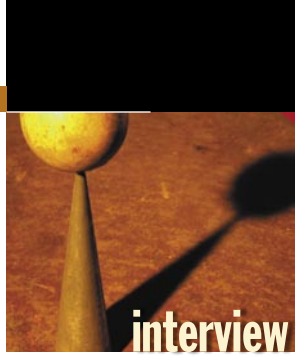

A Conversation with Jamie Butler

Photography by Steve Ruark

Rootkitting

OUT ALL EVIL

Rootkit technology hit center stage in 2005 when analysts discovered that Sony BMG surreptitiously installed a rootkit as part of its DRM (digital rights management) solution. Although that debacle increased general awareness of rootkits, the technology remains the scourge of the software industry through its ability to hide processes and files from detection by system analysis and anti-malware tools.

The best way to understand rootkits—how they work and how best to detect them—is to write one yourself. This month's interview subject, Jamie Butler, has done just that. Butler wrote the well-known FU rootkit, a proof-of-concept that illustrates vulnerabilities in the Windows and Linux operating system kernels. Butler also wrote a book on rootkits, a tome he coauthored with Greg Hoglund entitled *Rootkits: Subverting the Windows Kernel* (Addison-Wesley, 2005). Prior to that, the team collaborated on the rootkit.com Web site, a repository of rootkit information, code, and discussion. The Web site is controversial, with some security professionals bemoaning the fact that it provides executable rootkit code that could be exploited by miscreants.

Currently at Mandiant, Butler is the principal software engineer on the product development team. Prior to Mandiant, he was director of engineering at HB Gary and CTO of Komuko Inc., where he developed a low-level rootkit detection product.

Interviewing Butler is Matt Williamson, principal scientist at Sana Security. No stranger to rootkits himself, Williamson has spent his career inventing and integrating anti-malware technologies. At Sana, Williamson developed behavior-based malware detection and removal technology that identifies malware by looking at what the code *does* rather than what the code *is*. Prior to joining Sana, he worked at Hewlett-Packard Labs on a virus containment technology called Virus Throttling. Williamson has a Ph.D. in computer science from MIT.

MATT WILLIAMSON I think rootkit.com is a good place to start. A rootkit is software used to hide other software from the user and security tools, to evade detection. Rootkit technology is a common component of malicious software. Rootkit.com is a Web site where various aspects of rootkits are discussed. Do you know the early history of that site? Were you involved in setting it up?

JAMIE BUTLER Rootkit.com came along a few years before I got started, but I'm a close friend of Greg Hoglund, who established the site. I believe that his goal in starting rootkit.com, and much the same reason I got into this area of research, was to debunk the false sense of security in the security software market. Hoglund wanted to prove that the company he was working for at the time needed a more thorough solution than it was using.

MW Where in particular did the false sense of security come from?

JB Well, there was one technology used to identify malware and such, and Hoglund believed it was possible to hide from the detection algorithms or software that his company was using. It was more proof-of-concept.

MW So from the beginning, rootkit.com was more of a disclosure type of organization.

JB It's also an open community to discuss better ways to detect these things. The site also has some threads about malicious software.

MW Was the idea, then, to show publicly that these security tools weren't working well by giving examples of where they didn't work?

JB Yes, that was why it was founded. When I got involved, my role was to show that the technology at the time wasn't good enough for the level of threat that was really out there. Just because you buy something for $29.95 and install it across your enterprise doesn't mean it necessarily does everything the glossy tells you it does. There are circumstances where you aren't protected and perhaps you never will be, but the goal of rootkit.com was to try to bring those out into the public discussion so that they could be researched more in depth and solutions could be adopted by the vendors. It's free to both security vendors and malicious people.

MW I guess that's always going to be the case when you make things public: they can be used for good or for evil. But do you think that rootkit.com was successful in raising the attention of the security vendors?

JB I don't think it single-handedly changed the security environment that we live in, but I do think that along with vulnerability disclosure lists and other types of open information sharing, it has prompted the security environment to change quite radically over the past two or so years. And it's still evolving.

MW What about the converse? What sort of impact do you think it has had on the malware writers, the people who are using these techniques a lot more commonly? We perhaps need to highlight that in the past two to three years, the purpose of writing malicious software and distributing it on the Internet has really changed radically from amateurs writing it for fun, if you like, to professionals writing it to make money. The way they make money is by stealing and selling information from machines. Sana Security has a behavior-based detection product, and of the malware that we detect in the wild, a significant proportion has some sort of rootkit technology in it hiding files, processes, DLLs, and so on. In my experience, there has been a terrific take-up in rootkit technology out in the wild in the past two years. Maybe it's hard to chart these things, but do you have any inkling about how much of that is a result of the influence of tools such as rootkit.com?

JB I don't have raw statistics, but I do know from vendors that rootkits that are kind of self-packaged or don't require a lot of recompilation and so forth were adopted quite widely and used in everything from botnets to worms. Most of those rootkits, however, were software on rootkit.com, one of which was the FU rootkit, which I wrote. They weren't something that extremely malicious people would use if they really wanted to hide their presence. There were ways to detect them that were brought to light maybe a year or two after the fact. Not only that, there was no remote command and control system. There were no encryption modules for any communications. There were no self-destruct mechanisms within the rootkit so it would go away. There was no polymorphism. There was no data deletion or even data acquisition within the rootkit, so I believe that most of the better technologies that are discussed on rootkit.com, such as Shadow Walker, the FU rootkit, the FU-2 rootkit, and even the original NT rootkit, were more academic in that they showed the level of threat and were "demo-able." They weren't everything that a hacker would want to use, however.

MW I thought at least a couple of years ago that the FU rootkit was very popular because it was easy for someone to tack on to the bot, for example.

JB It was very popular, according to what I've heard from various vendors. It didn't try to hide itself, however. Therefore, I don't think it was that malicious. Then again, it wasn't malicious to begin with. It was whatever you chose to do with it.

MW Wouldn't it often be used by programs that were stealing data from the user?

JB That was possible, sure, but that would have required some integration.

MW You know, it's like nuclear technology: it can be used for good and for evil. Even the proof-of-concepts might not in themselves be openly malicious in the sense that you can use this software to immediately make money by stealing information. With the rise of such activity on the Internet, however, it appears that people have picked up technologies

It's definitely a technology that can be used either way, good or evil.

such as FU and have cut-and-pasted them into their malware to give it a little bit of an edge. It might not be the ultimate edge that would allow them to break into some high-security site, but it gives them enough of an edge to be successful in stealing information from your average computer user.

JB It's definitely a technology that can be used either way, good or evil, and it has raised the bar in the security community. I don't think, however, that rootkit.com and the whole open source disclosure arena have really changed things all that much because, going back to the nuclear analogy, the threat exists. Now, what are you more afraid of, the country that develops the threat in secret or the country that develops it in the open for the world to see? You're more afraid of what you don't know or don't understand than what is publicly disclosed. From the anti-virus point of view, as soon as something hits the site, the anti-virus companies are all over it. It actually trains their development base on what is possible and what may potentially come down the road. I think Hoglund and I would both argue that the threat is there; whether you know about it or not is the only distinguishing factor of rootkit.com.

MW I think that's true, but isn't there the issue of scale to consider? Someone privately in his or her basement with a nuclear weapon is very scary because you don't know what that person is going to do with it. You would rather that someone had it publicly because the assumption is that you would be able to contain the situation through diplomacy or other mechanisms. But isn't the issue really that by making technology available that's easily redistributable because it exists either as code or even precompiled executables, you're allowing a scaling of usage that in itself becomes a problem?

It's not just a capacity to do something that's private and scary, yet small, that we're worried about—it's the large-scale users who are difficult to contain. It's as if someone published a "how to make a nuclear bomb" kit that you could get for free, and then everyone would have one. That would be a different sort of bad problem than one controllable state or one private organization having it.

JB I would agree that the scaling does increase because of the "cut-and-pasteability" of the code if it's public. The analogy falls apart in that there are a limited number of places to hide. This has been discussed at several conferences over the past year. The rootkit technologies that are currently public have become well understood in the past year to year and a half. There are solutions. They are easily detectable if you know where to look.

So I would say that your scaling problem is almost the inverse, because at first you had an entire operating system to look at, and perhaps companies didn't quite know all the nooks and crannies where they could hide. Now that it has been exposed in public software, rootkit.com, and other places, they can actually see that there are only a limited number of ways to hide a process. Once you hide a process, you will see that there are some things that cannot be hidden. There's no way to hide the fact that it has a thread, as long as you want that thread to run on the system. The number of places to detect has grown smaller instead of larger.

The problem—or the deployment of rootkits—may have grown larger, but the ways to detect them are actually growing smaller because there's a well-known set of places to look, algorithms to use, and so forth.

MW Maybe this is the difference, then, because what we see is that there's a great scale of things where the technology isn't particularly sophisticated, but it's sophisticated enough for the attacker to find that piece of malware useful.

I would agree with you that from an academic perspective the limited number of places to hide means that these things should be easier to detect, but the sophistication of the malware we're seeing in the wild is not as high as that. This must mean either that people are making money quite happily without having to resort to those specialized techniques or that the security products aren't actually getting up to the level of sophistication to handle those cases.

MW I was talking to Greg Hoglund the other day, and he said that this move to malware for profit had actually changed the tone of rootkit.com quite significantly. Would you agree?

JB Yes, I would totally agree with that statement. As I mentioned before, originally rootkit.com was to show that we had some room to grow as far as making better security products, and that we should be using better algorithms and better techniques to find these things. That went very well for quite a while, but recently rootkit.com has become somewhat stagnant in its content and its contributors. This is partially because the subject matter is more complex than most people care to dive into, but also because there is a thriving market for this type of technology. When you have a spyware company that will pay anywhere from $20,000 to maybe $80,000 for a quality rootkit that will allow them to stay on the box just a few more days or a few more weeks so that they can hide from the security software, that becomes a large motivating factor for people who want the money.

I also see the flip side of that. I don't think there are many people selling their rootkits to these nefarious companies. There is some of that going on, but I think people fear their ideas being used for someone else's profit. They may not know exactly who's going to take this little tidbit of knowledge, incorporate it into some adware program, and then spam the world at two cents an ad.

MW What you're saying is that it's not exactly that the contributors on rootkit.com are now off earning between $20,000 and $80,000 per fresh exploit that they can produce. It's more that people now feel that if they put their ideas on rootkit.com, someone else is going to profit from them. But wouldn't that have existed three or four years ago? Do you think the polarization is a result of people making so much money off of the malware technologies in the criminal world?

JB I think the polarization has come about because of the larger amounts of money available in the criminal world. I don't know what is different other than that the criminal element wasn't so public back then.

MW I think that the change from amateur to professional writing of malware has had a tremendous impact in many areas.

JB The security arena is getting better. Before, you could talk about hooking system calls, which was like the original NT rootkit, and that lasted for maybe four years. That was the greatest thing. I just read a paper that was in last year's *Security and Privacy* magazine about rootkits that were hooking systems and causing malfunctions and so forth. This is very old technology, but it has lasted for many years. It used to be that you could talk about something, and it would exist for years. Now you talk about something, and it's gone in two weeks.

MW It's interesting that the security world has gotten more responsive to these things.

JB That was the original motivation for rootkit.com, so I guess it has been a self-fulfilling prophecy.

MW So far, we've painted a picture of the world where

people are writing malware for profit, and the exploits and technologies are entering into this underground economy. Given this world, what do you think the consequences are for security technologies and the enterprises that deploy them?

JB One of the detrimental effects of commercialization and the resultant lack of discussion is that enterprises are more vulnerable than perhaps they were when there was more public discussion. Because there are now zero-day exploits [an attack for which there is no warning and no protection], there's motivation to find a type of zero-day hiding technique or original compromise. There's not going to be any disclosure of it, and we all know that you're only as safe as your weakest link. Therefore, often you don't even find that you've been had until it's too late or until you accidentally stumble upon something.

MW What you're saying is that in the old days, you at least had a chance of everyone being in the open about what the latest attacks were, but now that's really long gone.

JB I think so. I would much rather face a world where I know that there's a certain element of malicious people out there who are willing basically to spam the world to get my data or whatever they want to take from me. I also know that they have a limited money pool to pull from, and there's no organized financially based incentive for them. But now that we've gone more underground, that's exactly what we're seeing. Maybe the malicious people don't have their own malware labs, but they do have money to spend, and now they can buy the very latest technology. And security is the hard thing because with security, everyone knows you just have to find one hole. To protect against these types of threats, you actually have to block all possible vectors.

MW So, when the public sharing goes away, the security companies appear to lose because they can't use that as a forum to get information anymore. But they must have their own private forums—networks between companies—that they can leverage.

JB That's true.

MW What's missing, then? Are you saying that in the old days the people who were gray or black would be contributing to these open cases, and now that information is not getting into the hands of the security companies?

JB Right. It's one thing to solve the problem. It's another thing to create the problem and solve it. For example, with companies such as Sana Security or Mandiant, if the public disclosure dries up, you have your labs and so forth, but you're left with them not just looking up the information, trying to figure out a little bit how it works, and then developing some protection or detection against it. Instead, now you're going to need additional resources not only to find and stop the threat, but also to create the threat.

MW And then even with that, you're only guessing that you've created the same one that someone else has created.

JB Right. You're poking holes in the boat and patching them at the same time.

MW And that's actually really hard to do from a prioritization point of view. You must know, too, from being in a software company, that when you want to put in a new feature to cover a hole, often the decision to put it in is based on what its impact will be. But what about the future? Where do you think the malware trends are headed?

JB With the money that's behind these things, attackers can now have a very focused effort. Often, they're not going to want to spam the world because they now have a targeted exploit. They bought it on the underground market, and they know that they have exclusive access to it and that the data they want is valuable enough so that they need to target only certain companies or individuals. They're not going to spend the time to go after everyone within the network.

So now you have a very targeted, very specific attack. And the more specific these attacks become, the more customized they are. When they're customized, by definition, they're harder to detect. It's definitely a much scarier scenario than when I knew there were a certain number of open source repositories, a certain number of open source vulnerability disclosure places, and there were maybe X number of bad guys willing to blast the world.

MW I tend to think of these things from the business point of view of the attacker. Capitalism is pretty ruthless. It seems to me that if you can make a lot of money with targeted attacks, if you can sell a lot of information very quickly, then those types of attacks will grow.

At the moment, however, what we tend to see are much more of the less sophisticated but more broadly applicable attacks, where the information that is being stolen is more generic—for example, logins, passwords, credit card numbers, and so on. They're not the kinds of targets that you need sophisticated technology to penetrate.

I think what's going on is just being driven by the economics of software development and the money you're getting from the information.

JB I would agree with that. I think it's the ebb and flow of the free market into the security world. It used to be that

if you were a security company and you were going for a low-cost $19.95-type of subscription model, all you would have to do is solve 99 percent of the problems. That other 1 percent, which is extremely leading edge and really cool and sexy to demo, wasn't your problem because you were not going to see it, your customers probably weren't going to see it, and to explain the market need to develop that would be hard. The return on investment was low.

Then open source sharing came more into play and pushed that 99 percent of the solution down to maybe 95 or 90 percent. Now there's about 10 percent of market share to be gained if you can be the coolest on the block as far as detecting the problem reliably every time.

So, security companies are getting better, and their products are better. The number of hits is going up because they're detecting all of these shotgun blasts from malware companies. When that source of revenue finally dies, or at least dwindles to a level unacceptable for the malware company or malware entities, then you'll see the more targeted specific attacks.

MW I have to say that our view is a rather grimmer picture than your 99 percent numbers. The numbers we are seeing are much worse than that. They are around 50 to 60 percent, maybe even lower. In fact, for zero-day stuff it's much lower—and that's even comparing one anti-virus vendor against another.

There is an enormous amount of malware out there that is exploiting readily available, easy information to steal. It's not using particularly sophisticated techniques, but it's winning against the signature-based approaches through repacking, mutation, recompilation, and using rootkit-type techniques. It's the number of variants of each individual type that is actually causing the problem.

The argument that the trend will move toward targeted attacks makes the assumption that people will buy the technology to fill the gaps, to catch the pieces of malware that are being missed. That gap—the pieces of malware that are currently being missed—is pretty big, and it might take a while for us to deal with that.

So, the move toward targeted attacks will be quite slow. I think it will happen in individual cases, but in the larger picture, it will be quite slow.

JB Do you think that maybe 60 percent is currently being detected, or getting past?

MW I think it could be either way. Q

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums