

Downstream Liability for Attack Relay and Amplification

[This article is an adaptation of a talk delivered at the RSA Conference 2002 in San Jose, California. All contents Copyright 2002 - Carnegie Mellon University, Pennsylvania State Police, and White Wolf Security.]

Disclaimer

Points of view or opinions expressed in this presentation do not necessarily represent the official position or policies of the Pennsylvania State Police, Carnegie Mellon University, White Wolf Security, or RSA.

Who are the authors?

- Scott C. Zimmerman, CISSP, is a Research Associate at the Software Engineering Institute, Carnegie Mellon University.
- Ron Plesco, Esquire, is the Director of Policy for the Pennsylvania State Police.
- Tim Rosenberg, Esquire, is the President and CEO of White Wolf Security (www.whitewolfsecurity.com).

The Scenario

To demonstrate the concepts involved, we will use a simple and hypothetical scenario in which four distinct entities are involved:

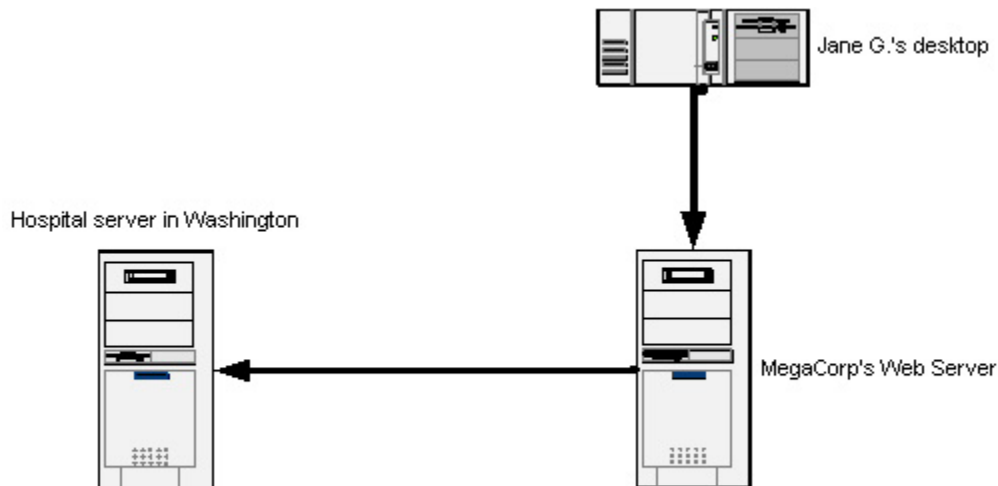
1. The first entity is Jane G. Jane is a network security administrator in the United Kingdom. She works for a company that does approximately US\$200M in business per year. Her yearly salary is US\$55,000.
2. The second entity is Megacorp's web server, a non-mission-critical machine accessible from the Internet. MegaCorp is a US\$10.4 billion/year public company. The server is hosted internally, and is physically located at MegaCorp's facility in Iowa. MegaCorp exercises complete control over all aspects of the web server.
3. The third entity is a web server that belongs to a non-profit research hospital in the state of Washington.
4. The last entity is Mr. Big Star, who receives medical treatment at the research hospital.

While accessing the Internet at work, Jane finds a six-month old vulnerability in Megacorp's web server. Exploiting this vulnerability, Jane is able to gain privileged

access to the system. From Megacorp's system, Jane then discovers a month-old vulnerability on the hospital system located in Washington state. She is able to exploit this as well and gains privileged access to the hospital server. Once Jane is a privileged user on the hospital's system, she is able to penetrate more deeply into the hospital's network wherein she finds a database server containing sensitive patient records. While browsing the database, Jane G. stumbles on Mr. Big Star's file and decides to download a copy.

Having finished her shift at work, Jane G. installs a Denial of Service attack tool on the MegaCorp server. She begins an attack against the hospital's web server to throw the administrators off her trail. She goes home and posts Mr. B. Star's file to a web site in Canada and sends it to her friends on IRC.

The chain of entities looks like this:



Parties Involved – Legal Issues

Before we can discuss the various legal theories under which the suits can be brought, we must first articulate which parties are involved in the case. The *plaintiff* is the person or entity that was harmed by the act and is seeking restitution. The *defendant* is the person or entity accused of committing the act. In this scenario, potential plaintiffs include

- MegaCorp
- The hospital
- Mr. Big Star

Note that this is not an exhaustive list, as we are focusing on the specific group of directly harmed individuals. Potential defendants include

- Jane G.
- MegaCorp
- The hospital

It may seem strange that the hospital, for example, may be both a plaintiff and a defendant, but in this case the hospital may seek damages from MegaCorp, and Mr. Big Star may seek damages from the hospital. Unfortunately, events such as these are akin to a multiple vehicle accident. We are presented with a large number of parties who have been harmed, none of which is exactly sure what happened. What will happen in both the multiple site attacks and the car accident is that all parties even remotely associated to the incident will be listed as possible defendants and/or plaintiffs. Once the case lands in court, it is up to the jury and the legal system to decide who did what to whom, who will pay, and how much.

The Legal Theories

We now have a series of possible parties to the case. The next portion of the analysis is identification of the legal theories under which the parties might be sued. This is a difficult process as the law is a very specific creature. For the purposes of the next section, we are going to focus on downstream liability. The crux of the downstream liability issue is *negligence*. Negligence consists of four parts: *duty*, *breach*, *causation*, and *damages*. We will approach each of these separately. Keep in mind that, in the real world, separation of these items is extremely difficult as they are all closely linked together.

Duty is simply defined as a prudent person's obligation to use reasonable care. A more detailed definition can be found in Prosser, Wade, and Schwartz's Cases and Materials on Torts: "requiring the actor to conform to a certain standard of conduct, for the protection of others against unreasonable risks". To use an automotive analogy, a driver has the duty to ensure his vehicle has fully functioning brakes and lights, good tread on the tires, and so forth. Furthermore, the driver of the vehicle has the duty to operate her car with reasonable care and not to drive recklessly. One of the most difficult aspects of showing

negligence is this: is there a clearly defined *duty*? In other words, regarding downstream liability, does an owner of IT assets on the Internet have a duty to keep his systems secure and not to be used to hurt another? We believe the answer to this question is a resounding yes.

Assume for now that the duty exists; showing negligence means there must be a *breach*. For a *breach* to occur, the plaintiff must show that the defendant failed to perform her duty. In the worst case, the defendant did nothing at all to address network security issues. In the less extreme case, the defendant could simply have failed to perform her duty to the appropriate standard. Either will suffice to show a breach in the duty, as long as the remainder of the requirements are met.

Causation means that the aforementioned *breach* caused the damages in the incident. In this case, you will have to show what each of the parties did (or didn't do) which led to some real damages. It is imperative for the plaintiff to directly link the breach in duty to very specific damages, and show that the damages which would not have been incurred *but for* the breach.

In order for *damages* to be awarded, something has to be harmed. Damages are broken down into three types:

- Nominal – just enough to say ‘you won’
- Compensatory – repayment for actual and real damages
- Punitive – Amount above compensatory to punish the defendant and make an example so as to deter similar conduct in the future

In our scenario, disclosure of Mr. Big Star's medical condition leads to termination of contract negotiations for a US\$15M lead role. This dollar figure defines the damages caused to Mr. Big Star by Jane G. through MegaCorp and the hospital. In some cases, the damages may not be as visible. Revenue lost through a disabled e-commerce site can be quantified, but what about loss of consumer trust?

What role does Jane G.'s employer play in the event? Her employer provided the computer and Internet connection to perpetrate the act. The legal world has created a theory of vicarious liability in this case, known as *Respondeat Superior*. Under this theory, the harmed plaintiffs may be able to sue Jane's employer for compensation. This is beneficial from the plaintiff's perspective as the employer typically has more financial resources than the employee. Under the theory of *Respondeat Superior*, an employer could be held vicariously liable for its employee's actions:

- Where an employee is acting within the scope of employment and doing something in the furtherance of his work; and
- The employer is or should be exercising some control; then
- The employer will be liable for the negligent acts of the employee

Jane G. is a network security administrator, and she conducted the attacks while at work, using her employer's resources. If her employer has published policies in place, and enforces them regularly, it will be difficult to hold Jane's employer vicariously liable. To make this determination, one will have to look at their employment practices and internal policies.

Jane G.'s employer may also have been negligent in its hiring practices (though we did not directly address Jane's background or character). If an employer hires a network security administrator who has a questionable background, one of two things probably happened:

- The employer did not conduct a thorough background check.
- The employer did conduct a background check but ignored the findings.

A similar situation would be that of a doctor who has committed malpractice at – and was dismissed from – his last three positions. Hospital #4 hires him without conducting a thorough background check, and the doctor commits malpractice yet again. The hospital would then be guilty of *Negligent Hiring*.

Keep in mind that negligence, vicarious liability, and negligent hiring all assume that a duty exists. Herein lies the difficulty: what is the due standard of care in a given situation? What are the accepted best practices? What, exactly, should MegaCorp have done to avoid being used as a conduit to the hospital intrusion? In general the duty is defined as the actions taken by “a reasonable and prudent person”. Unfortunately this definition provides a wide range of possibilities: one person's “reasonable” and “prudent” is another person's “overkill” and yet another person's “insufficient”. The problem often becomes the need to discover what these terms mean in a given trade or industry. However, a caveat applies: the tendency of an industry to be generally negligent in its practices does not mean that the court will - or should - use these practices as the de facto standard. Since our scenario deals with network security, the focus areas here will be architecture, patches, and personnel.

Architecture

One of the most widely-deployed network security measures is the firewall. In broad terms, this is a system that resides between the corporate network and the rest of the Internet, filtering traffic according to its configuration. Ten to fifteen years ago, firewalls were strange and almost unheard-of beasts. However, times have changed, and any organization that does not protect its network with a firewall is likely to be greeted with incredulity and dismay.

The Distributed Denial-of-Service attacks that affected prominent web sites in 2000 and 2001 contained thousands upon thousands of spoofed packets. Spoofed packets can be generated by freely available software tools, and contain an invalid or incorrect source address; the source address is not important as the flooding is meant to be a one-way communication. The DDoS attacks were made possible by the almost nonexistent use of

egress filtering by network-connected entities. Egress filtering is a simple concept: examine packets as they **leave** the corporate network to ensure no inappropriate or malicious traffic escapes into the world. For example, spoofed packets should not be allowed to leave the network because they do not bear a valid source address.

We would argue that an organization which owns/operates a connection to the Internet and does not filter traffic is already in breach of its duty to protect its assets from misuse and abuse. The first two elements of a negligent cause of action have been met. All that is missing is a hacker to come in and use the organization's resources to hurt another. That incident will provide the causation and damages.

Patches

As Mr. Bruce Schneier has stated, the cycle of developing buggy software and then rushing to develop patches does not work. However, until the software development process becomes as rigorous and precise as, for example, engine manufacturing, the patch treadmill is the best the industry can offer. Working within this constraint, there is a great deal of debate over the process of obtaining and installing necessary patches for applications and operating systems. On one side are the proponents who feel that all patches should be applied immediately. On the other side are those who cite any number of patches in recent years that fixed one problem but created three more, and so they feel that patching should be deferred until the patch is deemed safe and stable. Regardless of which side of the 'patch war' you take, installing patches is one of the best things an organization can do to protect itself against automated attacks.

Personnel

The personnel issue is a sticky – and expensive – wicket for most organizations. System and network administrators are often overworked because their employers cannot or will not hire additional personnel. In this situation, the system administrators must prioritize their tasks, and simply keeping everything running may fill 100% of their time. How many system administrators are enough? There is no clear formula like "one SA for every fifty accountants", so the needs and structure of the organization must be used to determine a suitable staffing level. In most cases, however, having only one person to cover any particular task is not a good idea: if only one person is on staff, what if he becomes ill or goes on vacation? Has the organization made arrangements to provide coverage for this employee's duties? Beyond the number of personnel, the roles of the individuals are quite important. Can any named defendant identify who exactly is responsible for security? Is this role documented?

This brings us to the topic of *due diligence*. In the area of network security, as everywhere else, due diligence is not a fixed point: it is a sliding scale. There is no magical line separating negligent from responsible, where an incremental move in a certain direction will cause a state change. Here are some clear-cut examples to demonstrate both sides:

- Negligent: a default operating system installation, with no firewall or patches, on a T1
- Responsible: a hardened operating system with post installation changes behind a robust firewall

Scott's Assessment of Due Diligence

This section is so named because the position taken in this section is Scott's; he is not speaking for any other personnel or organizations.

This section currently applies only to businesses, although it may eventually apply to individuals. It defines a minimum standard of conduct for a very important reason: placing a system on the Internet, where it can potentially affect the systems of others, entails a certain level of organizational responsibility

Due Diligence Statement 1 of 2

Installation of security-related patches, when potential exists to harm a third party:

These patches should be installed no later than ten (10) calendar days after release of the patch by the vendor.

Many individuals will think that this interval is too short or (probably) far too short. (There is at least one person who thinks it is too long.) Many of the reasons given for this include the fact that there are simply not enough personnel to handle the work. However, going back to the issue of organizational responsibility, the owner of the network has a duty to make sure the network is as safe as it can reasonably be made. This duty includes having access to the resources - i.e. personnel and equipment - needed to test and apply patches in a timely fashion.

Due Diligence Statement 2 of 2

Egress filtering should be enabled on the network perimeter.

As mentioned earlier, there is no legitimate business purpose for spoofed packets, and simple set of rules on the firewall or border router can block this traffic before it affects someone else. These rules could likely remain static and still do the job, which is as close as anything can get to "set it and forget it" in this arena.

This article has covered negligence and due diligence, but what happens if an organization *is* negligent? The results of negligence can vary widely:

- No incident occurs - business as usual
- Mild incident occurs - inconvenience
- Serious incident occurs - substantial financial damage

- Most serious incident occurs - life is lost

The DDoS attacks would be classified by most folks as a serious incident; eBay, CNN, and Yahoo! would almost certainly agree. However, a broad application of egress filtering could have mitigated the damage.

What about sites with sensitive information?

The value of information is generally subjective. If a company's trade secret – plans for a new and improved Super-Widget, for example – were stolen or corrupted, the company would have a difficult time quantifying the amount of loss: no one can predict exactly how much money would have been made through the sales of the new product.

What about sites with large amounts of bandwidth available?

Sites with large amounts of available bandwidth – or “big pipes” – are often targets of attacks because the fast network connection can facilitate a number of nefarious activities. The potential for damage can be more easily reckoned in this case: an OC-3 can flood a T-1, but not vice-versa. One may argue the point that sites with big pipes have a slightly greater responsibility to secure their networks, similar to the way that a tractor-trailer driver needs to pay more attention to the function and condition of his brakes than a person on a bicycle: if the tractor-trailer goes out of control, the potential for damage is much greater.

What about sites that offer Service Level Agreements (SLA)?

Any reasonable SLA must account for the fact that the systems require maintenance. One way around downtime is to have a load-balancing cluster of machines, and take down one at a time to install patches and so forth. The choice here is either to allocate a small amount of time for maintenance now, or to allocate a potentially much larger amount of time later when something untoward happens, be it an intrusion or a software bug that corrupts database tables.

Back to the Group Presentation

Questions to ponder:

- Should the plaintiffs go after the ISPs? Why or why not?
- Does anything change if Jane G.'s employer is an ISP?
- Evaluate the potential for damages; how much prevention could this amount have purchased?

Conclusion

Case law is just starting on these issues; to date no far-reaching precedents have been set. Most organizations will want to avoid being on either side of such a landmark case.

Please use this article to speak to your in-house counsel or other legal professional in order to dedicate more resources to the cause.