Best Practices for Distributed Denial of Service (DDoS) Attacks

featuring Rachel Kartch as Interviewed by Suzanne Miller

Suzanne Miller: Welcome to the SEI Podcast Series, a production of the Carnegie Mellon University Software Engineering Institute. The SEI is a federally funded research and development center sponsored by the U.S. Department of Defense and operated by Carnegie Mellon University. Today's podcast is posted on the SEI website at sei.cmu.edu/podcasts.

My name is Suzanne Miller. I am a principal researcher here in the Software Solutions Division of the SEI. I am very pleased today to introduce you to Rachel Kartch, who is a network analyst team lead in our CERT Division of the SEI.

Welcome Rachel, thank you for joining us today.

Rachel Kartch: Thank you for inviting me.

Suzanne: We are going to be talking about DDoS, Distributed Denial of Service, today. Tell us a little bit about what in your background led you to working in this network analysis area that sort of leads us into this.

Rachel: I am currently a team lead for a group of analysts in the CERT Division. We have a lot of different job duties, but one of the primary things that we do is we look at network traffic on our sponsor networks, and we are trying to understand what is happening: Is it normal? Is it abnormal? How do we know if we are looking at something abnormal, and what possible threats or bad conditions might be on the network?

Prior to coming to CERT, I actually worked as a network engineer for about 15 years. I worked in a number of different environments. I have worked at couple of ISPs [Internet service providers]. I worked for a hosting company. I worked for a very large global financial corporation, and I worked as a consultant. Just about every one of those jobs, first of all, obviously, there was a lot of traffic analysis involved when we are trying to figure out how to engineer the network.

But getting to the topic of today's podcast, I did sometimes find myself, one way or another, involved in a DDoS event. Especially when you are working for an ISP, you can find yourself responsible for the pipes for the attacker or the attacked part. You need to understand how to respond. If your customer is under attack, how do respond? If your customer is attacking someone else…? For better or for worse, I have had some of that experience.

Suzanne: I had not thought about being on the attacker side of things, so that is an interesting side of it.

Rachel: Yes, in some ways, honestly, that is a little bit worse, because you are not just in a defensive mode then, but you have got to deal with a bad customer.

Suzanne: Let's talk about DDoS and the recent event that makes this very timely for many of our viewers. Almost the whole East Coast was involved in this denial-of-service attack. It was one of those things, you mentioned before the podcast, one of the apps on your phone did not work, and that was kind of the first clue that there was something going on. People's Internet service was disrupted at the home level and businesses. Dyn was apparently the focus of the attack, but it was certainly a very well distributed in terms of reaching really pretty much all layers of the economy and into the consumers.

How did that attack wreak such havoc? What about that that made it possible for it to do so much damage in that short time frame, from your perspective?

Rachel: The reason why that attack in particular was so bad, I mean, we can just say objectively it was a bad attack. It was large and it did a lot of damage to Dyn's service. But the reason everybody cared about it, and the reason everybody was talking about it, is because the attack was actually targeting somebody who provides services that make a lot of parts of the Internet work. If Dyn was just some industrial [organization] building widgets somewhere, and they came under this attack, people would not have been talking about it all day. It would not have been on CNN.

The reason it mattered is because Dyn provides the underlying services that made so many major websites and major applications function. That is honestly very scary. If you read Bruce Schneier for example, in his blog just a month or two ago had made a comment about the fact that he was seeing that there had been a number of attacks that were happening over the past year where people seem to be targeting—he did not go into specifics, so I can't name anybody—but they were targeting the providers of Internet infrastructure. So you can imagine that an attack on Dyn would fall into that category. What is scary, there, is they attacked Dyn, but they took out lots of other stuff as well. Twitter was impacted. Reddit was impacted. My ability to use the Starbucks app was impacted, which is incredibly important.

Suzanne: Especially at 6:30 in the morning when all of us need that. That is one type of DDoS attack where you attack some element of the infrastructure. Just to give our viewers a little bit of idea [of] the breadth, what are some of the other categories of DDoS attacks and their likely impacts?

Rachel: There is no one standard way to classify DDoS attacks, but one of the systems that a lot of people who talk about this stuff use, is they divide them up into volumetric, protocol, and application attacks.

Volumetric attack, and I looked for numbers on this, so what percentage of attacks are volumetric? Everybody has got a different number, but what everybody agrees on is, that volumetric attacks for now and what we have been tracking in the past, volumetric attacks are the most common type. Definitely over 50 percent are volumetric.

Volumetric is what it sounds like, large volumes of traffic. It can be very crude. All you are trying to do is fill somebody's Internet pipe or just knock down their servers just because of a huge flood of traffic coming to them. But typically we are talking about filling pipes, so that is a volumetric attack.

A recent example of a volumetric attack, a very well-known one was in September, the attack on Brian Kreb's website. From what we know about that attack, from the details that have been published, that sounds like that was largely just a volumetric attack. In fact it was one of the largest that has ever been seen. It was something like 620 gigabits-per-second of traffic that was being thrown at him. It was something like twice as large as the next largest attack that had been seen. So it was huge. That is one category.

The second category of attacks is protocol attacks. These are attacks that target weaknesses or soft spots in how a protocol is implemented. An example of that would be a SYN flood. If you are familiar with the TCP three-way handshake, so the way the three-way handshake works, and this is what establishes a TCP connection, the client sends a SYN to the server saying, *Hi, I'm here and I would like to talk to you*. The server sends back a SYN-ACK, so it is saying, *I acknowledge you and I am sending a SYN back*. Then the client sends an ACK, three-way handshake is complete. Now you have got a TCP connection.

What happens with a SYN flood is the attacker will send many, many, many SYNs to the server. Those SYNs could all be coming from a single host with spoofed source IP addresses, so it looks like it is coming from thousands or millions of hosts. Or, it could actually be coming from thousands or millions of compromised hosts.

You hear people talk about botnets. If you have got a whole bunch of infected devices, all of those devices are now being told by the command-and-control server to send a SYN to the target

IP address. The server has now received a whole bunch of SYNs. What the server, if it is vulnerable to a SYN flood, what it will do is it will bind resources. For each one of those SYN requests, it is like, *OK, I've got a request for an open connection. I am going to bind these resources so that I can prepare to complete this connection.*

Suzanne: It is like making a reservation.

Rachel: It is like, *Hey, I've got a whole bunch of people who want to connect to me right now, I'm going to bind these resources*. The server sends back the SYN-ACK to all the IP addresses that sent SYNs. Now if those addresses are spoofed, the server is sending SYN-ACKs to hosts that never sent the SYN. The proper function of TCP is the host is going to drop it; the host that receives that SYN-ACK is like, *I don't know what this is, and I don't care, so I am going to drop it.* Even if it was not spoofed, the malicious hosts are not going to reply to the SYN-ACK with their ACK. They are going to ignore it.

The server now has all these bound resources, all these SYNs came in, it's sending back the SYN-ACKs, nobody's helping it out, nobody is completing that three-way handshake. At this point, if the attacker has done their job right, that server … all the resources are used up, it cannot accept any more SYNs. It cannot accept any more TCP connection requests. Therefore, you have now completed your denial of service because the servers unable to accept any new connections.

Suzanne: This is like having a restaurant with a bunch of empty tables because nobody showed up for their reservation. And nobody knows why. Now when you call them on their cellphone, their cellphone number is not right. You do not know why you do not have anybody in your restaurant.

Rachel: Then the third type of attack is an application attack, and this is an attack that targets weaknesses or soft spots in how an application functions. Frequently, but not always, when we talk about an application attack we are talking about an attack on a web server. We are talking about something involving HTTP, not always but many times. A good example of an HTTP attack is called Slowloris. The way Slowloris works, in some ways, it is a little similar to how a SYN flood works in the sense that the attacker is going to send an HTTP request to the web server, but it is not going to complete the request. It is going to start the request, and then very, very slowly, hence the name, on some periodic basis it is going to add another header to that request, but it is never actually going to complete it. Just like with a SYN flood, if the web server is vulnerable to Slowloris, it is going to be holding connections open. It is going to receiving a lot of these very, very slow connection requests. The connection is never complete. The next thing you know the web server connection table is full; it can't accept any new requests. So, as far as the world is concerned, your website is down. So that is the Slowloris attack.

The Dyn attack that we were talking about, from what we know of it, and we do not have complete information—they have released some of the information publicly about the nature of the traffic but not all of it—from what we know of that attack, it was at least partially an application attack, but in this case the attack was on DNS [domain name server].

It was on the DNS application and not HTTP, because part of what happened with that attack was that the attacker was sending what—here is a few different names for this—some people call it a random subdomain attack, which means that they were sending what looked like legitimate DNS requests, but they were for random nonexistent subdomains of domains that Dyn was hosting.

You just sort of prepend maybe a bunch of garbage characters to a legitimate domain name, send that request to Dyn. Dyn does not know that that is attack traffic. *Hey, this looks like a real DNS request to me, I better reply to it*. Send millions of those at once, and you can overwhelm any server.

Suzanne: I think the thing that, as somebody who does not work in this all the time, that gets me is that these are not complicated methods when it comes right down to it. OK, make me feel better now. How do we protect … do we know how to protect against these? What are the methods that prove promising? I will not say we have solved this, obviously, we haven't.

Rachel: I will make you feel worse first.

Suzanne: Okay, go ahead, make me feel worse.

Rachel: Something that people will ask me is, *How can I keep somebody from attacking me?* The answer is, go off the Internet. If you want to prevent somebody from trying to attack you, unplug your website and go home, and do not ever check your email, and do not worry about it.

The good news is that, I will not say this is a solved problem, but the good news is that there are a lot of tools available so that people can protect themselves at least from being completely overwhelmed, or protect themselves from being completely out of business.

Just like with any other kind of disaster that you would imagine happening to your business, planning is really key. You cannot tell yourself, *Well, if I am under attack I have got some really smart people, and we are going to figure it out*. You need to figure out in advance what you are going to do.

The first thing that you can do is look at your architecture, and by architecture, I mean a lot of different things. There is physical architecture. Look at how your network is put together: *Do you have only a single Internet connection? Do you have multiple connections, but they are all to*

the same provider? Are your servers all in a single data center? Do you have your servers geographically dispersed, and if so, how much?

All of those things are important, but in addition, going back to the Dyn event, some of the Dyn customers who were really impacted did not have secondary DNS providers. Dyn was their only provider. I think they counted on the fact that, *Well, Dyn has a lot of resiliency in their architecture. Dyn can protect themselves, so we are just going to use them.*

You need to look at even things like your DNS, and you have to constantly be looking at the worst-case scenario. What happens if our DNS provider is under attack? What happens if our data center blows up? Both looking for just what happens if something goes down in general, as well as *What happens if we are under attack?*

You want to look at the architecture first. If you see places where you have got bottlenecks, if you see places where you are maybe presenting a very rich target because you have got so many things in one easy-to-find location, you want to look at making your architecture a little bit more resilient, so that is number one.

The second thing, and here is some good news, there are certain really well-known types of attack that can be mitigated with network hardware. Network firewalls, web application firewalls, load balancers, many of those today have either a checkbox that you can check that says, *Enable SYN flood protection*, or has thresholds that you can configure and you can change to say, *How many half-open TCP connections will I allow open*? Obviously if you lower that number down—or if you lower how long you allow for a TCP connection to complete before you just time it out, you lower those—obviously, you are going to make yourself more resistant to SYN floods because you are going to be limiting the number of open connections they can create and how long those connections can hang around.

Similarly, for the application attacks, and in particular for something like a web application firewall, the manufacturers of these devices know the signatures of these type of attacks. They know things that you can set and that you can adjust that is going to keep your web server from holding all these connections open. Very commonly, a load balancer or web application firewall will not send an HTTP request back to the server until the request is completed. So, all those partially completed connections, they are never touching the server because the network device knows I am not sending this back there until I see it is a legitimate, completed request.

Then the other thing, in terms of network hardware, is there is purpose-built DDoS mitigation hardware that is designed to protect you from some of these types of attacks.

The downside of this is the network hardware is not going to able to help you too much if you are under a volumetric attack. Because the volumetric attack, again, somebody is sending 620 gigabits per second of traffic to your network, if your Internet connection is not 630 gigabits.

Suzanne: Mine isn't.

Rachel: Most people's are not. If your network connection is not large enough to handle the flood of traffic, it doesn't matter what kind of hardware you have got sitting at your front door that you are going to try to do something special to mitigate it, because your internet connection is just completely flooded. There is just nothing you can do. So, you are not going to be able to deploy network hardware to foil large volumetric attacks.

That gets us to the third suggestion that I have for things you can do. I will admit that this suggestion is not for everyone. There is only going to be certain organizations that can really consider this, but you can consider just scaling up your bandwidth really big. There are some organizations who we work with where this is one of the approaches that they are taking. They realize they are going to be potentially under attack, they have the money to be able to do this.

They are not necessarily going to be able to buy pipes big enough to foil the very largest attacks, but they can buy pipes big enough to foil maybe a less ambitious volumetric attack. They realize, *Hey, it is just easiest for us to just have really big pipes*. We have the money. We are willing to pay for this bandwidth even if we are not using it most of the time because it is that important to us that we protect ourselves.

Suzanne: This is a risk-benefit kind of a situation. When you were speaking earlier I was thinking that we have a lot of security risk analysis techniques that help people to understand this. The SEI has worked in that arena for a long time, and so this is just one more case where *If you have done your homework you can have business continuity* is essentially what we are trying to help people to achieve.

What sounds really important is looking at the entire chain. You cannot just look at your network. You cannot just look at your bandwidth. You have to look at your suppliers and the supply chain management aspect.

I have to imagine that a lot of businesses that are not of the really large size probably do not even know if they are supported by Dyn, or they are supported by this DNS provider or that ISP. Well, they know their ISP, but there has got to be gaps in their knowledge. How does the little guy deal with some of these issues that even though the method may be simple the solution is a little bit complex?

Rachel: That gets me to my fourth suggestion for how can you plan and try to make yourself less vulnerable, which is outsource. Honestly, this is going to be probably the best solution for all but the largest organizations. When I say *outsource*, I mean a few different things. Number one, everybody should be talking to their ISP and finding out what the ISPs' procedures are if you are under attack.

First of all, what can you do? [You can ask] *Hey, ISP, what will you do to help me if I am under attack?* Some ISPs, actually some ISPs have separate paid products and DDoS protection, but every ISP should be able to tell you *If you report that you are under DDoS attack, here are the things we can do*. You want to know not only what can they do for you, but what information do you need to provide to them when you are under attack. You want to find that out in advance so you are not like, calling them in a panic and saying, *We are under attack, and we have no idea what to do*. So, number one, talk to your ISP. Find out what they can do for you.

The second thing is there are companies that specialize in DDoS mitigation, what they call in the cloud. Cloud scrubbing is what many of these services are referred to. This is really the best way I think that we know right now for companies that are under large volumetric attacks to be able to do something about mitigating it. You are either going to have your ISP block it, or you can work with one of these cloud scrubbing services. You hopefully contracted with them in advance so that you are not suddenly in a panic, doing your shopping while you are under attack.

You contracted with them in advance, so when you are under attack possibly they have already detected it, depending upon what kind of monitoring you are doing, or possibly you are calling them and saying, *Hey, guess what, we are under attack.*

Then what happens is your traffic—your inbound traffic—is routed to the cloud service provider. Typically, they have many locations. So it may be routed to several different locations. It may just be routed to the one nearest you. It may be routed to whatever is closest to the source of the traffic. Either way, your inbound traffic gets routed to one of these cloud scrubbing centers. They do what is called scrubbing the traffic. They try to identify the malicious traffic, drop that out, and then send you the good traffic. You still hopefully will be getting your traffic but with the malicious traffic removed from it.

Suzanne: So, if Starbucks had that kind of a thing going on, your Starbuck application would have been slow because you have this extra stuff, but it would not have been denied all together, as a very simple example.

Rachel: It might have been slow. These services are not perfect. As you can imagine, there is effort involved in trying to identify which is the malicious traffic and which is the good traffic. We do know that depending on the nature of the attack traffic, some types of attack traffic are

easier to identify as attack and drop than others. There are other types where it is harder to tell, *Hmm, is this legitimate, or is this not?*

If we look again at the Dyn incident, those DNS requests may have looked legitimate. So how can you tell? I mean, that is Dyn's business. They are doing DNS. So how can you tell if any given request is part of attack traffic? Or, it is just one of the many legitimate retries that they were seeing as all the people who were trying to use their Starbuck app or get to Twitter or whatever?

Suzanne: *I will just reboot my system, and then it will be fine.*

Rachel: Yes. As people continually try to get to these sites that added to the traffic. The challenge, therefore, of a scrubbing service would be figuring out *which of this is legitimate and which of this is not and dropping the bad stuff?*

The other way that you would look at outsourcing would be—and this in particular is if your web presence is a huge part of your business—some companies just have a website and it is like, *Here is what we do*, but it doesn't really matter. For other companies, their website is their business. Again, think about someplace like Twitter. That is their business. People need to be able to get to Twitter. In that case, you could look at hosting your website or your application, or whatever it is at a hosting company where DDoS protection a mitigation is part of the service.

Some well-known brand names out there, you can speak to them and find out what do you do if my website is under attack, how do you protect me? Some of them can do things like dynamically spin up plenty of additional instances of the website in order to try to absorb the volume. They may provide scrubbing in their data center. There are a lot of different options, but again, look at hosting your stuff. Especially if you are a small or medium business, and your web presence is very important, look at hosting it someplace where DDoS protection is part of the service.

Suzanne: The interesting thing about this is the regional focus of this. I actually happened to be in California at the time that this occurred, so my Starbucks app was fine. But my brother works in the security domain, so he knew about this six hours before we saw it on any of the local news or anything like that. That is another aspect to this, is that you can have physical differences in where you are.

Some of these ideas of the cloud scrubbing that occurs in multiple places, really, it resonates with me because we were fine. There was not anything going on in our area, and yet all of these people at home were impacted in negative ways that you just did not even think about until you really start to understand what is happening here.

Rachel: And of course it happened that way because of the hosts the attackers chose to hit. Next week it could easily be California. You never know.

Suzanne: Hopefully I'll be here when that happens. You cannot always arrange your travel so that you are in the place you're not being hit. That is the truth of it. These attackers are going to continue to try bigger and better (from their viewpoint) attacks. At some point, I would expect we will see something that tries to attack across an even broader spectrum.

This podcast and other things that we do at the SEI are all trying to make the public more aware, and make organizations more aware of the security issues, the security risks, opportunities for mitigating those, planning for those. What are some the places that you would recommend on the SEI's web sphere that people look for this kind of information to help them get a better handle on—not just denial of service—but other kinds of risks that impact them or could impact them in terms of their security?

Rachel: Absolutely. First of all, what I would recommend is people take a look on the SEI website at some of the incident management resources that we have. We have a lot of great resources from our CSIRT [Computer Security Incident Response Teams] operations folks. Maybe you have got a company, and you do not even have any sort of incident management group at all. You do not even know where to start. We have resources for where do you start? How do you build this capability?

Suzanne: And training.

Rachel: Exactly, and training as well. Like I said, I wish that I could reassure you that DDoS is not a scary thing, and you can keep it from happening. Considering that you can't, the best thing that you can do is prepare yourself. Our incident management resources on the SEI website are really, really good to help people just prepare—not just for DDos—but for all sorts of other problems that can happen.

The other thing that I think would be a really, really useful tool is taking a look at our SiLK toolset that is available through the SEI website. The SiLK toolset includes free software that you can use to set up a network flow collection and analysis infrastructure. What that will let you do is, first of all, you can get a better understanding of your normal traffic profile. That is key for if you do find yourself under attack, and you are trying to figure out, *What is this attack?* If you know what your traffic normally looks like, it will be easier for you. You can look at *What is going on right now?* and you can see how it is different. *Is this notably different from what we normally see?* Even just in terms of the volume, that will help you figure out if you are under a volumetric attack, but also, *Is it significantly different protocols than we normally see?* That

collection infrastructure can really help you first of all to figure out if you are under attack and figure out what kind of attack.

Also, if you do have to call your ISP, or you do have to call your scrubbing company, you can let them know, *Here is what we are seeing. Here is the type of traffic*. That information is going to be really, really useful for anybody who is trying to help you to trace the attack back to its source, to block it, to scrub the traffic.

Again, the SiLK tools will be really, really helpful for you just to start understanding what your network traffic looks like, and as a resource that you can use if you do think that you might be under attack.

Suzanne: DDoS is one type of network analysis thing, what are some of the areas that your team is focusing on right now outside of DDoS that you think will eventually contribute to knowledge in this area?

Rachel: One of the big things that not just my team, but a lot of people in this field, are interested in right now is getting away from signature-based analytics and moving more towards behavior-based analytics or anomaly-based analytics. The idea being, it is fairly easy if there is a particular bad type of traffic or an attack, if you know exactly what it looks like, obviously you can develop a signature to detect it.

If it is like, well, you see three packets that look like this, and then you see two packets that look like this, and then it is coming from this IP address, it is very, very easy to design a signature to say, *Oh, that is a bad thing, because I know exactly what it looks like*. It is a lot harder to find the bad thing when you do not exactly know what the bad thing looks like. All you are really looking for is, *Hmm, did anything strange happen? Did anything that is out of the ordinary* [happen]*?* That is not as easy as it sounds because, of course, there is a whole lot of questions you have to ask yourself in order to define *strange*.

So, my team works a lot these days on trying to build analytics that are based on known signatures but are based on anomalies, that are based on changes in behavior. We are also just now launching some work where we are going to bring in some other folks and try to apply machine learning to this problem as well.

I think everybody is working on machine learning as well. I think it is really going to help us, because so often when we speak with network analysts, when we speak with the people who are out on the front lines doing this job, there is a lot of, *I do not know what it is is, but I will know it when I see it*. There is a lot of instinctive kind of knowledge stuff that is buried in the brain, and you cannot quite articulate why you know what you know. We do find that sometimes computers can make that task very easy. They can articulate it, and they can tell you…

Suzanne: Yes, and they can systematize it.

Rachel: Exactly. So we are excited about the possibilities of machine learning to help us to articulate those things that we know it when we see it.

Suzanne: I want to thank you very much for this very timely discussion. I know there are lots of people—I am not the only one out there that is worried about what is this going to do to me. But even as, not as an individual, as an organization it is a much bigger impact. I really appreciate your insights. I know I learned a lot about what to watch for. I have new questions to ask to my ISP, and I really thank you for that.

We will be taking care of giving out all the resource links that you have talked about, and any others that are germane to this will be up on the transcript for the podcast. [View a recent blog post by Kartch on DDoS Attack prevention and response]. So, we won't be giving URLs here.

I just want to thank our viewers today. I hope you learned some things about how to protect yourself from this imminent, and recent kind of thing. Start by looking at this podcast and the transcript. That will be available as it always is at sei.cmu.edu/podcasts, or go to the Carnegie Mellon University's iTunes U site it will also be there.

Again, I want to thank all of our viewers today, and as always, if you have any questions please do not hesitate to send an email to info@sei.cmu.edu. Thank you for viewing.