



Foreword

AN INTERVIEW WITH THE AUTHOR

Usually a book's *Foreword* consists of someone telling you a bunch of stuff about the book you're holding in your hand—either to encourage you to buy it, or to get you excited about the book before you read it. I don't know about you, dear reader, but if I'd picked up a book on network security by Richard Bejtlich, I really couldn't care less what Marcus J. Ranum also thinks about the book. I'm sure you're asking yourself, "Is this worth reading?" and you'd be pretty silly to take my word for it, in either case.

So what I thought I'd do, instead of the usual boring *Foreword*, is interview the author of the book. I'm an author myself, and it's been my experience that there's usually a lot of "why I wrote the book" kind of information, which you can't really put into the book itself, that would probably be pretty interesting. Without further ado, then, Richard Bejtlich, as interviewed by Marcus Ranum:

MJR: Richard, first off, thanks for taking the time out of your writing and teaching schedule to do this interview. I know you're a super-busy guy. So—last year you published your book on *network security monitoring*, and now it's *extrusion detection*. After reading both, I can see you're building a consistent worldview of how computer/network security should be done, and so far the underlying message I'm coming away with is "know what's going on, first and foremost." That really resonates with the old school security practitioners who basically felt



FOREWORD


that audit and change detection were one of the fundamental building blocks for secure systems. So you're leading us through this trajectory—any comments on what's next? What's the next lesson?

RB: My first book tried to alter the mindset of traditional intrusion detection system (IDS) users. I've found that too many security analysts rely on their IDS to identify compromised systems. Others believe that their so-called "intrusion prevention system" (IPS) has rendered the IDS obsolete. Unfortunately, it's not difficult to evade an IDS or IPS, despite the good work done by a variety of vendors and developers. A variety of technical problems, including lack of context and situational awareness, encryption, and various forms of fragmentation and application-layer obfuscation make it difficult for any network detection or prevention product to be completely effective, especially against expert attackers.

Beyond the technical limitations of security products, analytical and procedural obstacles frequently allow sophisticated intruders to evade detection and prevention mechanisms. Most vendors and analysts see an IDS alert, or an IPS block action, as the end goal of any security incident. They consider their job done if they take some sort of action based on the traffic they inspect. Unfortunately, when an IDS alert or IPS block action is reported, analysts on the front lines frequently ask themselves, "Now what?"

If analysts instead see alerts as the beginning of a security investigation, and not the end, then the IDS or IPS becomes a more useful tool. Analysts would then begin to wonder about other activities the intruder may have attempted that were not seen by the IDS or IPS. When one has the necessary data to move beyond alerts, then it is possible to detect and control sophisticated intruders. Accordingly, my first book provided theories, techniques, and tools to move "beyond intrusion detection" and its alert-centric data to incorporate full content, session, and statistical data.

Moving beyond intrusion detection does not mean adopting intrusion prevention. An IPS is certainly a helpful device that allows for more granular blocking actions. The two technologies serve fundamentally different functions in network security, even though both must be able to identify attacks or intrusions to accomplish their roles. An IPS is an access control device with a prevention function. The IPS should enforce a network security policy. An IDS is (or should be, if properly selected and deployed) a policy-failure detection device. The IDS should sound the alarm when router access control lists, firewall rules, IPS mechanisms, and host-based defenses fail to prevent an intrusion.



Those who accept the “inevitability” or “logic” of “converging” the IPS with the IDS into a single platform fail to appreciate the importance of separating the prevention and detection functions. The traditional audit community understands the need for separation of preventative and detective controls. A bank would go bankrupt if it employed a single person to authorize payments *and* detect fraud. Why should we expect network security to be any different?

This new book tries again to change the way security architects and analysts build and watch the network. Shortly after the first book was published, I discovered and responded to a bot net in a client’s enterprise. A bot net is a collection of systems under the control of a remote intruder. This client presented a minimal Internet footprint; essentially, its only public IP belonged to a gateway/firewall/router (GFR). Despite not offering any services to the Internet, this client suffered multiple internal intrusions. I realized that watching inbound traffic to the public IP address was not very useful for this client. Traffic initiated by remote hosts, destined for the GFR, would be dropped. Instead, it was much more interesting to watch traffic leaving this client. Hence, the idea of “extrusion detection.”

While not a novel term or concept, no one else had devoted much print to the subject. This book is designed to fill that gap. Thus far I’ve concentrated on inbound traffic in the first book, and outbound traffic in this one. Traffic that never leaves the intranet is a more difficult problem. A threat model that consists solely of internal traffic, with no communication with the Internet, means activity by rogue insiders. Internal traffic load also dwarfs the bandwidth used in the perimeter. Additionally, vendors like Microsoft are pushing for ubiquitous deployment of Internet Protocol Security (IPSec) internally.¹ I think what that means is that the next place to watch is each host—not the traffic passed between hosts.

MJR: You talk about trying to change the way security analysts build their networks—this is a possible problem, isn’t it? I know I’ve seen a lot of networks in the last five years, and they’re built all wrong, from a standpoint of security and survivability. A lot of the ideas you’re trying to put in front of network administrators are definitely the kind of thing that would be vastly more effective if they were built into the network from the get-go. If you were talking to a network administrator who’d just gotten tagged with security, where would you tell her to spend her first \$10,000 and her first weeks of effort?

1. See <http://www.microsoft.com/windowsserver2003/technologies/networking/ipsec/default.mspx> for more information on IPSec in Microsoft networks.

FOREWORD

RB: I would begin by assessing the degree to which the administrator's enterprise is a defensible network. A defensible network, as explained in Chapter 2, is an information resource that is monitored, controlled, minimized, and current. Those operating a defensible network have the best chances of resisting intrusions. If and when any compromise does occur, a defensible network is best postured for rapid intrusion identification and efficient incident response.

The four defensible network components are ordered by ease of implementation. Begin with monitoring. At the very heart of any defensible network is the idea of figuring out what is happening in the enterprise. If you have no idea how your network is being used, by authorized and unauthorized parties, it is difficult to know how to move forward. Unfortunately, lack of knowledge of network use and abuse does not stop many organizations from implementing the security silver bullet *du jour*.

Assume the administrator has no spare equipment to begin monitoring. With \$10,000, the administrator could buy one or more decent server-class systems to host an open source NSM suite like Sguil. She may need to buy one or more taps or perhaps an enterprise-class switch. I would also recommend buying one or more books from my recommended reading lists (<http://www.bejtlich.net/reading.html>) to guide her analysis process. There's no point deploying equipment and inspecting traffic if it cannot be deciphered!

I suggest conducting a traffic threat assessment, as described in Chapter 6, to get an idea of exactly what sort of activity is entering and leaving the enterprise. Based on her monitoring findings—and there will be findings of some unpleasant sort—she may find it easier to justify additional expenditures. From there, continue with control. Open source solutions like the Pf firewall on BSD and the Squid proxy can begin to limit inbound and outbound traffic. Minimizing and updating software will be costly in terms of time, but hopefully not in financial expenditures. Of course, a large enterprise may require a commercial patch management solution.

Incidentally, I originally wrote Chapter 2 to help reduce the amount of traffic an analyst must inspect. Just as it is impossible to prevent intrusions on an indefensible network, it is nearly impossible to detect them. When any traffic is allowed to pass to any host in any direction, how can an analyst decide what is normal, suspicious, or malicious? Implementing a defensible network architecture provides preventative benefits and assists detection operations. Entire books could be written on good network infrastructure. The purpose of Chapter 2 is to narrow the amount of traffic analysts must investigate, particularly in the outbound direction. The main focus of the book is *extrusion detection*, but



extrusion prevention is well-served by implementing a defensible network architecture.

MJR: I've noticed you're a fan of Bruce Lee! It's interesting to me how a lot of us security guys find parallels between computer/network security and the martial arts/art of war. Remember Lee's great "It's like a finger pointing away to the moon" speech? What do you think would be the equivalent for a student of computer security? What do you think Bruce would tell us?

RB: I am indeed a fan of Bruce Lee, and I've practiced several martial arts. I even asked Jackie Chan, in person in 1998 at a book signing, to sing at my wedding!

I remember hearing Bruce talk about not anticipating an opponent's actions. I think he would see parallels in the way many security practitioners rely on IDS signatures or watch for known patterns of malicious activity. That sort of behavior is similar to facing an opponent known for his powerful punching techniques. You might wait for him to position his hands as a sign that a punch was coming. You could focus all of your attention waiting for that one indicator and totally miss the barrage of kicks he throws your way.

I advise that intruders should be viewed as smart (sometimes smarter than you) and unpredictable, and able to beat your defenses. Bruce would probably agree. He would train to be ready for whatever his opponent would deliver, and he would have techniques in place to deal with the consequences of not blocking an initial punch or kick. Rather than failing catastrophically when an opponent lands a blow, Bruce would take advantage of the attacker's proximity to initiate a different sort of counterattack or improved defense.

Bruce also based his fighting style upon what he found to work in the real world. I once heard a story about Bruce and a contemporary martial artist, American Kenpo founder Ed Parker. The two martial arts pioneers are reported to have enjoyed dressing and acting as drunks outside bars in rougher parts of the city. They would wait outside the door late at night with money hanging from their pockets. When local toughs stepped out of the bar and decided to "take advantage" of the supposedly drunken duo, Bruce and Ed would try out their latest punching and kicking combinations!

This reliance on real-world experience helped Bruce and Ed develop techniques that were efficient, compact, and effective. While theory and beliefs were important, they were not the sole basis for the pair's fighting systems. A book called *The Visible Ops Handbook* by Kevin Behr, Gene Kim, and George Spafford (Eugene, OR: Information Technology Process Institute, 2005) would approvingly call their approach "management by fact." In comparison, too many security personnel



FOREWORD

seem to “manage by belief.” *Visible Ops* coined that phrase for those who act without real-world knowledge. All of my books try to emphasize that gathering information on threats is crucial. Traffic threat assessments and network forensics (covered in this book) are ways to determine how an enterprise network is really being used. My company’s motto, “Know your network before an intruder does,” exemplifies the importance of management by fact.

MJR: Different counterattacks or improved defenses. . . . So, really, you’re advocating a war of maneuver. Static defenses don’t work against an opponent that is inventing new attacks; we need to invent new defenses. And knowledge is the most important weapon in our arsenal for doing that. So you seem to be pretty firmly in the school of “get your hands dirty and learn stuff” rather than “run out and buy something that does it for you.” I’m guessing you’re not a big fan of outsourcing security?

RB: I do believe in investing in training one’s people to meet organizational goals. For example, I personally do not have a problem with hiring someone who can configure and deploy open source solutions. In contrast, some organizations prefer hiring people that administer commercial solutions, because management believes knowledge of commercial products is more widespread and visible.

I don’t think security can ever be “outsourced,” since the victim bears the ultimate responsibility and consequences of any incident. However, competent managed security service providers (MSSPs) offer three main advantages to their customers. First, some MSSP personnel are deep security experts. Their teams cover multiple disciplines. It is difficult for a multi-tasked enterprise administrator to find the time to stay as current with security issues as a dedicated MSSP analyst.

Second, properly staffed MSSPs ensure experts are available on an around-the-clock basis to monitor and respond to security incidents. This response time closes the window of vulnerability and may reduce the damage caused by an intrusion. Third, MSSPs responsible for a decent number of customers have a wide field of view of the Internet. The MSSP can see activity affect one client and use that knowledge to warn all other clients.

The problem with most MSSPs is that they subscribe to a failed model of intrusion detection. Most do not collect NSM data (alert, full content, session, and statistical data) that would allow the MSSP to detect and contain high-end intrusions. Some MSSPs seem to be nothing more than “worm catchers.” Other MSSPs consider it advantageous to never inspect traffic and to rely on system and



event log messages. Besides the value of log aggregation, I think log-centric MSSPs deliver limited value to their clients.

MJR: So where do you see the “next big thing” on the offensive side coming from? What piece of badness are you most concerned about?

RB: This is an excellent question. I’ve largely given up trying to figure out what comes next. It is probably fashionable to talk about attacks against non-PC yet IP-enabled devices like smart phones, personal digital assistants, cars running Windows Automotive 5.0, and the like. All of this will happen, if only because “owning” someone’s car will be one of the most interesting exploits of the decade.

Rather than try to appear smart by making predictions, I fall back on my NSM principle that says intruders are smart and unpredictable, so prevention eventually fails. The security industry could spend a lot of time and money on what it thinks is the “next big attack.” Suddenly, a smart person in a remote part of the world unleashes an exploit or technique that rocks the foundations of the Internet.

MJR: You keep coming back to that notion—since the attacks are going to be unpredictable and change, preparedness and flexibility are the keys to defense. I couldn’t agree more. So the general recommendation for dealing with the next big attack is likely to be “know as much as possible about what’s going on in your network”—there’s no silver bullet, though, is there?

RB: That’s exactly right. If we don’t—and in many cases, can’t—predict what’s going to happen, we should put in place people, processes, and products that are equipped to handle unknown problems. In the monitoring world, we must ensure that at least some of our data collection techniques are content-neutral. In the past I’ve used the term “network audit,” but that is becoming a loaded phrase now that traditional auditors are taking the reins away from security staff. I now say we should perform transaction logging wherever possible. At the wire level, collecting session data is a great way to log network transactions. At the host level, event logs perform similar functions.

In some ways, it’s like dealing with a new disease. You can’t possibly immunize everyone against every disease ever to affect any person. Instead, you watch for indicators or symptoms of a serious disease in a few people. They obviously and tragically suffer, but they provide the knowledge and hopefully the early warning that spurs the medical incident response process into action. It’s a “Centers for Disease Control” model rather than a “high castle wall” model. Of course prevention still has a role, but the prevention can only be really effective against known threats. There’s no sense fortifying your castle wall because you think that’s the



FOREWORD

enemy attack vector when he's planning to tunnel under that wall. In some rare cases, it may be possible to eliminate an entire class of attack via preventative measures. If that is truly the case, it may be worthwhile to devote resources to removing that threat. In most cases, however, I prefer to balance prevention, detection, and response.

MJR: Richard, thank you!

RB: You're entirely welcome.

There you have it, dear reader—the “view from behind the book,” as it were. Personally, I really like the way Richard thinks about security. He's conservative about fundamentals, but he's not afraid to challenge your preconceptions, either. I've enjoyed reading this book, and I've learned from it in the process. I hope you will, too.

Marcus J. Ranum
Chief Security Officer
Tenable Network Security, Inc.
Morrisdale, Pennsylvania

