, ask the attacker's ISP to cut off the attacker's service; then the ordering would be such: FBI (1), Our ISP (2), and Attacker's ISP (3). This is because shutting down the attacker at the source ISP would have a greater likelihood of success than handling it from our ISP or through the FBI – if our goal was an immediate removal of the threat, if the goal was more permanent, then the FBI would probably be higher.

6.3.8. Example Action Evaluation

Using the previous scoring chart (Table 1. Example Scoring Chart for a University), and asset evaluation (Table 2. Example Asset Evaluation), an example of Foo Bar University's action evaluation is given. As actions depend on the environment, the network topology of Foo Bar University is also given for easy reference.

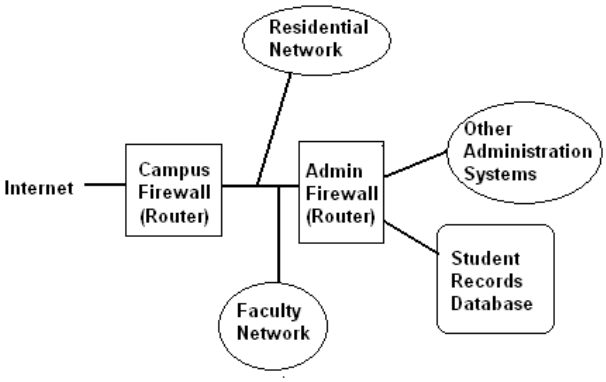

Figure 1. Network topology of Foo Bar University

Threat: TA-1	**Goal:** Stop the ongoing DoS attack while preserving access to the database behind the campus firewall		
Stage 1 Actions			
Act1:	**Risk Score:**		**Success Order:**
Legal	*Score*	*Prob*	*S * P*
National Security			
Financial			
E. Consequences			
E. Actions			
Stage 2 Actions			

Table 4. Example Action Evaluation

[COMPLETE THE TABLE]

7. Escalation Ladder

So far, this paper has presented the first two parts of the model, asset evaluation, and action evaluation. By doing this, an organization now has two tools with which to scientifically analyze their risks with respect to active defense. This has answered the question: what risks are involved for my organization if an active defense policy is initiated. The question still left to answer is: if faced with a threat against an asset, how does my active defense policy describe what an organization should do?

The escalation ladder answers these questions of how to proceed and what actions to perform. It does this by representing the threat, and all possible actions to mitigate the threat in a graph representation. Then, we use a standard graph algorithm to find the path through the graph with the least risk (i.e. shortest path). In the end, the algorithm will give us a path through the graph that will tell us the best actions to perform to minimize our risk and maximize the successfulness of our active defense.

7.1. Graph Representation

The graph representation for the escalation ladder is straightforward. It is a weighted, directed acyclic graph with potentially negative edge weights.

1) Every threat is a separate graph
2) Each vertex in the graph represents an action.
3) The weight of each vertex shall be: Risk_Score(Action) – Risk_Score(Threat) – Success_Order(Action)
4) The graph is only feed forward, once an action is performed, there is no edge backwards.
5) The start vertex u will always be a non-action and have no value.
6) All vertices in the last stage will feed into the final stage, non-valued vertex v.
7) Each vertex will feed into each vertex in the next stage.

Additionally, if an organization defines in their active defense policy that one action must always precede another, then that creates a single path through the graph. This could be found when the notification of some person must always precede a stage x action. The organizational structure or laws will define these rules.

7.2. Shortest Path

After the graph has been created, then it is simply a matter of traversing the graph such that we get from the start vertex u to the end vertex v with the minimal cost using the well-defined shortest path algorithm.

7.3. Contingency Plan

But what if we cannot complete an action, or for some other reason, an action is not available that is on the chosen path? Then a contingency plan is created. The algorithm simply backs up to the last successful step and then selects the next least-cost path and continues to do this until it succeeds in finding an action that can be completed.

7.4. Example Escalation Ladder

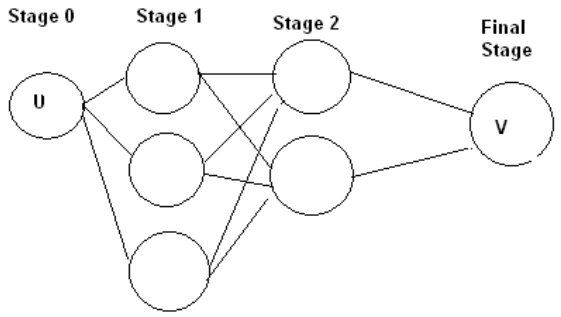

7.5. Evaluation

Before each action is taken, an evaluation must be made to ascertain whether the previous action has been successful in mitigating the risk to the satisfaction of the protection goal. If the protection goal has been met, then naturally, the next active defense action is not taken. If the goal has not been met, then the algorithm continues to execute active defense actions.

[DISCUSS TOTAL RISK]

8. Analysis

At this point it is necessary to look at the model objectively and to determine whether it has satisfied the goals stated in section 4.1 To do this, each goal will be examined in turn.

8) **Generalizable**: The model should allow any organization or individual the ability to create an active defense policy and escalation ladder.

The model is generalizable because it does not discriminate towards any particular organization and can also be used by individuals. Since the power of the model rests with the organization's ability to identify their threats and assets, as well as the ability to rank their risks (which are relative to the organization), the organization holds complete control over the model and is therefore fit for any organization.

9) **Useful**: The model should be practical and useful to any organization contemplating active defense.

This goal can only be shown to be true when organizations actually attempt to adopt the model. However, the extensive examples in the paper should give a proper empirical argument for the usability of the model.

10) **Expandable**: The model should allow organizations to include elements that are not included in the model with only limited changes to the model in general.

Since the organization that is developing the active defense policy can determine the categories, assets, threats, risk charts, and all other aspects of the model, the model can be expanded as large as necessary to accommodate any organization necessary.

11) **Mitigates legal risk**: Allows an organization to prove that they have practiced due diligence with respect to active defense in the face of any legal challenge.

[INSERT ARGUMENT]

12) **Consistent**: Every element in the model should be consistent with every other element in the model.

[INSERT ARGUMENT]

13) **Thorough**: The model should allow any organization the ability, with the proper time investment, to create a complete assessment of risk and benefit for each potential active defense action.

Since the model requires that the organization fully enumerate all of their assets and risks that will be covered by the active defense policy, none should be missing and the model is complete. However, if the organization missed any assets or risks, then the model becomes only as good as the organization's approximation. And since the model has no control over the organization's ability to enumerate these necessities – then from the model's perspective, it is as thorough as it possibly can become.

14) **Automated**: The model should allow explicit analysis and action by automated methods.

The model was designed with this goal in mind. It can easily be implemented in a contemporary intrusion detection system because it is only a series of tables used to create a graph which autonomous agents can analyze easily using well-known algorithms.

9. Conclusion

Admittedly, the creation of the active defense policy and escalation ladder presented requires a tremendous resource commitment on the part of any organization. However, the questions regarding what one should do in an active defense situation are astounding and require such a commitment to explore the real ramifications of an active defensive position.

As it has been shown, this model can be used by any organization to explore an active defense position with regard to their assets, threats and protection goals. The model is generalizable and expandable to be useful to any organization. It also helps to alleviate some of the legal risk involved as well as provides for a way to incorporate their active defense policy into existing protection technology such as an intrusion detection system or firewall. With these goals met, the model should be a great tool for any organization to explore the boundaries of information assurance and to take the war to our enemies.

10. References

[1] D. Frincke and E. Wilhite, "Distributed Network Defense," presented at IEEE Workshop on Information Assurance and Security, West Point, NY, 2001.

[2] B. Endicott-Popovsky, "Active Defenses to Cyber Attacks," A. Group, Ed., 2003.

[3] D. L. Drake and K. L. Morse, "The Security-Specific Eight Stage Risk Assessment Methodology," presented at 17th NIST-NCSC National Computer Security Conference, 1994.

[4] R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives*. Cambridge, Massachusetts: Cambridge University Press, 1976.