Snort
and the future of IDS

NICK VEITCH talks to *Snort*'s creator about Sourcefire and network security.

Anyone who knows that IDS stands for Intrusion Detection System is probably aware of *Snort*. This GPLed software has not only become well respected in the area of IDS itself, but has also spawned many other projects. In addition, it has formed the key technology at the heart of Sourcefire, a company set up by *Snort*'s creator, Marty Roesch. Sourcefire builds on Snort's IDS capabilities by creating standalone network sensor devices, backed up by training and support.

We managed to track down the hard-working author of what is believed to be the most used IDS in the world and asked him about his particular take on *Snort*, Open Source and the future of network security.

LXF: **You have a lot of history in the computer security arena, and I guess that's why you embarked on the Snort project in the first place. Did you think there was something specific missing from the tools that did exist?**

MARTY ROESCH: I started writing *Snort* because I got irritated with *tcpdump*. I wrote it for a few reasons, but one of the first things I wanted to do with it was use it as a debugger for another program I was writing. It was a piece of network software and I needed to see the payload of the packets, so I wrote *Snort* to do that.

It's really a classic 'scratch your own itch' Open Source story – I had a need for a better sniffer so I wrote it.

LXF: **At what point did you realise there was a lot of interest in *Snort* in its own right?**

MR: The first couple of months *Snort* was out, it was just a sniffer, and I started adding features to it to let it start behaving as a network intrusion detection system. Once I did that, people started to get interested in it. At the time there were no network IDS in the Open Source world, and intrusion detection was still somewhat an arcane, black art. Really Snort started hitting its stride about a year later, after I cemented in the architecture we are using today which actually makes Snort not just into an IDS but a flexible software framework for performing network traffic analysis that we just happen to use primarily as a NIDS.

You can make Snort do all sorts of tricks because it's extensible and has a plugin system.

LXF: **I guess that's one of the reasons it's been so popular – because you can use it for all sorts of things.**

MR: That's it. The research community liked it because *Snort* provided an API and interface to a decoded packet stream that they could do all sorts of rapid prototyping of ideas – so we saw *Snort* getting picked up by universities and government environments really quickly.

It really gave people a solid platform for network traffic analysis with. Using it as an IDS is the primary reason it got popular, but once people understood all the interesting things you could do with it, it got popular for more than that – it became popular for being a neat piece of software that you can kinda use as a Swiss army knife.

LXF: **It must have been gratifying for you that it was so popular. I don't suppose you imagined so many people would end up using it when you started?**

MR: No, I was really surprised actually. It was a big shock to me. I put *Snort* out there and people started to use it, and that was as expected. But I didn't expect as many people to use it as are using it today, and also I didn't expect the organisations that are now using it to pick it up – we have large financial institutions, government and military institutions – really people who have some very specific and critical needs using the software. It was surprising to me that it was that useful. Or that I had done that good a job! The good thing is that it's so extensible that maybe they could trim out all the bad stuff I did and fix it to further their own goals!

LXF: **I guess having people like that using the software, with very specific needs and aims, must**

The Snort.org website is hosted by Sourcefire.

have generated a lot of feedback on what you could actually do and how things might develop from there.

MR: It's interesting. We have a lot of different viewpoints on Snort in the world at large – we have the Open Source world view which is that Snort is neat and can do a lot of interesting things, and then we have the commercial competitors view of Snort. They try to box it in and minimise it. One of the things I enjoy is that some of the people who market products against *Snort*, they like to pick different versions when they do the comparison – they pick one version that didn't do something, another version that was missing something else, or a particular release that had a bug in it.

It is interesting how the whole political scene has evolved around it. It's definitely been a real education. Interestingly enough Open Source projects, once they get more than two people working on them – all of a sudden it's an exercise in politics and community management as much as it is in writing good software.

LXF: **So the reason competitors do like to pick holes in it is because security is big business? Obviously with Sourcefire you are involved in that side as well. I would guess that you never dreamt about building a company off the back of *Snort* when you started it, but at what point did that become a feasible option?**

MR: Yeah. If I'd thought where we wanted to be today and started writing *Snort*, I would be a lot smarter than I really am!

Basically in the fall of 2000 I came out of a startup that didn't do too well, where I was an engineer. I was looking around for what to do next, and because *Snort* was out there and I'd built a reputation around it. I didn't have a problem getting job offers, but I wanted to pick carefully because I hadn't had much fun at the startup.

I kicked around the idea of doing a company for a while, then another security company came along and made me a very nice offer to move to their company and bring *Snort* with me. That was the deciding point – I thought if it was worth that much to them then it must be really worth a lot more so I started Sourcefire. Capitalism at its finest!

LXF: **These days, as shown by the current state of the market, it seems to be a very popular idea to build a company around an Open Source project, and there is more of a general roadmap of how that might work. At the time though, did you have difficulty in trying to base your business strategy around a piece of Open Source software?**

MR: Well, I had a clear idea of what I wanted to do, but convincing other people that I was rational or sane was a difficult process. The investment community, when I was looking for funding in 2001, the investors were very down on Open Source. They didn't believe you could make money on it, they pointed to the collapse of VA and other hyped Open Source IPOs and said "Open Source is dead, don't you read the papers?" They basically said, " call us back when you make some money". So we did.

The idea with Sourcefire was that we would build proprietary systems around Open Source cores and we'd bring value to the Open Source system that would. It's kind of a basic idea. One of the problems corporations have with Open Source is that they are a little skittish about the notion of being supported by a community, so they like to have commercial support behind the products and we provide that as well as all the other things we provide here – all the technology we develop here to enhance snort to do Intrusion detection and what we call Intrusion Management better. So really it's a logical evolution. A lot of people have flirted with different business models and I think this is the best of both worlds. You have the good parts of Open Source and the good parts of commercial development and as long as you maintain the separation of church and state and you maintain dedication to your Open Source components, I think it's a very good way to do business.

"I started developing *Snort* as an Open Source system just to see how the development methodology would work, and found it's a really great way to develop software."

MARTY ROESCH

LXF: **Obviously, when you started Sourcefire, Snort was already quite popular in Open Source. Was that 'kudos' easy to transfer to Sourcefire itself - in the commercial space perhaps many weren't very aware of Snort. Was it easy to use the popularity of Snort to gain ground early on with Sourcefire?**

MR: Yes, actually it was. Our original marketing campaign was my .sig file at the end of my emails. I was very active on mailing lists and that's how we got our first sale to Price Waterhouse, which was big win for us. The community is large and diverse enough, and I personally have a high enough profile out there that it wasn't too bad. Even if the business community didn't know about us, the technical community and the people doing security do know about us, so we were able to get their attention immediately, when we started operating.

That's how I ended up getting funded. These guys told me 'come back when you make some money', so I sold about $300,000 worth of product from my living room, primarily to Fortune100s. Then all of a sudden we were very popular!

LXF: **It seems that the answer to my next question should be obvious then; but the continuing development of *Snort* is important to you and to Sourcefire?**

MR: Oh yes. We continue to update and improve it. We had a new release a few weeks ago and there is another one coming up shortly. As we come up with new things for *Snort*, both here and community contributions, we add them into *Snort* and push it back out to the community. We spent a lot of money improving *Snort* and buying testing infrastructure and things like that to make sure it's a good system and hits performance marks. All the improvements and changes that we have made go back out to the Open Source community.

"If you work with the Open Source community, you can be successful in the face of large, entrenched competition."

LXF: **So there is a definite value to Sourcefire in *Snort* being Open Source?**

MR: Sure, There's lots in it for us – we have a very large and diverse group of users who provide excellent QA for us. We have a huge installed user base, which also means a huge potential customer base. We have the ability to speak directly to our market via our Open Source presence. I have a big soapbox I can stand up on any time I want to and a lot of people will see it when I do. For us, it has some distinct advantages.

Also, with development, the quality of the feedback we get and the rapidity with which we get it when we turn out releases is much much faster than with proprietary commercial systems, which means we can evolve the centre software much faster than you would typically see in an organisation of our size. We were able to compete toe-to-toe with the biggest companies in this space when we were only 20 people. To use a military term, it's a pretty incredible force multiplier. If you have solid technology and you build smartly around it, and let the Open Source community continue to do the great things that communities can do, and you work with that community, then you can really be very successful in the face of large, entrenched competition.

LXF: **With Sourcefire attracting ever-greater numbers of customers and *Snort* now getting a wider user-base and , is it easy to divide your time between the pure Open Source progression of Snort and the obligations and demands of your business?**

MR: No, not really. It's actually very hard. In large part I've had to take a much more management-oriented position within Sourcefire – we've hired many of the core *Snort* developers here, so I kind of provide direction for them, and a leadership role, but day-to-day I'm not writing tons of code in *Snort*; but I am writing code on a product we are building in-house here. I'm busy with that, and also supporting our sales team and other efforts that involve the success of both *Snort* and Sourcefire.

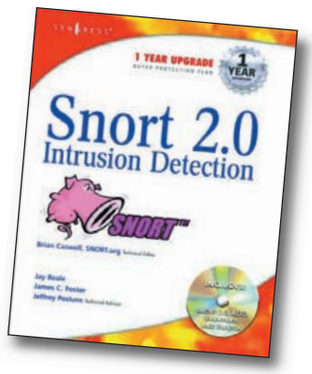

There are several good books available on using *Snort*.

LXF: **Do you have regrets about not being so hands-on with *Snort* as you used to be, or had you reached your limit with it anyway?**

MR: I miss being as involved with it as I used to be, but I have enough things to keep me busy right now. I'm working on some pretty cool technology right now. It's another idea that I had that we're turning into a product here. I miss it, but I don't miss the politics and a lot of the other stuff I had to deal with on a day-to-day basis. Yes, and No.

LXF: **If we talk about Intrusion Detection in general for a moment – a few years ago there probably weren't too many people investing much time or thought into it, apart from the obvious candidates who had to take it seriously. Nowadays it seems to be much more on the agenda for any business with a network presence. Do you think that enough thought is being given to the whole principle of IDS?**

MR: I think a lot of thought is being given to it, the problem that I suspect is really plaguing people is that, you can think about IDS a lot, but unless you can implement things that are reasonable from a usability standpoint – from the notion of just improving the state of security on a network. A lot of people seem to fall down on that – they don't seem to understand it thoroughly. I think a lot of people who design systems aren't really practitioners, and maybe they never have been. Certainly their marketing departments haven't. One of the things that is really neat about Snort is that a lot of guys involved with the development have been practitioners in the past or are currently. Myself, I used to work for the government doing network analysis – we have other guys here who have worked for government or universities, where they have handled very large IDS deployments, and I think it's very important that you have some of this background to be able to put together solutions that are really useful to people.

LXF: **I suspect that some of the problem is that these days IDS is one of those things you ought to have, and people go out and buy any old solution just to say that they have it, irrespective of whether it actually fits with what they really need to do. In order to get people to understand the problem, I guess there's a lot more education needs to be done for the customers first**

MR: Certainly. One of the big problems is that, when the user goes to analyse the data, typically the data they are getting is so technical that the number of people who know how to take what they are looking at and turn it into any reasonable response is very small. It was often said of *Snort* that it's a great system, but when you go to look at the data you have to be Marty to understand it. To an extent that's true. When I wrote *Snort* I was doing forensic analysis for large government, military organisations. I really wrote *Snort* for me, so the expectation was when I wrote it that I would be the one doing analysis. I can look at raw packet dumps and tell you what's going on, but the number of people out there that can is maybe measured in the tens of thousands or maybe hundreds of thousands, but definitely not in the millions. That's a problem – if you want to roll out IDS to

people's homes for example, they're going to have a tough time getting anything useful out of it because the skill in interpreting the data is pretty severe.

***LXF:* Would you say that was the major challenge to IDS? Solving intrusion problems is about not merely reporting the data or even identifying attacks, but actually trying to make some sort of sense of the data, acquiring some intelligence and working out what is going on on behalf of the user?**

MR: Yes definitely. This is where the whole correlation space comes from so people are trying to do correlation so they can take micro-events and turn them into macro events. That's interesting and there is a lot of interesting work being done there. I think that ultimately we get to having better IDS, and this is what my next project is based on, by having more context about the network environment we are operating in.

In other words, if I see an attack on my network, the way I figure out whether that attack was even capable of succeeding in the first place is that I go an either look it up on my network diagram or I look at the box and see what services are running and test if it's been hacked. So, IDS needs to get smarter and have the context about the network – the classic example is you get a 'Code Red' notification out of *Snort*, and you look at the box that got attacked and see that it's a Linux box which can't even be vulnerable to that attack. The IDS doesn't have that context, that one piece of information about the network – this IP address is running Linux which means that it can't be vulnerable to the following 500 things or whatever. I'm developing a system that will generate this contextual information and give it to the IDS. Then the IDS will do its job in term of the targets on the network instead of doing it purely in terms of the traffic.

***LXF:* that sounds really interesting. When is it going to be ready?**

MR: It will be ready this fall. It is a product, not an Open Source system. Pieces of the idea will show up in *Snort* though, because *Snort* has to get smarter to understand the contextual data. So we'll have an automated system for producing all this data and *Snort* will be able to understand it, but if *Snort* users want to manually input the data they can do that as well.

***LXF:* We recently got a press release about a product, Real-time Network Awareness?**

MR: That's what I'm talking about. The idea is if we are truly going to make intrusion detection better, then IDS has to start defining the targets on the network, not just defending the network traffic. So the only way to do that is if I know what's on my network. There are all these ambiguities that IDS has to deal with now – it doesn't know what OS you are running, it doesn't know how many hops it is from a host, it doesn't know the path MTU. If I don't know those three things there are a ton of things that can be used to evade me or make me report false data. If I can generate all that information, even that basic information, I can reduce the evadeability of the IDS and reduce false positive. If I can take the next step and say I know these vulnerabilities are available at specific IP addresses, if I see those vulnerabilities exploited then I'm very interested – if I see anything else I'm not so interested – so I can prioritise appropriately.

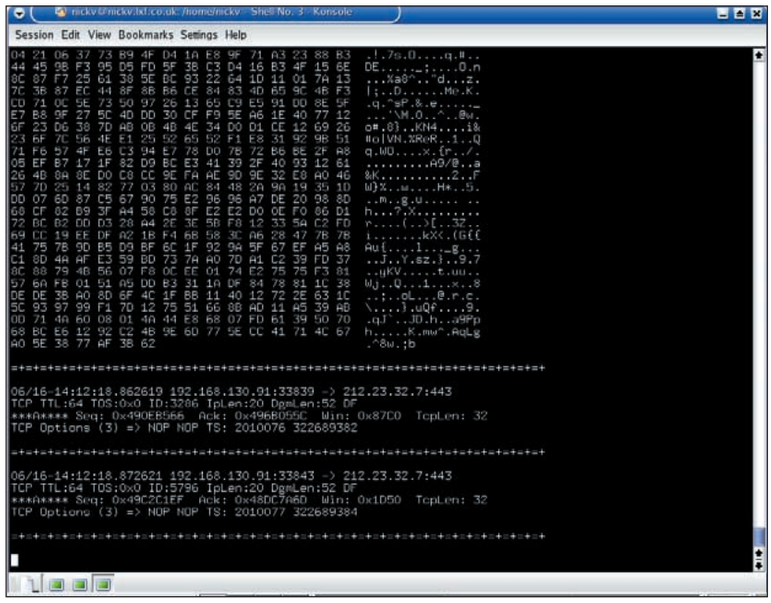

If you know your network packets, you might be able to work out what's actually going on here...

***LXF:* Do you think that false positives are a real problem currently? IDS generally report so much that users take them with a pinch of salt...**

MR: Well certainly. That's one of the primary problems of IDS. Without the contextual information we give the users a lot of work to do – we force them to turn into detectives and try to figure out if what the IDS is telling them is true, and if it is true, then if it's relevant or not. We are seeking to eliminate that and say 'these are the things that are truly interesting – here are attacks that try to exploit vulnerabilities that you are actually vulnerable to'. Plus we can do things like detect change so I can say "your webserver is running *Apache*, I saw the *apache* chunked encoding buffer overflow and then 30 minutes later your webserver started offering IRC as a service as well, is that interesting to you"?

"To make intrusion detection better, IDS has to start defining the targets, not just defend the network traffic."

***LXF:* I think that covers everything we wanted to ask – is there any other advice for our readers?**

MR: I think from the perspective of an Open Source developer, I started developing *Snort* as an Open Source system just to see how the development methodology would work, and I found it's a really great way to develop software. To all of your readers I guess I would say if they are thinking about starting an Open Source project, I would highly recommend it – it's very interesting and very fun. ■