WHITE PAPER

ixia

Cyber Range: Improving Network Defense and Security Readiness

Real-World Attack Scenarios for Cyber Security Training

915-6729-01 Rev. A, August 2014

Table of Contents

Executive Summary

To harden the resiliency of vital government, military, and commercial infrastructures, organizations are deploying cyber ranges, vast test beds that allow war games and simulations aimed at strengthening cyber-security skills and defenses. Traditional cyber ranges require significant, costly investments in hardware and personnel — and even then cannot scale effectively to address today's growing network traffic volume and ever-more complex attack vectors.

But a cyber range is really nothing more than realistic environment that is used for cyber technology development and cyber warfare training. Cyber technology development applications might include network performance evaluation, application performance evaluation, security development and verification and product assessments and bake-off to name just a few. Training applications on the other hand include flag exercises, cyber competitions and red, blue, white team training exercises.

We want to show you how to achieve cyber range testing and training without the need for the quickly-outdated and vast test beds of past cyber ranges. In this paper, we'll discuss architectures and scenarios associated with typical cyber range deployments and how you can achieve the most benefit out of your cyber range.

A cyber range is really nothing more than realistic environment that is used for cyber technology development and cyber warfare training.

What is a Cyber Range?

The word range in the term cyber range is military wording for a typical targeting or kinetic range where you can send troops to hone their fighting skills with a variety of realistic well-planned exercises that may include interactions with guns and ammunition, tanks, war planes, war ships, and so on. In that manner, the war fighter is able to train as they would fight.

Similarly, a cyber range provides your network and IT technology personnel with a realistic platform for training related to network attack and defense scenarios, mimicking real-world scenarios in the lab so that they too can train as they would fight. Cyber range training focuses on how to evaluate situations and apply the correct policy/configuration for specific real-life attack situations.

It's probably fairly obvious why militaries want to train their cyber assets. They need to ensure that their mission-critical war-fighting infrastructure performs adequately in a time of war. This holds true for all pertinent elements of wartime architecture, from the radio-equipped soldier on the ground, to command and control elements, to various service elements, and the often times extraordinarily complex networks in between.

It's not enough just to make sure that things work. They also have to be resilient against attack for an enemy who might be hell-bent on disrupting mission-critical infrastructure. Many companies are realizing a bit late in the game is that this is not applicable just to militaries and governments, but these scenarios are actually playing out every single day around the world, in everyday life, and in private industry. Even in modern "peace times", state-sponsored attacks are prevalent. Our connected world is under constant attack from criminal organizations, activists, and rogue nation states.

A cyber range enables training with real-world, physical and virtual services and networking, such as:

- IT Services:
  - Virtual machines
  - Microsoft servers
  - Microsoft server applications
  - Linux and UNIX servers
  - Database applications
  - Authentication services
- IP Networks:
  - Routers and switches
  - Security devices (firewall, IDS, IPS)
  - Application optimization (SLB, proxy, etc.)
- Critical Infrastructures:
  - Simulation of a real power-plant, water system, train system, airport, etc.
  - SCADA – a protocol used for critical systems communications, not only IP

Who Needs a Cyber Range?

Let's discuss who actually needs a cyber range. This is a really important question to ask and to answer.

Every single day, the news headlines scream out about new attacks on the financial industry, financial fraud, credit card fraud, identity theft, data leakage of corporate secrets, defense contractors being bombarded by penetration attempts, public power and water network intrusions, and politically- or ideologically-motivated cyber attacks. The simple fact is that no industry is immune to this new and prevalent attack culture.

If you have a significant online presence, an attack that causes your services to go down even for short periods of time can cause tremendous amounts of lost revenue. Perhaps reaching into the hundreds of thousands or even millions of dollars of loses for certain industries. If you were on a closed network serving public infrastructures such as a power grid, a water treatment plant, a hospital and a USB-introduced or mobile-phone-introduced piece of malware infects your systems, people could die.

If your business relies on virtual private network, a VPN technology for your day-to-day operations for a diverse workforce across a nation or around the world, and a network attack scenario degrades or brings down elements of your VPN architecture, you'll wish you had trained for the eventuality in a cyber range environment before it hit.

If you have a significant online presence, an attack that causes your services to go down even for short periods of time can cause tremendous amounts of lost revenue.

If you're an oil company whose critical databases are adversely modified or destroyed by a cyber attack affecting your ability to ship your product or to drill, then you will absolutely wish that you had analyzed that eventuality in a cyber range environment to have been able to effectively mitigate it ahead of time.

If you're a service provider and you can't handle or mitigate a massive-scale botnet aimed at your network by heavy-handed nation, state, or political entity, your customers, their service level agreements (SLAs), and their business bottom-line are impacted.

If you are a network equipment manufacturer (NEM) and your hardware fails or performs exceedingly poorly under a targeted attack, not only will your customers be in an uproar, but you will almost certainly be held accountable and that will pave the way for more competition to emerge in the shadows and take business right out from under you.

Who needs a cyber range? Well, it's pretty clear that practically every organization needs one. No one is immune to these issues and breaches. Every single day it seems attacks are being divulged in organizations under constant attack.

Is it Really That Bad?

We continually hear of the constant barrage of new security attacks and the prevalence of security holes. It just seems that there's no end in sight. The nature of most of the connected world makes securing our networks and data in extremely difficult problem to tackle.

As good as the largest software and hardware development corporations are, there seems to be no end to the number of discovered security holes, with security vulnerabilities being uncovered daily across the globe for network systems IT and services.

There seems to be no end in sight. These problems are not going away. They consistently seem to get worse for the simple reason that network infrastructure and computing complexity is growing over time. Where complexity exists, security problems will exist. We live with a very complex computing and networking infrastructure all around us. Where complexity flourishes, security problems will flourish – and the data proves this out.

According to the Secunia Vulnerability Review 2014, the trend for security vulnerabilities for the top most used software programs has gone up significantly over the course of the last five years (see figure). It's clear that the best software companies in the world have huge vulnerability problems and these problems are just not going away.

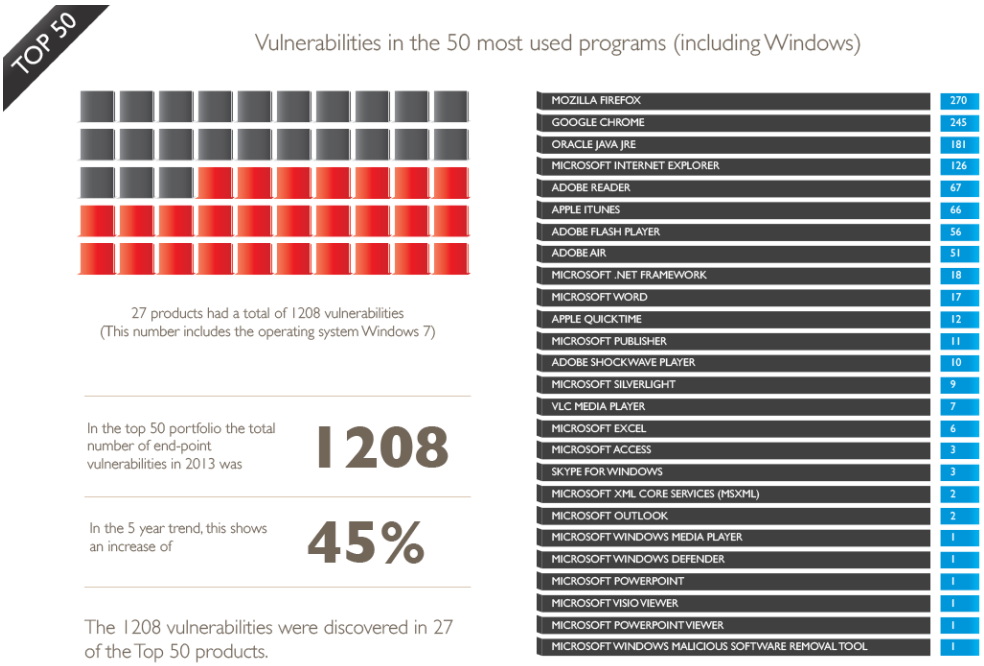

Source: http://secunia.com/vulnerability-review/vulnerability_update_top50.html

How Can We Tackle These Security Issues?

There's been an interesting set of discussions going on for years with competing viewpoints as to how to solve these security issues in both the short term and the long term. There are discussions about the benefits of open architectures, benefits of closed architectures, a brand new supposedly "secure from the ground up" network and competing architectures, use of firewall, intrusion and prevention systems, zero-day attack recognition, unified threat management systems (UMTS), and a plethora of other security architectures and discussions.

Those discussions are necessary, but most neglect something that has been very true since the beginning of recorded human history — education and training are what makes human beings good at something and provide a means for us to adapt to new surroundings. Even protégés must learn and this is where cyber ranges come into play. You cannot expect your IT organization to be able to defend against a complex cyber attack that occurs if you are not adequately training them. It's like sending someone who has never thrown a punch in their life into a fight against the championship boxer.

The novice doesn't stand a chance. He'll lose every single time. Yet that is precisely the situation that most organizations find themselves in today with regard to securing their critical infrastructure. That brings us the two obvious questions.

1. How do you architect and construct a realistic cyber range
2. How do you make most use of that range once it is constructed

Education and training are what makes human beings good at something and provide a means for us to adapt to new surroundings.

Architecting a Cyber Range

The cyber range can be architected in many ways but in general you can group them into three categories, physical, virtualized, and hybrid. Let's take a look at each range type, and their advantages and disadvantages.

Physical Cyber Range Architecture

In a full physical range, you duplicate your entire physical network infrastructure, your switches, routers, firewalls, servers, endpoints, etc. You use that duplicated entity for your training. This is great from a realistic perspective because you can't get it any more realistic than that because you're using basically everything that you're using in your range is real.

However, this approach has some significant disadvantages and these are really important:

- One of the biggest disadvantages is the often exorbitant financial and staffing outlay to replicate your typical complex live network environment.
- Another disadvantage is the setup and teardown time to craft new training scenarios in that environment.
- Another is physical management of the environment, ongoing operational cost considerations for things like cooling and power consumption of such a large complex physical range.
- Finally, and not be stated lightly, is the difficulty in cleaning a pure physical range after an exercise completes. This is important because many scenarios running on your range will involve complex attack vectors that can leave undesirable artifacts on your network after the fact and that's bad.

One of the biggest disadvantages of a physical cyber range architecture is the often exorbitant financial and staffing outlay to replicate your typical complex live network environment.

Virtual Cyber Range Architecture

In a purely virtual range, everything is simulated. Each component is emulated with virtual machines. This approach offers some distinct advantages. The initial capital cost and the ongoing operational cost are significantly less expensive than a full physical reconstructive range. The set up and teardown are typically simpler as well and are typically known quantities in terms of the time they take.

The range hardware becomes simpler and that's because everything is virtualized to run on relatively common off the shelf hardware and less expert human resources, or fewer expert human resources, are typically needed than for a fully physical notation. The attack artifacts are also more easily disposed of. For example, by reverting to known good snap shot of your virtualized infrastructure.

However, there are also severe disadvantages to this approach as well. Basically, what you gain in simplicity and lower costs, you pay for in performance. Purely virtual infrastructures will never perform to the same level as physical systems, regardless of what virtual vendors tell you and sell you. This can actually be a critically important issue in training as you fight. Since in a purely virtualized environment, you may not be able to model certain attacks scenarios that you will find in the real world, purely on the basis of physical performance constraints of virtual architectures.

A great example of this is realistic distributed denial-of-service (DDoS) attacks. In many cyber range environments, DDoS analysis mitigation by a response team can work only at low-scale because of the limits of virtualization. In the real world, DDoS attacks are some of the most prevalent attacks you could find and most effective. Quantity has a quality of its own.

Two more examples of why pure virtualization can cause issues are network throughput, which is always lower in a virtualized environment, and firewall IPS/IDS performance, which are greatly constrained in terms of performance when those elements are purely virtualized.

Yet, another issue is that in virtual environments, it is not always possible to mimic precisely what real physical elements can actually do. For example, creating a virtual warplane that can actually drop a precision bomb thousands of physical miles away is practically impossible. In a corporate environment, creating a virtual webcam to max the capabilities of high-end video surveillance gear can be equally as impossible. These are things that you may absolutely need to model correctly in your range.

We've looked at physical and purely virtual architectures. What about the third possibility? The hybrid cyber range architecture.

Hybrid Cyber Range Architecture

This third architecture type, the hybrid range, is the only range environment that deserves merit. In a hybrid environment, you use virtual elements when and where they make sense and then mix in physical elements when and where they make sense. For example, you can virtualized your desktop environments and your Windows, Linux, or other servers, but you would use physical device connections for your high-end video surveillance or certain other hardware elements that you're looking to protect such as perhaps network printers, Voice over IP (VoIP) phones or adapters, and probably most importantly your actual security equipment such as your firewall, IPS/IDS, etc.

In a hybrid environment, you use virtual elements when and where they make sense and then mix in physical elements when and where they make sense.

The reason is because attempting to virtualize those things can invalidate your range because it eliminates the realism needed to ensure that you train as you fight. For example, if you're using a high-end firewall in your production network but you resort to using an open source type firewall on a virtual machine in your cyber range, you aren't training as you fight. Not to mention that you won't be able to train on the effective techniques that your high-end gear provides you in many network attack scenarios, which you cannot replicate in a virtualized scenario.

Using virtualized network defense techniques in a cyber range environment also typically constrains the amount of traffic that can flow in those networks. This means that you won't be able to evaluate the proper functions need for an event such as a real-world, high-scale service attack.

Regardless of your topology, if you go with a hybrid solution, you'll still have some certain limitations in your range. But in general the hybrid solution provides the best mix of cost performance and scale towards the goal of providing the best environment possible to train your human cyber warrior assets and evaluate the resiliency of your network defenses, basically allowing yourself to train as you fight – which is your end goal.

Effective Use of Your Cyber Range

We know we want a hybrid environment and we have a general idea of what we want to virtualized and what we want to remain physical. Because the various components can be virtual or physical depending on the training scenario requirements, we won't cover the actual implementation of a cyber range in depth.

You may have a variety of physical and or virtual endpoint equipment on a couple of subnets connected to a few switches, with routing and firewall resources interspersed as is common on many topologies. Some of these switching and routing functions maybe virtualized or they may be physical, depending on the exercise requirements. The same thing goes for the firewalls.

A network demilitarized zone (DMZ) is an external enterprise service such as web hosting and mail services. This might be your publicly facing web and mail servers for example.

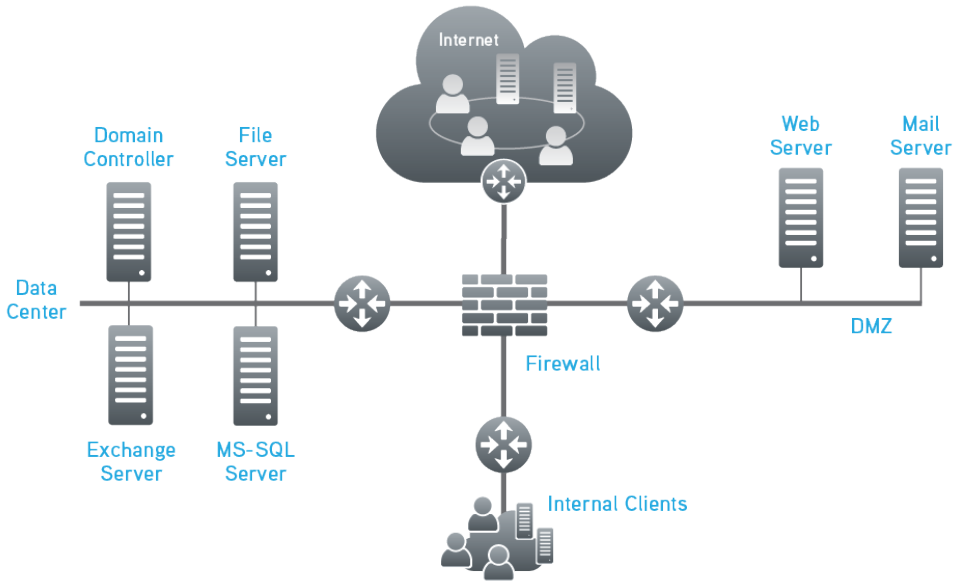

In a hybrid cyber range, the various components can be virtual or physical depending on the training scenario requirements.

The routed communication links could be anything as well: ISP connections, direct point-to-point cabling, satellite communications, cellular, or other. The point is that whatever you're modeling should be as realistic as needed to meet the constraints of the exercise and its goals. Maybe you've got a couple of servers hanging off of a network demilitarized zone (DMZ). A DMZ is an external enterprise service such as web hosting and mail services. This might be your publicly facing web and mail servers for example. They may also have special access through the firewall to other corporate resources and servers, making them very probably targets for attack by external networks.

The figure above shows a sample high-level cyber range topology that includes:

- Data-Center: Internal enterprise services (DC, DB, FileServ, Exchange etc.)
- DMZ: External enterprise services (web hosting, mail services etc.)

- From external clients:
  - Internet Users accessing DMZ
  - Internal Clients accessing from Internet via VPN
  - Clients from Internet trying to penetrate / harm the organization
  - DNS and Mail SYNC
- From internal Clients:
  - Internal clients accessing enterprise resources from the data center
  - Internal clients browsing the WWW
  - Surfing the Internet

How do you make effective use of the hybrid cyber range? This is actually a really important question. It is one thing to say you got the cyber range all built, but how do you actually use it?

Technology Development Assessment

Let's first talk about the technology development assessment case. At first glance, this case appears to be very straight forward. You're using the range to assess performance in a particular device or service under constraints that you set. Using the range provides you with a valid laboratory environment, mapping to real-world use cases that you can ensure your device, service, configuration, bug fix, etc. will work as expected once deployed to production. Same thing would go, if you're comparing multiple products in a product bake-off. Maybe a more interesting example would be if you're testing out a second physical video camera solution alongside an existing solution. You need to see how it behaves and compares to your existing infrastructure.

Use the range to assess performance under constraints that you set for a particular device, service, configuration, or bug fix, ensuring they will work as expected once deployed to production.

Even in this environment though there is a plethora of critical considerations. For example, how real is your range construct? Did you test the device's protocol behavior and network bandwidth in isolation or in the midst of realistic background network traffic? These things make a huge difference in the validity of your experiments on your cyber range.

It is important to note that realism is very necessary in a cyber range. If you don't have realistic scenarios that involve realistic background traffic and realistic security attacks, then your range is practically useless.

Red Team/Blue Team Cyber Warfare Training

In training scenarios, it is very easy to see how things can quickly get quite a bit more complicated. We use what we call red teams and blue teams to respectively attack and defend the network, servers, and applications as part of the cyber range, given rules set up by white team.

White teams are extremely important. They set the goals for the exercise. These could be red team-based goals, blue team-based goals, or both. They manage the training exercise, have full visibility into the exercise, and set the rules of engagement. The white team also ensures the constraints in which the training exercise will occur. They are kept aware by the red teams and the blue teams of their progress during the exercise.

For example, a red team may claim that they have subverted a particular network resource. The white team once informed will have a means to verify the truth of the claim. A blue team may claim a successful defense against some known or unknown attack vector. The white team can verify that as well. The white team can control various aspects of realistic background traffic during various stages of the exercise to challenge the red team, the blue team, or both. Without a white team to collectively control the exercise, your training results will be severely limited and you will not accomplish the goals of the exercise to train as you fight.

Let's look at relatively simple red/blue cyber scenario architecture:

The white team would construct the range and assign the resources to the various teams and functions.

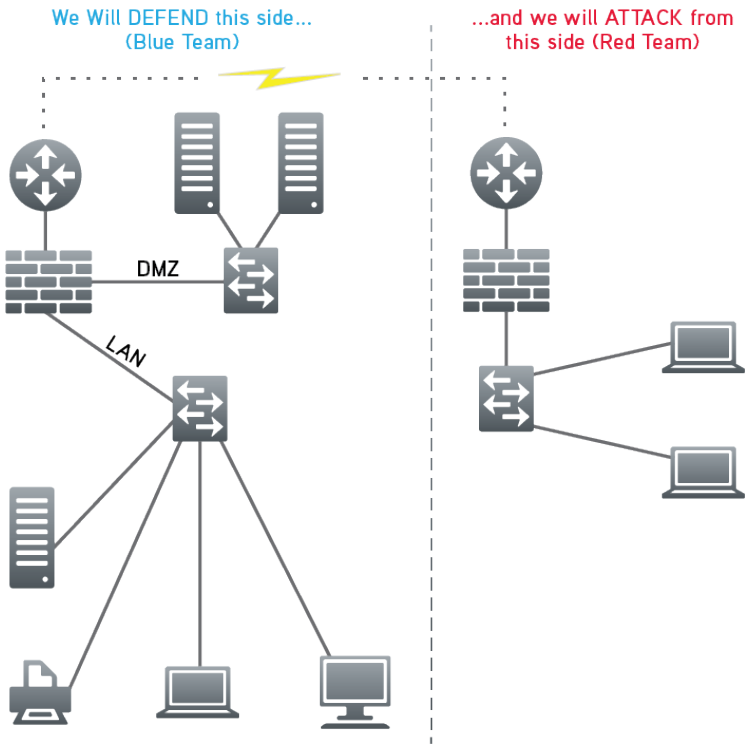

A simple red/blue cyber range scenario

Here you can see we've drawn an arbitrary line right down to the middle got the network. We're going to be defending one side and attacking from the other. The white team would construct the range and assign the resources to the various teams and functions. Then you commence your training or simply assess your existing defenses from a technology standpoint.

For example, let's say the goal of one exercise is to train your blue team forces to ensure that mission-critical video camera systems can still function under the presence of two simultaneous massive-scale DDoS attacks against firewall and DMZ infrastructure in your protective network.

You could assign your defense from IT blue team resources and assign the attack in red team forces. The red team forces can be internal resources if you have such expertise on staff. They could be third party personnel that you hire on a temporary basis to perform the attacks in your controlled range environment, operating according to the constraints set up by your white team.

The red team attacks according to whatever rules of engagement that the white team has configured. In this case, they're told to issue two particular large-scale DDoS attacks simultaneously. The blue team attempts defense and the white team monitors the proceedings.

At the end of the exercise, all the results are analyzed and lessons are learned by all involved, as directed by the white team. In this case, the blue team learns what they did well, what they did not do well, and where they are vulnerable. This allows them to perform suitable remediation and mitigation for future exercises, as well as to better secure their actual live mission-critical corporate infrastructure, which in this case, we defined as video surveillance.

In the previous diagram, we arbitrarily drew a line for the red team or blue team divider what represented what a white team may have setup as demarcation points for a particular exercise. That line can be anywhere the white team needs it to be to achieve the purpose of the particular cyber range scenario that you might be training for.

For example, you could decide on an exercise modeling an insider threat. Your white team has crafted a scenario where you have a security breach inside your to-be protected network (see following diagram). The goal is to see if the blue team can detect malicious activities coming from the insider threat that could result in critical data going across your external communication link to an enemy command and control server.

This usually requires the defenders having an internal firewall and/or network visibility solution in place to look for strange communications pathways that don't normally exists on the network, and then analysis of deep packet inspection (DPI) engines where advanced signature detection techniques separate benign behavior from anomalous and malicious behavior.

The red team attacks according to whatever rules of engagement that the white team has configured. In this case, they're told to issue two particular large-scale DDoS attacks simultaneously.

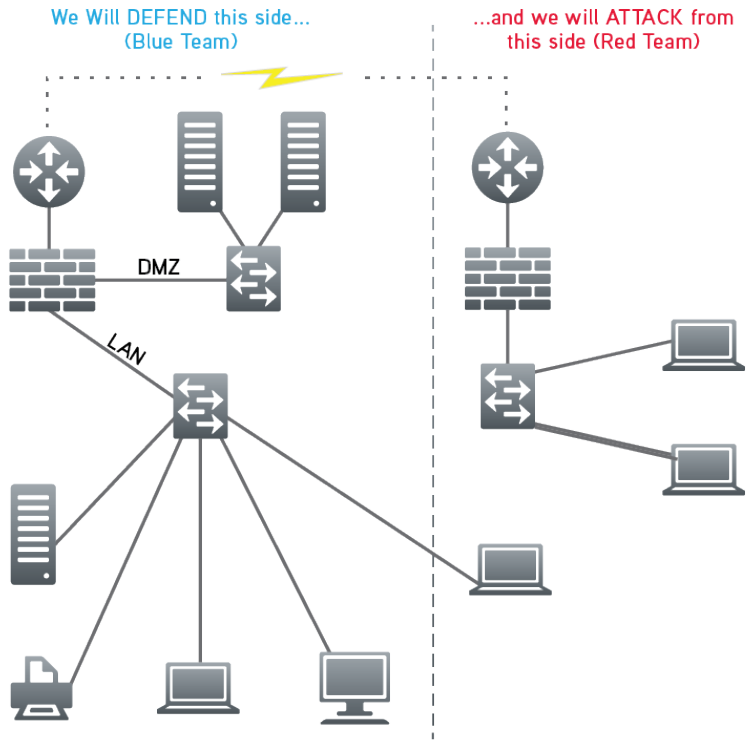

Insider threat and behavioral modeling

However, this can become exceedingly complex. A good range environment will allow you to perform just this type of what we call behavioral modeling training in the event that you find that you require it. Why is this so important? Because the insider threat model is precisely the model that applies if you've been compromised by what are termed Zero Day attacks. Most organizations have been compromised with these types of attacks at one point or another.

Zero Day Attacks

Zero Day attacks are the ultimate attack vectors because by definition what they refer to are attacks that have not yet been discovered by security vendor experts. They are propagated by malicious interests to gain control of what you thought were well-protected network resources. If these attacks are successful, and the vast majority of Zero Day attacks are, they often result in the installation of a variety of malicious software (malware) to take control of your network and your data, allowing your important data to be siphoned off to external entities. They can also cause a disruption or destruction of elements of your network or your data.

Zero Day attacks are the ultimate attack vectors because by definition what they refer to are attacks that have not yet been discovered by security vendor experts.

In some cases such malware can even lay in wait for a precise time when they're instructed to wreak further havoc. You may never know if or when that occurs. They can also delete themselves, morph themselves, and do all kinds of other nasty things in an environment. If you have any truly mission-critical infrastructure in your environment and you are not training for these types of insider threat scenarios using behavioral modeling techniques, it's only a matter of time before you reap the repercussions and they will be severe.

Creating and Managing Realism in the Next-Gen Cyber Range

There are many potential cyber range configurations and scenarios that a well-designed range can accomplish. We've gone through a several scenarios already. What you quickly realized is that each and every configuration requires different human elements, scale, and attack and defense paradigms. What you'll also find is that these things consume a large amount of human resources and set-up, preparation, and analysis time, which equate to time and money. This may cause you to constrain your exercises to the point that they become cumbersome and don't yield the results that you want.

You can easily get overwhelmed with these critical factors to the detriment of your exercises, thereby reducing your efficacy and prohibiting your assets from being able to truly train as they fight. For example, let's say you have a thousand people who use your computing resources in your organization. How are you going to model appropriate background traffic mixes that represent their day-to-day operations to scale without requiring all of them to actually participate in the range exercise itself?

How are you going to craft properly scaled, realistic attack vectors when you don't have the competing resources of an international botnet? How will you gain the expertise that you need to effectively attack your range constructs in the first place to assess your network defenses when you might not have that expertise. Will you rely on proven internal security personnel? Will you outsource it? If you do initially outsource it, you'll probably quickly realize the importance of developing your resources internally; how will you manage that without an explosion in personnel and computing requirements.

Automated Application and Security Testing Tools

Basically, the bottom line question is how do you manage realism? That's a really important question that most organizations neglect to answer at the beginning of creating or using a cyber range. The question comes up once they realize a need to ask it, which usually occurs after the first exercises fail miserably. It's all about realism here.

How do you manage these things? You could choose to outfit your range exercise with hundreds of thousands of people going about their daily business by using range computing resources. That obviously doesn't scale and nobody wants to tie that many human resources up for long periods of time. You could instead attempt to capture your live network traffic and attempt replay that on your range networks. That can help to a degree, but by itself it is inherently flawed because it replays only a snapshot in time of what you're live network was doing. It's highly demanding of computing stored resources required to play back that data which can be a tremendous of data in the terabytes in most cases.

Even worse, such environments are unrealistic since they do not have randomness as an innate quality. Your real world networks and the data they carry change minute to minute, hour to hour, and day to day. You simply cannot model that effectively and realistically by replaying snapshots in time of your production data. It does not work.

Ideally, you'd have some automated mechanism where you could quickly and efficiently pick and choose what you want in your background traffic and attack traffic to look like, making it as repeatable or as random as you want it to be for a given exercise.

Ideally, you'd have some automated mechanism where you could quickly and efficiently pick and choose what you want in your background traffic and attack traffic to look like, making it as repeatable or as random as you want it to be for a given exercise.

Similarly, if you're attempting to protect against a massive-scale DDoS attack, you need equipment that can simulate real-world botnets that can approach tens of thousands or more end points, simultaneously attacking your range network. You quickly realized that you can't build out that many physical or virtual endpoints. Ideally you'd have a tool that you could quickly and officially configure to look like as many external endpoints as your range exercise scenarios dictate, and then use those configured endpoints to attack your network with a built-in library of realistic security attacks.

Or you could do the smart thing and use a hybrid approach to both of these mechanisms and in doing so, create a true next-generation cyber range. This is precisely how the best and most effective cyber ranges operate. They didn't get there on day one. It took years of operation to figure this out. They started with people, but realized it didn't scale. Tried to replay network traffic, but realized that's not random enough and didn't scale.

They finally evolved to a hybrid approach of using automated tools whenever and wherever possible to ensure realism and scale and, for the highest granularity or certain crucial operations, using actual human beings sitting in seats controlling physical or virtual infrastructure just they would in our day to day job functions. The combination of those things allows your range to scale and maintain realism and that is a critical factor, probably the most critical, in a successful cyber range environment.

Now, what is immediately obvious is that without highly capable automated tools in place, you will never be able to scale your range to meet real-world conditions. Having such tools at your disposal provides the missing link for a comprehensive next-generation cyber range environment, bringing you ever closer to the ultimate goal of allowing your teams to train as they fight.

The following diagram shows you where an automated application and security testing tool would fit in to a hybrid cyber range. In the following figure, with the cyber range interconnections depicted as the dotted lines, you're able to configure the tool to generate any mix of good traffic and/or nefarious traffic that you'd like to best match your range scenario requirements. These tools perfectly complement your hybrid cyber range and your human assets and allows for finalizing the construction of a variety of both simple and highly advanced range scenarios.

The hallmark of a good cyber training exercise is realism. You may have a beautifully designed hybrid range, but if the right types of realistic background aren't flowing to match your real-world scenarios, and if you can't generate mono-realistic attack patterns at high scale, then you aren't training as you fight.

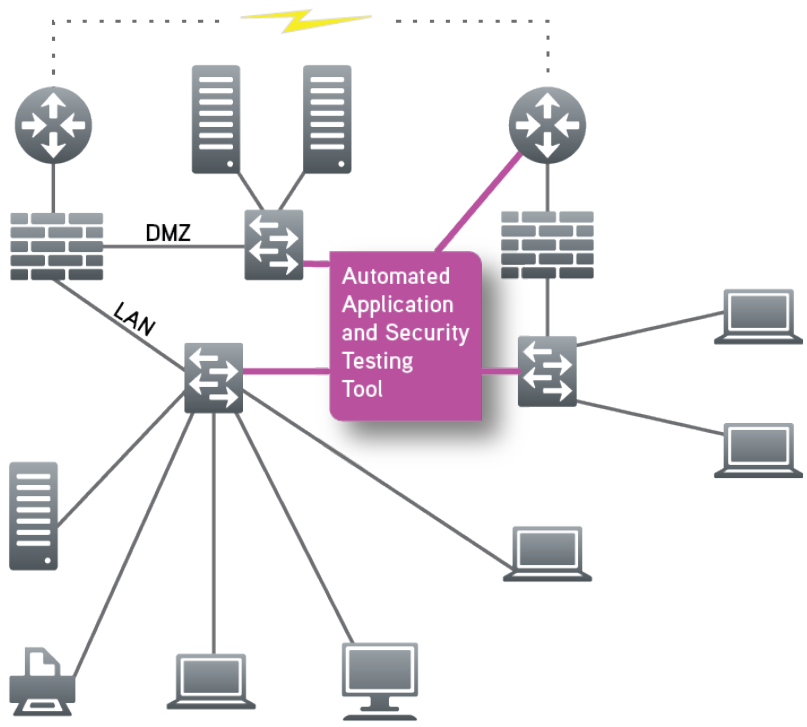

Application and security test tools simulate real-world application protocols and attack traffic at line rate

As previously discussed, the hallmark of a good cyber training exercise is realism. You may have a beautifully designed hybrid range, but if the right types of realistic background, your mission-critical network traffic, aren't flowing to match your real-world scenarios, and if you can't generate mono-realistic attack patterns at high scale, then you aren't training as you fight.

For example, it will be ridiculously easy for a defending blue team to detect the network attack if there was no meaningful or realistic background traffic flowing at all because at that point, the vast majority of network traffic would be known to be a formulation of the attacking red team.

The challenge of the blue team is often to discover the needle in the haystack, and for that you have to have a realistic haystack. For example, can your network defenses detect that one piece of malicious JavaScript in the presence of a slew of realistic randomized background traffic containing lots of good JavaScript? Can you detect that while your systems are constantly bombarded by a competent DDoS attacker? That's real world.

Selecting a Cyber Range Application and Security Test Tool

A make or break for your cyber range is having an easy-to-use, robust application and security test tool to simulate real-world application protocols and attack traffic. Because this is so important, you'll what to know what to look for when evaluating application and security test tools.

Usability is Key

Your tool should allow you to quickly craft a very complex set of range exercise requirements. Let's start this discussion with a real-world configuration example for using an application and security testing tool in your cyber range. Start with a lot of good background traffic and add two different DDoS attacks – in this case we'll use a UDP fragmentation attack and a DNS flood attack. Also add a variety of a live security attacks along with highly specialized attack invasions that target certain internal systems by traversing firewalls with some non-stereo techniques like IP fragmentation, Base64 encoding, or other nefarious activity.

Also include a sample of proprietary protocols that may be on the network that you can't generate natively, things like proprietary corporate protocols or even classified protocols in military environments, things that your tool can't generate on its own because perhaps they're top secret.

Lastly, add some intense network protocol fuzzing. Protocol fuzzing is a great attack strategy. It takes the traditional rules (as defined in various protocol specifications/RFCs) for a protocol and purposely circumvents the rules to see if strange and illegal data values can cause a security flaw to present itself.

Your tool must have a simple user interface to allow you to quickly do that level of configuration. It must simplify the complexity. You have a very complex set of scenarios with different attacks and realistic background traffic, but the configuration of them should be very simple, very straight forward, and very intuitive in your tool's GUI.

In this example, your tool should operate as both sides of the conversation, sending out the client side as well as acting as the server. Alternatively, it should operate as just the clients and hit your real servers and your network. This is really useful in certain obfuscation scenarios.

Speed is essential in your test tool, and typical exercises should initialize in seconds…not minutes, hours, or days. For example, your exercise may require 16 million IP addresses and 775,000 sessions. As a point of comparison, that kind of scale would equate to anywhere from about 77,000 to 750,000 end nodes depending on how much work those nodes were actually doing in any given time. It sounds extraordinary, but your test tool should typically take just 30 or 40 second to initialize, even for a very complicated range scenario like the one we described above.

Think about how long it would take to set up and tear down 77,000 to 750,000 virtual machines or real physical machines, days perhaps even weeks or longer.

Speed is essential in your test tool, and typical exercises should initialize in seconds…not minutes, hours, or days.

When the scenario finishes initialization, your tool should provide a real time stats view as a scenario plays out, displaying statistics for interfaces, TCP, SSL/TLS, IPsec, applications, clients, attacks, GTP, and other resources. You should be able to see a variety of failures occurring due to certain attack vectors such as network layer fuzzing. Other real-time stats should allow you to hone in on a variety of things that may interest you. For example, you should be able to check advanced encryption rates for SSL/TLS and IPsec, look at real time application performance, and check attack statistics details.

Your tool should provide obfuscation and evasion mechanisms to be employed in generated attack vectors.

Obfuscation and Evasion Mechanisms

Your tool should provide obfuscation and evasion mechanisms to be employed in generated attack vectors. For example, if you are modeling a JavaScript-based attack, you should be able to actually change the JavaScript each time the attack runs while maintaining the integrity of the attack. This allows a cyber range assessment to perhaps weed out poorly written firewall rules.

Let's say a JavaScript attack might use a variable name of Fu one time, but the next time that might change to Bar. If a poorly written rule relies on a JavaScript element that can legally change, like a variable name, the attack would still get through. On the other hand, a well-written rule would stop the attack regardless of the obfuscation used. Similarly, you should also be able to configure the attacks to use some very devious evasions, such as to force IP fragmentation, use Base64 encoding, or embed white space to defeat protection equipment; or to ensure protection equipment can correct and detect those evaded attacks.

This is absolutely critical because a highly skilled attacker will not just give up on an attack that used to worked once they realized the defending team has found a way to block it. Instead, they will see if they can still slip by those defenses by making minor tweaks in the attack using the same mechanisms that are available to you in your range now using the application and security test solution. With this solution, you will finally have the ability to model what real-world attackers actually do, ensuring that you can properly defend against them. You can model these attacks to scale all the way to line rate as needed.

Continually-Evolving Application Protocols and Security Attacks

Your application and security test tool must supply you with the application protocols and security attacks found in today's highly-dynamic networks. This is not a static product that is delivered at a point in time. You'll need a simple way to download updates to your toolbox of application protocols and security attacks on an on-going basis from your test tool GUI. The attacks of old are still around and have been added-to daily with new attacks and an ever-growing number of applications. We all know that realism is critical in a cyber range. That realism comes from having access to the applications used on your network and the very latest attacks found across the world.

Look for tools from companies that take this function very seriously and employ a team of experts to continually add security attacks, live malware, obfuscations, evasions, and application protocols. You should be able to model attacks ranging from viral attacks like 1999's infamous Ping of Death attack through the latest complex attack vectors that include models of the notorious Stuxnet and Flame malware; and DDoS attacks ranging from simple SYN floods to the complex DNS reflection attacks that devastated Spamhaus and caused localized Internet problems in parts of Western Europe.

Comprehensive Reporting

Your test tool should provide comprehensive reporting so you can investigate anything you'd like across a wide range of measuring criteria including latencies, frame rates, applications transaction performance, RTP jitter for voice and video, throughput, TCP performance, and many more.

Comprehensive reporting allows you to verify with certainty that your network performance with respect to background traffic is acceptable while your network undergoes a variety of attack conditions. Reports and graphs can help you understand when and how things start to degrade. The white team can use this information to cross-check red team and blue team claims on the overall effects of attacks on critical network elements protocols and feeds.

Conclusion

In summary, the ongoing flood of innovative applications and evolving security attacks quickly outpaced the costly, static, labor- and equipment-intensive traditional cyber ranges of the past. At the same time, the dramatic increase in cyber threats to enterprises must be addressed, driving the need for compact cyber ranges that can fit within an enterprise IT budget.

A well-designed cyber range is the first step in enabling your organization to perform complex technology assessments, as well as level today's playing field against complicated cyber attack scenarios by training cyber warriors. A valid cyber range uses a hybrid physical/virtual construct and includes application and security test tools to assist with realistic mixes of highly scalable mission-critical background traffic and attack traffic while providing detailed analysis in real time and via comprehensive reporting tools.

With a cyber range, you'll improve the capabilities of your network defense, thereby ensuring the resiliency of your networks and applications, and that your blue teams and red teams are prepared for real-world cyber scenarios by allowing them to train as they fight.

How Ixia Can Help

Ixia develops market-leading application and security test systems used as the foundation for enterprise, government, and educational institution cyber ranges. If you need a quick start, our Professional Services personnel have the expertise to work with you to create your cyber range and train your cyber warriors. We also have specific Cyber Range Training educational courses available to show you how to use our BreakingPoint test system to create an optimal cyber range and operational scenarios.

A well-designed cyber range is the first step in enabling your organization to perform complex technology assessments, as well as level today's playing field against complicated cyber attack scenarios by training cyber warriors.

Ixia Worldwide Headquarters
26601 Agoura Rd.
Calabasas, CA 91302

(Toll Free North America)
1.877.367.4942

(Outside North America)
+1.818.871.1800
(Fax) 818.871.1805
www.ixiacom.com

Ixia European Headquarters
Ixia Technologies Europe Ltd
Clarion House, Norreys Drive
Maidenhead SL6 4FL
United Kingdom

Sales +44 1628 408750
(Fax) +44 1628 639916

Ixia Asia Pacific Headquarters
21 Serangoon North Avenue 5
#04-01
Singapore 554864

Sales +65.6332.0125
Fax +65.6332.0127

915-6729-01 Rev. A, August 2014