



Fundamentals of Capabilities-based Attack Tree Analysis

Copyright © 2005 Amenaza Technologies Limited – All Rights Reserved

Amenaza[®], Secur//Tree[®], as well as the  and  Secur//Tree symbols are registered trademarks of Amenaza Technologies Limited

Amenaza Technologies Limited
Suite 550, 1000 8th Ave SW
Calgary, Alberta
Canada T2P 3M7

1-888-949-9797 toll free US & Canada
+01 403 263 7737 International

E-mail: info@amenaza.com
Web: www.amenaza.com

Table of Contents

Fundamentals of Capabilities-based Attack Tree Analysis

Background	1
Hostile Risk Theory	3
Risk	3
Models	4
Prerequisites of an Attack	4
The Origins of Attack Trees	5
Attack Tree Vulnerability Models	5
A Sample Attack Tree	7
Attack Scenarios	8
Behavioral Indicators – Hardness Metrics	9
Impact Indicators – Logical Effects	10
Predicting Attacks	12
Adversarial Capabilities	12
Identifying Probable Attacks	13
Combining Impact with Incident Probability Yields Risk	15
Estimating Motivation	16
Modeling Countermeasures	17
The Need for Analysis Tools	17
Attack Tree Analysis versus Traditional Risk Analysis Methodologies	17

List of Figures

Figure 1 – Four Quadrants of Risk	3
Figure 2 – Quadrant II	3
Figure 3 – Goal Oriented Tree	6
Figure 4 – Approaches to Burglarizing a House	7
Figure 5 – Attack Scenario Example	8
Figure 6 – Behavioral Cost for Garage Attack Scenario	10
Figure 7 – Garage Attack Damage Cost	11

Fundamentals of Capabilities-based Attack Tree Analysis

Background

Each and every day we all make decisions that involve the assessment of risk. In fact, it is hard to think of any choice that doesn't involve some risk. Should I have the tuna sandwich for lunch or the yummy cheeseburger? The fish might be more prone to spoilage if not handled correctly, but the burger has a higher fat content that, over time, may harm my heart. Price is also a factor.

Most of our risk assessment decisions are informal. Through experience we intuitively come to understand which choices result in more risk than we are willing to accept and take steps to avoid, reduce or share it. In general, humans are fairly adept at correctly estimating the risks associated with events that are within their experience. There are exceptions to this rule – people are notorious for overestimating the risk of spectacular catastrophes (such as airplane crashes) while underestimating mundane events (the drive to the airport).

It is now common to use statistics instead of intuition as a more reliable predictor of events. Although no single individual may possess the necessary experience to predict the frequency of a particular event, society as a whole may. Historical, event-related data, collected from a broad range of sources, forms a body of knowledge representing society's collective experience. The accumulated information can be used effectively to make useful predictions. For example, statistics on the frequency and duration of power failures make it possible to design backup power systems that will handle most outages without wastefully purchasing excessive capacity – all without an in-depth knowledge of the electrical grid.

Unfortunately, using past events as a basis for risk decisions is not adequate in all situations. This is particularly true when considering the risks associated with deliberate, hostile attacks. The 9/11 terrorist attacks on the World Trade Center were unprecedented. The only previous incident involving an aircraft of significant size flying into a skyscraper in New York City occurred in 1949 when a bomber, lost in the fog, collided with the Empire State Building. After the first World Trade Center tower was hit on September 11th, one might have concluded (based on statistics) that skyscraper-aircraft collisions in New York City occur about once every fifty years – yet the second Trade Center tower was struck only seventeen minutes later. Statistics are poor at predicting human behavior (which can change abruptly).

Finding a systematic, disciplined way of understanding the risks from deliberate, malicious activity has never been more important. The lever of modern technology allows an individual, or small group of individuals to inflict damage wildly disproportionate to their numbers. The World Trade Center terrorists created human guided missiles that killed thousands of people and destroyed a \$40B skyscraper that took thousands of person years to build – all on a budget of less than \$500,000. When the ongoing cost of the US-led retaliation is included in the calculation, the terrorists cost their victim over \$200B, or a 40,000 times return on investment (ROI)!

At the more mundane level, adolescent computer hackers regularly, with a few hours of

work and almost zero cash outlay, create viruses that cost businesses tens of thousands of hours (and millions of dollars) of repair work. These examples demonstrate that people of modest means are now capable of inflicting damage that would previously only have been attainable by nation states. This situation is unique to our era.

It is essential to find a way to manage hostile threats intelligently. The resources available for defenses are finite. Although it is possible to implement protective measures against almost anything, it is not practical to protect everything. If something bad happens (and we didn't prepare for it), how do we show that our preparations were not unreasonably (and miserly) frugal? If the bad things that were predicted do not come to pass, how do we demonstrate that it is due to our excellent preparations, and that we did not squander resources through paranoia? Answering these questions is what risk analysis is all about.

The balance of this paper discusses an intelligent, methodical approach for assessing hostile risk. Before proceeding we invite you to review the definitions given in *Appendix A* as they will be used hereafter.

Hostile Risk Theory

Risk

Although the concept of risk exists in many disciplines, it is always based on the same premise

$$\text{Risk} \equiv \text{Incident Probability} \times \text{Incident Impact}$$

That is, *risk* is a combination of the likelihood that an event will occur combined with the amount of damage that will result if it does.

It is possible to plot the range of these two variables on a Cartesian plane (see **Figure 1**). This results in four quadrants, each with certain characteristics.

Quadrant I represents events that are highly likely to occur and of high impact. The need for risk mitigation in this quadrant is usually obvious. In fact, if the system under consideration exists and is in operation, then the mitigation must already have been performed. If not, then the system would already have suffered serious or fatal damages.

Quadrant IV involves incidents that occur frequently, but cause no serious damage. If allowed to recur enough times the cumulative damage might be significant. In most cases, however, the nuisance becomes too great to tolerate and somebody does something to stop it.

Quadrant III deals with incidents that are of only theoretical concern. These incidents occur rarely and cause little pain if they do. Action is seldom required.

The good news is that dealing with ¾ of the risk plane is, if not easy, then at least straightforward. Deciding what to do for problems involving quadrant II is more difficult.

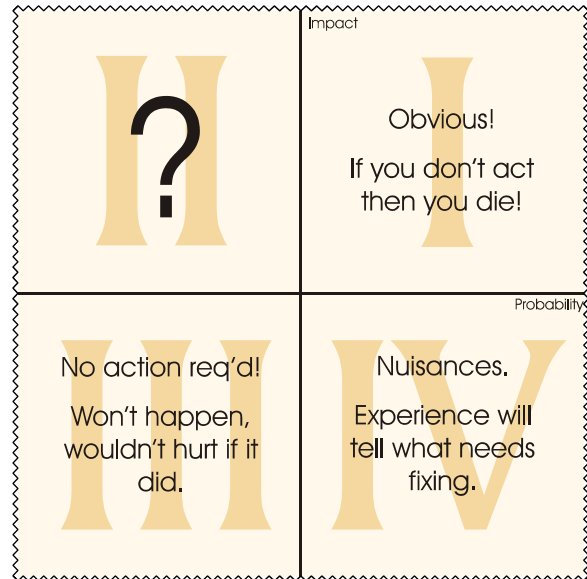


Figure 1 – Four Quadrants of Risk




Figure 2 – Quadrant II

Quadrant II (**Figure 2**) can be thought of as being divided into two regions. On the lower left side are the events that, while ostensibly fatal, are so rare as to be of no practical concern. The upper right half of the quadrant is not so easily dismissed. In fact, psychologists tell us that the events it represents are those that humans are poorest at judging.

People tend to underestimate the risks associated with dangerous activities that can be performed multiple times with no apparent ill effect. Cigarette smoking is a prime example. Intuition incorrectly suggests that since the user has smoked cigarettes numerous times

without ill effect, that there is no reason to expect any different results in the future. The failure by many people to wear seat belts is another example. Intuitive responses are often inappropriate for quadrant II situations.

Applying statistics to quadrant II type events is also difficult. The infrequency of quadrant II events means that long sampling periods are required to create a representative sample. However, accumulating data over a long period of time is problematic since real world conditions may change over the sampling period.

While the analysis technique we are about to demonstrate works in all quadrants, it is most valuable in aiding understanding in quadrant II. It can provide guidance about which types of events fall in each region.

At least in the case of hostile risk, it is more difficult to calculate the probability of an incident than to estimate the impact. Impact can be estimated by considering the immediate damages caused by the hypothetical attack as well as the consequential effects on higher level operations (often called business impact). Business impact is usually found by identifying who will be negatively affected by an event and asking them to describe the costs associated with it.

Models

The world is a very complicated place. Events are influenced by an almost unlimited number of factors. Trying to consider all of these influences is an insurmountable task. In many cases most of a system's behavior can be understood by examining a small set of important drivers. We call this simplified worldview a *model*. Models can be very helpful in understanding the real world but it should never be forgotten that their simplifications are valid under many, but not all, circumstances.

Prerequisites of an Attack

Three conditions must be present in order for an attacker (*threat agent*) to carry out an attack against a defender's system.

1. The defender must have **vulnerabilities** or weaknesses in their system.
2. The **threat agent** must have sufficient resources available to **exploit** the defender's vulnerabilities.
3. The **threat agent** must believe they will benefit from the attack.

Condition 1 is completely dependent on the defender.

Whether condition 2 is satisfied depends on both the defender and the threat agent. The defender has some influence over which vulnerabilities exist and what level of resources will be required to exploit them. Different threat agents have different capabilities.

Condition 3 mostly involves the attacker. It represents the motivation to carry out the attack. The defender may have a role if their actions provoke a threat agent to carry out an attack.

As can be seen, the threat agent and the defender jointly determine whether an attack occurs. Our method of analysis will examine these three conditions in an attempt to predict the behavior of adversaries and the impact on the victim. Analysis will also provide insight into effective

ways of preventing attacks.

The Origins of Attack Trees

We will study the modeling of attacks through the use of a graphical, mathematical, decision tree structure called an *attack tree*. There is reason to believe that *attack trees* originated in the intelligence community. At least one intelligence agency is known to have used tree-based attack modeling techniques in the late 1980s. In 1991 Weiss published a paper¹ describing *threat logic trees*. In 1994 Amoroso² detailed a modeling concept he called *threat trees*. More recently, Bruce Schneier³ (a noted cryptographer and security expert) popularized the idea, although he called it *attack trees*. Other researchers have continued to develop the idea of tree-based, threat analysis models^{4,5}.

Amenaza Technologies Limited has taken inspiration from all of these tree-based security models and enhanced them. Amenaza refers to its approach as *capabilities-based attack tree analysis*. Amenaza's SecurITree[®] software supports this form of analysis. Capabilities-based attack tree analysis was first applied to the field of information technology security. However, **attack tree analysis is effective for understanding almost any type of system** and now enjoys usage in defense, critical infrastructure, health care and banking sectors.

Attack Tree Vulnerability Models

Attack trees are constructed from the point of view of the adversary. Creating good attack trees requires that we *think like an attacker*. Initially, do not think of how to defend a system; think of ways to defeat it's security.

Like most mathematical tree models, attack trees are represented by a diagram with a single *root* node at the top. The root branches downwards, expanding through forks and more branches. This is similar to the *decision trees* used to help with business decisions or the *fault trees* used to

¹ J.D. Weiss, A System Security Engineering Process, Proceedings of the 14th National Computer Security Conference, 1991.

² Edward G. Amoroso, Fundamentals of Computer Security Technology, pp 15-29, Prentice-Hall, ISBN0131089293

³ B. Schneier, Attack Trees, Dr. Dobb's Journal, v. 24, n. 12, December 1999, pp. 21-29.
B. Schneier, Attack Trees: Modeling Actual Threats, SANS Network Security 99 – The Fifth Annual Conference on UNIX and NT Network Security, New Orleans, Louisiana. Wednesday, October 6th, 1999, Session Two, Track One - Invited Talks
B. Schneier, Seminar session given at a Computer Security Institute conference in November, 1997. See also <http://www.counterpane.com/attacktrees.pdf>

⁴ Moore, A., Ellison, R. and R. Linger, "Attack Modeling for Information Security and Survivability", March 2001, <http://www.cert.org/archive/pdf/01tn001.pdf>

⁵ Shelby Evans, David Heinbuch, Elizabeth Kyle, John Piorkowski, James Wallner, "Risk-Based Systems Security Engineering: Stopping Attacks with Intention", November/December 2004, IEEE Security and Privacy

understand the reliability of machines and automated processes.

In an attack tree vulnerability model, the topmost (*root*) node represents an objective that would be a potential goal of one or more threat agents. The root also represents a state that has a negative consequence for the defender⁶. If the goal is chosen carefully it is usually possible to analyze a system completely with a single attack tree. In some situations a particular adversary may have several different goals, or different adversaries may have their own unique goals. These situations will require multiple attack trees to carry out a complete analysis.

From the attacker's point of view, the root goal is so lofty or broadly stated that it lends little understanding as to how it may be achieved. As with most endeavors, it is often helpful to break a high level goal into smaller, more manageable steps. It is possible to formulate a number of different strategies that could be used to achieve the overall goal. These strategies can be expressed as a series of intermediate objectives that singly, or in combination, lead to the realization of the root goal. This decomposition process continues, breaking the intermediate goals into ever finer grained activities. This is easily shown in a graphical format (see **Figure 3**).

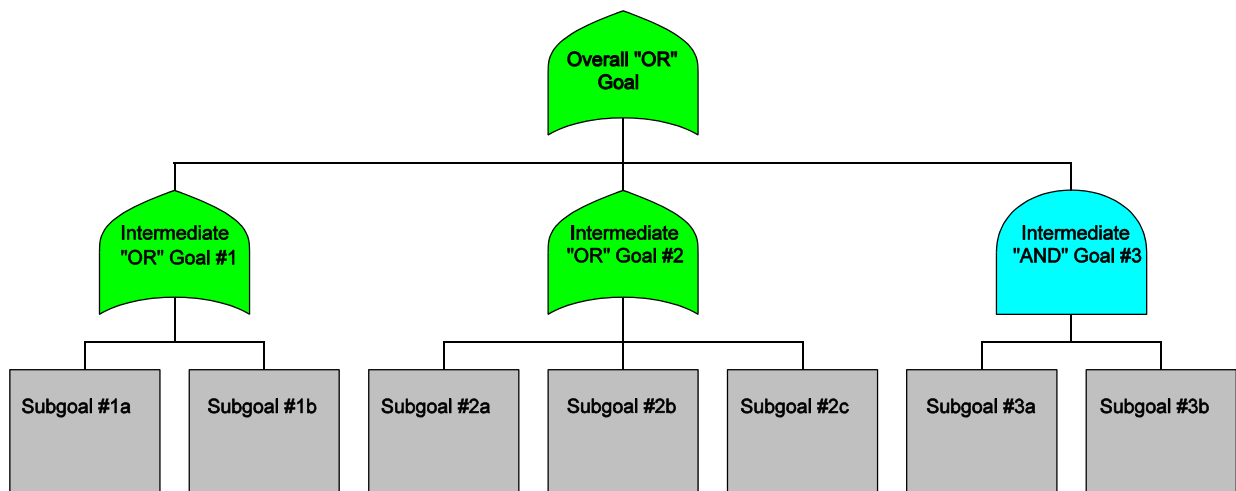




Figure 3 – Goal Oriented Tree

The topmost symbol in the tree represents the adversary's highest goal. It is often referred to as the *root* of the tree. The *root* in this particular example is depicted by a green symbol . The diagram shows how high level goals decompose into increasingly precise subgoals as we descend through the tree.

The *OR* symbol  (whose shape should be familiar to students of Boolean algebra) indicates that the root *Overall "OR" Goal* can be attained by achieving *Intermediate Goal #1 OR*

⁶ If there are no negative consequences on the defender from an attack, there is no reason spending effort in preventing it.

Intermediate Goal #2 OR Intermediate Goal #3. These in turn are further decomposed. For example, *Intermediate “OR” Goal #1* is achievable by attaining *Subgoal #1a OR Subgoal #1b.* The grey rectangular shapes, called *leaf nodes*, represent atomic activities which cannot, or need not, be decomposed further.

Intermediate Goal #3 is represented by a cyan AND symbol . This indicates that both *Subgoal #3a AND Subgoal #3b* must be completed in order to attain *Intermediate Goal #3.*

This somewhat abstract discussion will make more sense with an example.

A Sample Attack Tree

To illustrate the concept of a **capabilities-based attack tree**, let us imagine a hypothetical system we are trying to defend. Consider the home security challenge faced by the residents of a typical, suburban home. While few middle-class home owners would do a formal security risk assessment on their home, the subject was chosen as being one to which all readers could relate.

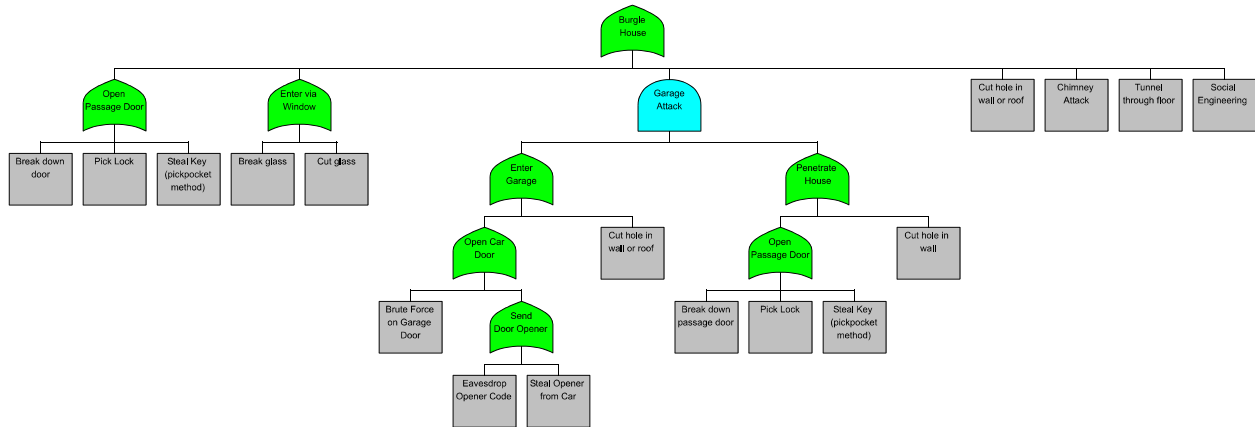


Figure 4 – Approaches to Burglarizing a House

The house we have in mind is a middle-class dwelling, complete with attached garage. The incident that concerns us is the possibility of the house being burglarized (see **Figure 4**).

After some consideration, we can think of seven possible ways in which a thief might enter the house to commit burglary:

1. Passage doors (i.e., the front and back doors normally used for entry).
2. Windows.
3. Via the attached garage.
4. Walls (including the roof – it is essentially an angled wall).
5. Chimney.
6. Floor (attacking from beneath).
7. Social engineering (convince the resident to allow entry to the attacker).


These attacks, which have been partially decomposed into more detailed steps, are shown in **Figure 4**. To simplify our example, we have restricted the decomposition to the *Open Front/Back Door*, *Enter via Window* and *Garage* attack vectors. Obviously, greater detail could also be added to the *Cut Hole in Wall or Roof*, *Chimney Attack*, *Tunnel through Floor* and *Social Engineering* attacks.

As can be seen in the diagram, there are three types of passage door attacks. The doors can be physically broken, the locks can be picked or the key can be obtained through theft. Similarly, an intruder can either cut or break the glass in the windows. To enter via the garage, the burglar must first gain entry to the garage and then enter the house (either through the wall or by penetrating the passage door leading from the garage to the house).

Decomposition of higher level events into smaller, more precisely defined events could continue almost indefinitely. For our purposes, it need only continue to the point where further decomposition will not increase the understanding of the intended viewers of the model. For example, the *Break glass* leaf node could be decomposed into the steps of picking up a rock and throwing it at the window. This is unnecessary since almost everyone knows how to break a window. On the other hand, the leaf node that deals with *Eavesdrop opener code* should be decomposed into smaller steps to enhance the analyst’s understanding of the actions to be performed by the burglar. We have not described this attack to that level for purposes of brevity.

It is important to note that the adversaries’ interaction with the system they are attacking takes place entirely at the leaf nodes. For that reason, some people call the leaf nodes *attack stabs*. All of the non-leaf nodes in an attack tree represent logical states that the attacker achieves through their efforts at the leaf nodes.

Attack Scenarios



An attack tree shows a logical breakdown of the various options available to an adversary. By performing the exploits associated with one or more leaf level events which have been carefully selected to satisfy the tree’s AND/OR logic, the attacker can achieve the root level goal. Each minimal combination of leaf level events is known as an *attack scenario*. The combination is minimal in the sense that, if any of the leaf events are omitted from the *attack scenario*, then the root goal will not be achieved.

Associated with each *attack scenario*’s set of leaf nodes is the collection of intermediate nodes that are activated along the path (or paths) to the root goal. Strictly speaking, these intermediate nodes are not part of the *attack scenario*, but we usually include them in graphical depictions to make it clear how the attack will take place.

The complete set of attack scenarios for an attack tree shows all of the attacks that are available to an attacker who possesses infinite resources, capabilities and motivations. One particular attack scenario from the house burglary tree is shown in **Figure 5**. It consists of two leaf level events: *Brute Force on Garage Door* and *Break down passage door*. Both events

Figure 5 – Attack Scenario Example

are required due to the AND node (*Garage Attack*) several levels above.

Behavioral Indicators – Hardness Metrics

To this point, our attack tree shows how attacks could occur, but not whether they will. Intuitively we know that a burglar will choose to break a window rather than digging a tunnel underground. We suspect this is because it is easy to break windows and difficult to dig tunnels. It seems reasonable to suppose that an attack's level of difficulty affects the behavior of an adversary. What is required is an objective, quantitative way of measuring difficulty

Performing the exploit activities associated with leaf nodes requires the attacker to possess and expend resources. These resources may include money, technical skill, time and a willingness to suffer a penalty (jail time, injury, death) for the carrying out the exploitive actions. **Even a highly motivated threat agent can only carry out a given attack if their resources meet or exceed the resources required to perform the exploit(s).** Therefore, the scarcity of resources constrains a threat agent's behavior. Capabilities-based attack tree analysis incorporates resource metrics into attack tree models to determine the likelihood of attacks. We call these resource metrics *behavioral indicators* because they influence the behavior of adversaries.

To add *behavioral indicators* to the attack tree model, the analyst selects the classes of resources, assets or traits that are required to carry out specific exploits. These factors constrain, to a greater or lesser degree, the ability of an adversary to do the exploit. These values are seldom available from published references. Instead, the analyst estimates the amount of each resource required. This may require consultation with subject matter experts. This process must be performed for each resource type and every leaf node in the model. Although there is some uncertainty in the estimated values, they should at least be of the correct order of magnitude.

Ideally, the indicators should be **orthogonal**. This means that the behavioral influence of one indicator is independent of another. For example, an attacker's bank balance is largely unrelated to their willingness to be apprehended in an attack. Therefore, *Cost of Attack* and *Noticeability* are orthogonal indicators. Sometimes, however, there are unavoidable dependencies between indicators. For instance, in many cases it is possible to use money to buy technical skill. Thus, *Cost of Attack* and *Technical Skill* are not completely orthogonal. Complete independence is not always achievable. Indicators should be chosen that minimize dependencies as far as possible.

For the house burglary model we might choose to define three behavioral indicators: *Cost of Attack*, *Technical Ability* and *Noticeability*. *Cost of Attack* is the amount of money (expressed in USD) that the adversary will need to spend. *Technical Ability* uses an arbitrarily chosen rating scale of 1-100. *Noticeability* ranges from 0 to 1.0 and is a metric of how obvious the attack is to the homeowner and neighbors.

The resource estimates for the *Cut Window* leaf node are shown in **Table 1**. For this particular exploit, the cost is low and technical skill required is moderate. Cutting glass, when done well, is fast and quiet. We believe that it is unlikely to be noticed due to the short time required.

Resource	Value	Description
Cost of Attack	\$5	Glass cutters are inexpensive.
Technical Ability	40 (out of 100)	Cutting glass is tricky, particularly when vertical.
Noticeability	0.2 (out of 1)	Cutting glass is quiet and only takes a few seconds.

Table 1 – Behavioral Resources Required to *Cut Glass*

The *Cut Glass* exploit’s resources can be compared with the requirements for the *Eavesdrop Door Opener* task shown in **Table 2**. Executing this radio playback exploit is expensive and complicated, but very stealthy.

Resource	Value	Description
Cost of Attack	\$5,000	Requires receiver, computer and transmitter.
Technical Ability	70 (out of 100)	Considering the field (house burglary) this is hard!
Noticeability	0.05 (out of 1)	Park across the street and wait. Quite stealthy.

Table 2 – Behavioral Resources Required to *Cut Glass*

Note that the total amount of resources for a given attack is dependent upon the collection of leaf nodes that need to be performed by the adversary. The *Cost of Attack* for the *Garage Attack* scenario is shown in **Figure 6**.

Indicators are assigned to the tree by the analyst as seems relevant to the problem at hand. Typically three or four indicators are used. Too few indicators leads to a flat, one-dimensional understanding of the forces that drive incidents. An excessive number of indicators may lead to such complexity that the forest is lost in the trees⁷.

Impact Indicators – Logical Effects

Whenever an adversary launches an attack on a target, there are consequences to their actions. If there were no negative effects then the defender wouldn’t care about the attack and, without the hope of some benefit, the adversary would be unlikely to carry it out.

When an exploit happens, the victim may suffer some damage. However, the damage that results directly from the exploit itself is often insignificant. The damage accumulates and grows as the attack propagates upward through intermediate nodes towards the

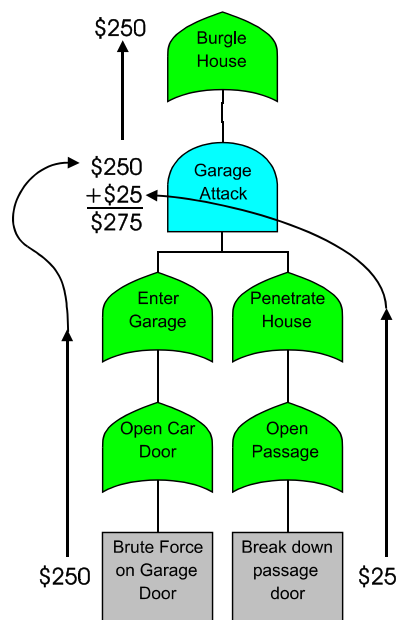


Figure 6 – Behavioral Cost for Garage Attack Scenario

⁷ Surely you realized that, somewhere in this document, that pun would be used?

root node. While there may be various paths that lead the attacker from the leaf node(s) to the overall root goal, the total negative and positive impacts are path dependent.

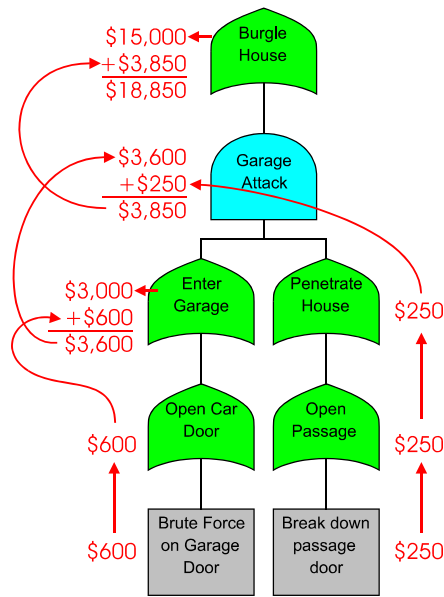


Figure 7 – Garage Attack Damage Cost

In the case of our house burglary model, there are immediate damages to windows, doors and other items at the exploit (leaf) level. If the attacker chooses an attack that passes through the garage⁸, they will find tools and sporting goods waiting to be stolen. In the house described in our model, the value of the garage items is \$3,000. The house itself contains \$15,000 of items that can be easily stolen. So, depending on the attack used, the cost to the home owner could range from \$15,000 to \$18,000 in stolen goods, plus a little extra in damage to the house itself (see **Figure 7**).

Since attacks require the investment of time and resources on the part of the adversary, it stands to reason that they expect to benefit in some way. In the case of a house burglar, the benefit is fairly obvious. The burglar obtains goods which can be sold for a fraction (say $\frac{1}{3}$) of the replacement value. In the case of the garage attack (shown in **Figure 7**) the attacker's earnings would be $\$18,000 \times \frac{1}{3} = \$6,000$ ⁹.

To represent these impacts within our models we need to create *impact indicators*. Impact indicators are very similar to behavioral indicators. One major difference is that values for behavioral indicators are only input at the leaf nodes. Impact indicator values can be input at any level in the tree. Usually, the largest business impacts occur high in the tree. For victims, impact indicators reflect things like loss of money, damage to reputation or even casualties.

The leaf-level events associated with a specific attack scenario are useful for detecting when a particular attack is underway. It is possible to build an attack detection system which compares leaf level events detected by sensors with attack scenarios and sounds an alarm when an attack is underway. Although this may be overkill for recognizing the relatively obvious house burglary attacks, it can be useful in more complex situations. When the attack tree is large, and there are hundreds or thousands of events to monitor for, it is useful to detect those combinations of events that are significant. This could conceivably eliminate many of the false positives present in existing intruder detection systems.

Predicting Attacks

A very simple premise can be used to understand attackers' behavior:

⁸ Of course the intruder might get directly into the house, then go to the garage. We are ignoring this possibility in this somewhat oversimplified example.

⁹ Note that the burglar receives no reward for breaking doors and things. So the \$850 of property damage experienced by the homeowner is a side effect of the attack that provides no benefit to the attacker.

IF they want to AND they can THEN they will

In other words, if there exist adversaries that are motivated to harm the system, and they possess the capability to carry out the exploits (along with a willingness to accept the consequences¹⁰ of their actions), then, sooner or later, they will carry out a successful attack on the system.

There are a few situations in which this premise might appear to fail. The adversaries might not attack if they are unaware that the system exists¹¹. They might not attack immediately if there are so many other systems with similar (or greater) defects that they simply haven't gotten around¹² to hitting the one we are concerned with (yet). However, unless all of the capable adversaries are caught before they can get to you, your number will come up.

Adversarial Capabilities

As stated earlier, whether or not attacks occur depends as much on the adversary as on the system being assaulted. Our focus, which has been directed to the system under consideration, now shifts to modeling the adversaries.

Who are our enemies? Most defenders have some instinctive idea of the types of people with a plausible interest in harming their system. Who these adversaries are depends greatly on the nature of the system being defended. Common types of adversaries include vandals, *script kiddie* computer hackers, opportunistic thieves, professional thieves, serious computer hackers, industrial spies, irate customers/partners, insiders (disgruntled employees), terrorists and foreign governments.

The defender should identify the types of people who might benefit from an attack (threat agents) on the defender's system. A table can then be prepared showing the resources estimated to be available to each potential adversary for each behavioral indicator defined in the attack tree model. For our house burglary situation, we might identify *juvenile delinquent* and *cat burglar* as our threat agents. Our estimates of their capabilities are shown in **Table 3**.

¹⁰ Our definition of *capability* includes the attacker's tolerance for embarrassment, financial loss, personal harm or even death. We are unable to think of a more appropriate term.

¹¹ If the secrecy surrounding the system is a significant part of the system's defenses then the ways in which the system could be discovered should be modeled in the attack tree. One could argue that one of the capabilities required to attack any system is knowledge of the system's existence.

¹² More formally, we might state that, attacking the other targets either brings greater benefits to the attacker or their resource costs for attacking the alternate target are lower.

Threat Agent	Budget	Willing to Attempt Attacks that Have a Noticeability of	Technical Capability 1 - 100 scale
Juvenile Delinquent	\$50	50%	25
Cat Burglar	\$5,000	10%	70

Table 3 – Capabilities of House Burglar Threat Agents

We believe the Juvenile Delinquent is an angry (probably male) youth who doesn't have much money to spend on burglarizing houses. He (or she) is not worried about getting caught. A misspent youth has prevented our miscreant from developing technical skills.

The Cat Burglar, on the other hand, is a pro. Our felonious filching feline views burglary as a type of employment. He is willing to spend money to make money. He is prepared to spend up to \$5,000 on tools. Like any professional he has studied his subject well and is quite capable of picking locks, deactivating simple alarms and jimmying windows. The one thing he is not willing to do is go to jail.

These profiles are assumptions based on the information available to us as well as expert opinion. An attack tree will show you the logical outcome of your assumptions. It forces you to explicitly state your assumptions and exposes them for review and critique by other professionals. **The accuracy of predictions depends on the correctness of the assumptions about the threat agents and the attack tree model.**

Identifying Probable Attacks

IF they want to AND they can THEN they will

This simple statement sums up the core tenet of attack tree analysis. So far, the attack tree model we have built tells us what resources are required to carry out each attack scenario. It follows that, if we understand what resources are available to our adversaries, we will know what they can and cannot do.

When you have eliminated the impossible, whatever remains, however improbable, must be the truth.

Sir Arthur Conan Doyle

As was so elegantly stated by Sherlock Holmes, one way of predicting what has (or will) happen is to eliminate everything that cannot.

Recall that

$$\text{Risk} \equiv \text{Incident Probability} \times \text{Incident Impact}$$

As the discussion has pointed out, capabilities-based analysis compares threats to vulnerabilities to find the *incident probability* term in the equation. In other words,

$$\text{Incident Probability} = \text{Threat} \times \text{Vulnerability}$$

Recall that our underlying premise (IF *they want to* AND *they can* THEN *they will*) states that it is the motivation (*they want to*) in combination with capability (*they can*) that determines probability (*they will*). So, the magnitude of a threat can be expressed as

$$\text{Threat} = \text{Capability} \times \text{Motivation}$$

Making the substitutions allows the risk equation to be rewritten as

$$\text{Risk} \equiv (\text{Capability} \times \text{Motivation} \times \text{Vulnerability}) \times \text{Incident Impact}$$

Our capability table contains estimates of the resources available to a set of adversaries. These adversaries were chosen with the belief that they had some reason to want to harm the system. At this stage we have not attempted to quantify which of the adversaries are most highly motivated. For the present we will assume that their motivation is significant and roughly equal. If motivation is constant, the risk equation reduces to

$$\text{Risk} \equiv (\text{Capability} \times \text{Vulnerability}) \times \text{Incident Impact}$$

This means that we can estimate the probability of a given adversary carrying out a particular attack scenario by comparing the adversary's capabilities with the system's vulnerabilities. This is straightforward using the attack tree model we have created. We need only remove the attack scenarios that have behavioral resource requirements beyond the threat agent's capabilities. Simply compute the attack resource requirements for each and every attack scenario in our tree and eliminate those whose resource requirements exceed those of the chosen threat agent. For example, **Table 4** shows which of the scenarios from the house burglary tree are within the reach of a Juvenile Delinquent (based on **Table 3**).

Scenario #	Attack Scenarios for House Burglary Tree	Cost of Attack	Noticeability	Technical Ability
1	{Break down door}	25	0.3	10
2	{Pick Lock}	250	0.15	65
3	{Steal Key (pickpocket method)}	2	0.2	70
4	{Break glass}	1	0.3	2
5	{Cut glass}	5	0.2	40
6	{Brute Force on Garage Door, Break down passage door}	275	0.64	45
7	{Brute Force on Garage Door, Pick Lock}	500	0.608	65
8	{Brute Force on Garage Door, Steal Key (pickpocket method)}	252	0.68	70
9	{Brute Force on Garage Door, Cut hole in wall}	251	0.68	45
10	{Eavesdrop Opener Code, Break down passage door}	5,025	0.145	70
11	{Eavesdrop Opener Code, Pick Lock}	5,250	0.069	70
12	{Eavesdrop Opener Code, Steal Key (pickpocket method)}	5,002	0.24	70
13	{Eavesdrop Opener Code, Cut hole in wall}	5,001	0.24	70
14	{Steal Opener from Car, Break down passage door}	30	0.28	10
15	{Steal Opener from Car, Pick Lock}	255	0.216	65
16	{Steal Opener from Car, Steal Key (pickpocket method)}	7	0.36	70
17	{Steal Opener from Car, Cut hole in wall}	6	0.36	45
18	{Cut hole in wall or roof, Break down passage door}	275	0.64	45
19	{Cut hole in wall or roof, Pick Lock}	500	0.608	65
20	{Cut hole in wall or roof, Steal Key (pickpocket method)}	252	0.68	70
21	{Cut hole in wall or roof, Cut hole in wall}	251	0.68	45
22	{Cut hole in wall or roof}	250	0.6	45
23	{Chimney Attack}	25	0.01	98
24	{Tunnel through floor}	20,000	0.2	60
25	{Social Engineering}	50	0.85	40

Table 4 – House Burglary Attack Scenarios Impossible for a Juvenile Delinquent

All impossible attack scenarios have been eliminated from consideration. The remaining three scenarios (#1, #4, #14) are within the capability of the Juvenile Delinquent. Note that this analysis does not determine the exact probability of the viable attack scenarios¹³. It only tells us that for scenarios #1, #4 and #14, the value is clearly > 0 for cases where the threat agent is a juvenile delinquent.

Combining Impact with Incident Probability Yields Risk

In order to assess the risk of the probable attacks identified above we must now incorporate the corresponding impacts into our analysis. To do this we reexamine the viable Juvenile Delinquent attack scenarios, this time including the scenario's impact value on the victim as computed from the model (shown in **Table 5**).

¹³ Later we will show how to refine our assessment of the probability of each attack.

Scenario	Attack Scenarios	Damage Cost	Cost of Attack	Notice-ability	Technical Ability
14	{Steal Opener from Car, Break down passage door}	18,500	30	0.28	10
1	{Break down door}	15,250	25	0.3	10
4	{Break glass}	15,150	1	0.3	2

Table 5 – Impact of Attack Scenarios Available to Juvenile Delinquent

Note that **Table 5** has been sorted by *Damage Cost*. Since we have eliminated all impossible scenarios, and have assumed that those that remain have roughly equal probability, **the risk level is determined entirely by the impact**. Hence, **Table 5** is a prioritized list of the risks to the house from a juvenile delinquent.

Estimating Motivation

In many cases the assumption that all of the selected threat agents have the same level of motivation is sufficient. **For cases in which we do not have a good understanding of the adversary’s mind it may be the best we can do**. Sometimes, however, it is necessary and possible to provide a better estimate of motivation (and therefore, risk).

As noted earlier, the adversary doesn’t (usually) care what the victim loses¹⁴. Rather, they are looking to gain one or more benefits from the attack. An adversary’s decision to execute an attack scenario is made based on some sort of cost-benefit assessment.

Scenario	Attack Scenarios	Attacker Gain	Cost of Attack	Notice-ability	Technical Ability
14	{Steal Opener from Car, Break down passage door}	6,005	30	0.28	10
1	{Break down door}	5,000	25	0.3	10
4	{Break glass}	5,000	1	0.3	2

Table 6 – Attack Scenario Benefits for Juvenile Delinquent

The model we have created can easily show the attacker benefits for each scenario. For example, the juvenile delinquent benefits monetarily in varying amounts in each house burglary scenario (see **Table 6**). The items stolen can be resold (at a discount) on the street.

At first glance, the burglar might choose the most lucrative attack (scenario #14). Very quickly they would realize that earnings is not as important as profit. I.e., each scenario has costs as well as benefits. In certain cases they might even focus on the cost benefit ratio (ROI).

Monetary costs are not the only criteria used by an attacker. They might use some formula which weighs costs and benefits in choosing attacks. Although this sounds very formal for a simple house burglar, it is quite likely that they unconsciously use such a weighting scheme.

¹⁴ The exception is where the adversary seeks revenge. Even there we maintain that the actual loss is of little importance to the adversary. They value the suffering the loss will cause the victim.

Analysts interested in these sorts of advanced analysis techniques would be well served to investigate the research by Evans and Wallner⁵ cited previously.

Modeling Countermeasures

The risk analysis techniques discussed identify attack scenarios beyond the comfort level of the defender. Attack tree modeling can also be used to find and test solutions to problem areas.

One approach is to examine all of the leaf nodes that form part of an adversary's attack scenarios for ways in which the vulnerabilities can be hardened. In some cases, this is necessary. However, since trees tend to grow exponentially descending from the root, there may be a great number of leaf nodes to deal with.

A better solution is to find portions of the attack tree that are common to multiple scenarios. Oftentimes it is possible to introduce an architectural change high in the tree which will resolve issues with many descendants. This typically involves the creation of an AND node. Beneath the AND node is placed the section of the tree that is easily compromised. Additionally, some new technology, control or process is introduced which forms a sibling under the AND node. If the change is unattainable by the adversary then the attack can never progress beyond the AND node.

Proposed changes can be modeled before they are implemented. This gives the analyst the opportunity to test their effectiveness by repeating the pruning and attack scenario generation. If they are demonstrated to be effective, then the cost benefits of the change can be verified to ensure that there will be a positive return on investment.

The Need for Analysis Tools

The techniques described earlier are based on simple concepts. The examples given have, for the most part, been small enough that they could be carried out by a person with pencil and paper. However, these operations quickly become unwieldy when applied to attack trees of sufficient complexity to describe meaningful problems. All desire to experiment with the model by changing the tree or altering the assumptions about the threat agents would quickly vanish once the analyst realizes the effort required to prune or calculate attack scenarios.

The answer to this problem is to provide a software tool capable of performing these operations with the click of a mouse. Just as a spreadsheet program removes the tedium of performing the cascading calculations caused by updating a spreadsheet cell, an attack tree analysis tool can free the analyst to use his or her insight in understanding the system.

Amenaza Technologies has produced such a tool. Secur//Tree[®] is the world's first commercially available attack tree modeling and analysis tool. With Secur//Tree you really can see the forest through the trees!

Attack Tree Analysis versus Traditional Risk Analysis Methodologies

A conventional, statistics-based risk assessment might tell you how likely you were to have your house burglarized and what damage you might expect to suffer. Unfortunately, it would not provide anything more than general guidelines (so called, "best practices") on how to be more

secure.

The attack tree analysis allows us to see the underlying forces that channel the attacker's behavior. For example, the home owner might notice that, simply by removing their garage door opener from the car in the driveway, the riskiest attack from juvenile delinquents is eliminated.

Although statistics might be available for house robberies, this is not true of many other types of illicit activity. Without statistics, the conventional risk analysis process is unable to convincingly predict which attacks are likely to occur. This leaves the analyst to rely on guesses which, even if correct, are not supported by evidence. If the analyst chooses to alter the system to mitigate certain risks, the conventional risk assessment methodology provides no suggestions as to which changes will be most effective. The result is a series of highly subjective decisions for which the reasoning process is undocumented. Sooner or later, problems arise – and no one can remember the rationale behind the recommendations. This leads to indefensible due diligence positions and legal exposure.

Attack tree models are largely self documenting. The assumptions about the systems vulnerabilities are captured in the tree itself. The assumptions about the threat agents are stated in the table of threat agent capabilities. The conclusions are reached through the mathematical operation of applying the threat agent's profile to the model (pruning). This is much more reliable than depending on an analyst's memory.

Conventional approaches to risk analysis are also very time consuming. This makes analysts reluctant to update them when the system or environment changes. As a result, only a snapshot of the system at some point in time is considered. By the time the analysis is complete it is no longer relevant because the system being studied has changed. Attack tree models can (with the proper tools) be updated and reevaluated in minutes.

Appendix A

Risk Theory and Definitions

The definitions of risk related terms vary slightly from author to author. Here are the meanings that we will use throughout this document.

Risk – Traditional Definition

Traditionally, the **risk** associated with a particular event can be defined as:

$$\text{Risk} \equiv \text{Incident Probability} \times \text{Incident Impact}$$

While this formula is correct, in many situations it is not useful. Although it is usually straightforward (if tedious) to estimate the potential damage caused by a hypothetical incident, it is not always obvious how to find a value for the *Probability of the incident* term. That term's value (usually expressed as a number between 0 and 1) is a result of many factors, some of which may not be easily quantified. We will shortly show an alternate form of the equation that is easier to work with.

System

Whenever we consider risk, we have to establish what is included in the scope of our analysis. Philosophers would probably argue that an injurious event ultimately affects everyone in the world, both today and down through eternity. Most other people limit their concern to things for which they have a direct responsibility or that affect them directly. The area of consideration is usually called, the **system**.

Webster's dictionary defines the word **system** as a *regularly interacting or interdependent group of items forming a unified whole*. Although a *system* almost certainly contains a variety of physical components (such as computers, buildings), systems may also include the people that interact with the components and the processes they use to do so. An early step in risk analysis is to decide which components make up the **system** being studied.

Vulnerability

All systems suffer from one or more **vulnerabilities**. A **vulnerability** is a weakness in a system. It is a mechanism by which a system could be damaged, its resources used in an unauthorized way or caused to enter an undesirable state. For instance, a computer system that authenticates users via passwords is vulnerable to password guessing. The classic Greek hero, Achilles, was only vulnerable to injury in his heel.

Threat

A **threat** is a potential source of danger to a system. A **threat** is something that might act on a specific vulnerability or set of vulnerabilities. The presence of a threat does not guarantee that the threat will act on a vulnerability. Rather, it shows the potential for this to occur. A threat may originate from a hostile, intelligent agent that desires to inflict damage or it may be a product of random conditions in the environment. For example, the possibility of metal stakes being driven into trees is considered a threat to safe logging operations. The potential of

lightning strikes are a threat to people who enjoy walking outdoors during thunderstorms.

Threat Agent

A class or group that embodies a threat to a system is known as a **threat agent**. Using the example given earlier, radical environmentalists are one instance of a **threat agent** willing to drive metal rods into trees to interfere with logging. However, a disgruntled ex-employee hoping to cause financial damage to his former employer would also be a valid **threat agent**. If you are protecting a computer system that contains trade secrets then both industrial spies and adolescent script kiddies are plausible threat agents.

Strictly speaking, every class of individuals that has the potential of carrying out an attack on the system constitutes a **threat agent**. However, it is more useful to consider as **threat agents** only those entities that perceive a benefit from hostile action. The extent to which a **threat agent** believes they will gain something of value through an attack corresponds to their level of **motivation**. Generally speaking it is only necessary to consider **threat agent's** that have motivation to carry out an attack.

This means that, by our definition, the US Army is not a plausible threat agent against a New York bank. Although the US Army certainly has the ability to invade the bank and steal the money, we can think of no reason why it would be motivated to do so¹⁵. To emphasize that our usage of the term includes the concept of motivation, we will often speak of plausible or viable threat agents. A **threat agent's** level of motivation is related to the benefits they believe they would achieve through an attack.

Exploit

Whereas a threat is an abstract way to take advantage of a vulnerability, an **exploit(n)** is the detailed procedure for doing so. The term is frequently used in connection with attacks on computer systems. For example, it may be known that a software application suffers from a *buffer overflow* vulnerability. That is, if excessive data is supplied to the program's input, the program may behave in a way that is inconsistent with its design. An **exploit** that makes use of this vulnerability would consist of the exact procedure needed to cause the program to misbehave. It would include the method by which the data would be transmitted, the sequence of characters that would be used and any other details needed to make use of the vulnerability. When used as a verb, **exploit(v)** means the act of carrying out the malicious procedure.

Incidents and Events

Again citing Webster's, an **incident** is *an action likely to lead to grave consequences*. In essence, when a threat ceases to be a merely hypothetical possibility and a vulnerability is acted upon, it becomes an **incident**.

In many cases, an **incident** occurs as a result of a sequence or combination of other, contributing **events**. For example, a car may have an accident as a result of a flat tire. At first glance, the flat tire is the **incident** that caused the accident. However, the flat tire was the result of a puncture

¹⁵ If the government wants money they can simply print it!

by a nail. The nail fell from a passing construction truck. This occurred because a construction worker neglected to put the box of nails in the toolbox. Each of these events is an **incident** that caused a subsequent **incident**. Some of the **incidents** caused a degree of harm immediately while others only led to unpleasant consequences.

We tend to use the term **event** for lower level actions which may, or may not, have immediate consequences. The term **incident** is usually used to describe the event or events higher in the causal chain that are more directly associated with the resulting *grave conditions*. This distinction is largely artificial since almost any event can be decomposed into more detailed events. **Incident** and **event** are practically synonyms.

Attacks and Mishaps

There are two types of **incidents**. An **incident** that is caused by a conscious, deliberate application of an exploit is called an **attack**. **Incidents** that result from unintentional or random events are called **mishaps**. The attack tree methodology focuses primarily on **attacks**.

Victim Impact and Attacker Benefit

Incidents and events generally cause damage to the system involved. If they did not, there would be no reason to try to prevent them. The damage is called **victim impact** or, more simply, the **impact**. This damage is usually expressed in monetary terms, but may be tallied using other metrics – e.g., casualties.

As mentioned earlier, when the incident occurs as the result of a deliberate attack, it was with the expectation that there would be a positive impact or benefit to the attacker. This is called the **attacker benefit**. Many different metrics may be used to measure **attacker benefit**.

In some cases, the **victim impact** of an attack is equal to the **attacker benefit**. For example, if an attacker steals \$1000 then one party loses and the other gains \$1000. However, this is the exception rather than the rule. Vandalism illustrates this point. The victim may suffer significant financial damage while the attacker gains nothing of monetary value.

Depending on which vulnerability in a system is targeted, and which exploit is used, the **victim impact** and **attacker benefit** will vary. All other factors being equal, **threat agent's** will choose attacks that yield the highest reward. Analysts interested in this type of advanced analysis would be well advised to examine the work of Evans and Wallner.