Eye on Security

The Making of an M-Trends Report

Transcript

Luke McNamara:

All right, welcome to another episode of the Eye on Security podcasts, I'm your host Luke McNamara. Joining me today to talk about the latest M-Trends report. We have Steve Stone, Senior Director for Advanced Practices and Regina Elwell, Senior principal Threat Analyst on the Advanced Practices team. Steve, Regina, how are you today?

Regina Elwell:

Good, good. Thanks for having us.

Steve Stone:

Good. Luke, how about you?

Luke McNamara:

I'm doing well. Where I think makes sense to begin with this conversation about the latest M-Trends report, newly released, is maybe first talk a little bit about the process. I think to do that well, it'd be great to hear a little bit about your team, the Advanced Practices team, what you guys are working on, and your role specifically in putting this report together.

Steve Stone:

Regina and I both work on the Advanced Practices team inside of FireEye. So, we are an internal team designed to provide direct support, and then really do across all of FireEye combination of decision making based around threats and technology, and then active pursuit and illumination of a range of adversaries. Our overall team mission is pretty straightforward, it's just know more about the adversary than anyone, and make that knowledge actionable for FireEye.

So, really what we're here to do is work with each of the distinct business units, IR Managed Defense product, those directly helping our clients around preventing, responding, and remediating to intrusions, we provide direct support to the incident response and manage defense engagements. Then we also provide, back to those organizations, some context and larger picture of what they're seeing in the specific engagement. We try to bring all of FireEye's intrusion knowledge to those engagements, and then take those engagements back, make that part of the larger FireEye knowledge store, and really what we do in Advanced Practices, is we focus on the intrusion deltas.

What are we seeing is at all known before, have we not run across it? Do we need to expand the attribution surface area for the technical intrusions? Do we need to deploy weak signals and surface under reported intrusions? Close those deltas, and then move on to the next intrusion event. So, really M-Trends is, putting near and dear to our heart, because this is really a manifestation of our daily work that we do with all these other teams across FireEye, and it's really neat as a company to then get to work with those same teams putting together M-Trends and be a part of that holistic view.

Luke McNamara:

I think it's certainly the most quantitative intensive report that we produce every year. So, I always enjoy reading that when it comes out and seeing the bigger things that we identified as notable trends across the threat landscape and broken down in all the different ways that you do. So, we'll get into that a little bit. But I'm curious, like the actual process of putting this together, because I know this is a month long endeavor, these cumulative threat landscape reports are becoming commonplace, I think, across the industry, but this is a report that we've been producing since 2011.

Luke McNamara:

How do you start with looking at what are the datasets we're going to include in this? This is obviously focused on breach data. But when you think about all the different parts of the report, the case studies, how does that all come together in terms of the planning and the process that you think about with the execution of this report?

Regina Elwell:

To address the timeline, M-Trends is an every day, year long effort for Advanced Practices, like Steve said, it's a direct manifestation of our day-to-day work. The report itself takes probably four to six months to compile, but throughout the year, we're categorizing and classifying all that frontline incident response data to be able to pull metrics after the reporting period is complete. Right now we're already seven months into the M-Trends 2022 reporting period, to put that into perspective.

To the data itself, we are providing that community with information that we're seeing on the front lines, and all the metrics are derived from actual incident response investigation. So, meaning consulting investigations, so direct work product of that. It takes a huge effort for us to distill all that information, to vet it, to analyze it, to categorize it, and then put it back into our company knowledge bank. Then when it comes to writing the report, we're looking at, what are the trends? What's change year over year? Is there another subset of data that we should be looking at that we maybe haven't delved in before?

And also looking ... like now we're doing the process of looking for next year, was there something that maybe we wanted to account for this year that we didn't have a way to track that we should be tracking for next year. I guess ransomware is a good example of that for this year, that we saw that ransomware was much more prevalent this year compared to the 2020 reporting period, compared to previous years. So we broke out ransomware specific metrics to show those numbers for that.

Luke McNamara:

Yeah, that's something I definitely want to get into is some of the, I think, unique insights from this past year, and a story that I think you guys had a very excellent job in telling where there was a lot of nuanced language and breaking down things like the dwell time, but separating out how ransomware has played a part in that. But first, one of the questions I have, again, back around the process aspect of this, how have you seen this report change over the years? This is a report that obviously we've been producing for some time.

I went back actually a bit before this and looked up the very first, the 2011 report, and there's some things that some of the language and aspects of it that certainly are still pertinent today, but the amount of data that you see in this report now is by far and away a lot more than the kind of early versions. But I'm curious, any insights that you've seen, well, just how the report itself has morphed over the years?

Regina Elwell:

It's definitely an iterative process. I mean, you can see the evolution report to report and every year we end up adding some extra metric, it seems like it very rarely do we drop a metric. It's just like every year we're adding more and more, trying to give the community more insights into what we're seeing, to be able to have a better grasp on what is truly a trend year to year, what is an anomaly, and it has been interesting to say like, "This year, adding that ransomware metric because it is a very important component to that threat landscape and staying with the most relevant metrics for the community." Yeah.

Steve Stone:

I think also, just piggybacking on Regina's comments, one of the really interesting things as we've worked M-Trends, has been the ability to take things that we see and really provide visibility into that. I think one of the real guiding principles we have for M-Trends as a company is, this is a show our work kind of product, we don't want to just say, "Hey, you should think about this, or Hey, what about that?" We want to actually provide context on why we are saying this is important, here's some things to look at, here's some priorities, some of those aspects. And ultimately, I think is, Luke, to your comment about the volume, especially on the metrics has just increased year over year, that's been very intentional.

Again, back to Regina's earlier comment about this as a daily thing across different teams where we're really trying to drive workflow changes forward and use M-Trends to do that. So as we've gotten better at things as a company, you can see products like M-Trends get more robust. But all while keeping that kind of mentality of Kevin, and Jurgen in particular, are huge proponents of this is our chance to be very open with the larger industry and our clients and an organization that don't have anything to do with us. You can really hopefully see that come through this document. This is definitely designed to be a show your work kind of manifestation of the company's efforts.

Luke McNamara:

You mentioned this is an ongoing process, and this is something that your team is obviously tracking and focused on every day, not just the period where you're putting together the report itself. But I'm curious, are there any particular data sets that maybe are more intensive or difficult to sift through and collate when you're putting the report together itself, where that starts to be in some of the more intensive parts of the work?

Regina Elwell:

The most difficult part of this process is that, for a Mandiant consulting investigation, no two are identical. So, you're having to distill that information and categorize it, and you have potentially super disparate pieces of information that you're trying to pull these trends and be able to distill that information down from. So, not knowing what customer blogs are going to have, or what is provided during the investigation or what's discovered, because it could be anything and anything under the sun that we then have to standardize and put in some uniform format to be able to have that metadata to pull metrics from. So, just the threat landscape itself and how different it is, makes the whole process of pulling metrics on that difficult.

Luke McNamara:

So you mentioned ransomware, and this, I guess, gets into some of the takeaway pieces of the actual report itself. I think probably I'm not alone in this, but every time this report comes out, I think the first thing I virtually flipped to is the dwell time metric and see how that's changed. I think that's probably one of the pieces of this report that's probably most well known. Can you talk a little bit about what unique about that this year? Because it seems like there's a little bit of good news, bad news, the bad news being how ransomware has impacted that, but then some positive news, at least compared to last year in terms of internal detections trending back up.

Regina Elwell:

Yeah, so for internal detection, absolutely, companies are doing a much better job of detecting those incidents without that first notification coming from that external source. They're doing it in house, they're developing those capabilities, but also something that is in those numbers that we had to tease out is that ransomware is a factor in that improvement number. Because ransomware threat actors have no desire to stay quiet in the environment, it's destructive, it's loud, their goal is to get paid. The only way they're going to get paid is if they draw attention to that breach.

So definitely, there's some percentage of that that is impacted by that ransomware. And we see that in the dwell time too. There's no need for a ransomware threat actor to stay super stealth for extended periods of time in an environment if their goal is to get in, get access to as many systems as possible, encrypt everything and then demand money. It's a completely different type of threat that you're defending against compared to some espionage long term intrusion, that their goal is just to be in there, collecting information over extended periods of time to stay still.

So, seeing that dwell time number drop drastically, definitely is impacted by a huge spike that continued surge of ransomware in the proportion of investigations that we're responding to.

Steve Stone:

I think too, as we talk about, especially at the global level, the global media and dwell time, one of the things as we talk about that a lot internally, and even then we do follow on specific engagements around how to interpret and use that data, the biggest thing I think we get sometimes is, this really is designed to talk about everybody, the global scale of that. So, we get the question all the time of ... we've actually gotten a few of these this year. As a company, we're saying, "Hey, we think this 24 days global median dwell time for this M-Trends is significant, we think it speaks to both ransomware, and also clients improvement, industry improvement, whole of government improvement."

I mean, I can't speak for every one of the presentations we've done, but all of them that I've been in, we get the question of, "Well, that's too long and attacker can do a lot in 24 days. Shouldn't that be under insert timeframe?" Our position is, "Yeah, absolutely." But those two things are very different. What we're not saying is, "Hey, 24 days for an attacker in an environment is great." What we are saying is this, year over year, over year improvement, going from 416 days down to 24 in this time frame is significant, and it does speak to a holistic, larger improvement.

I think that's one of the things as we talk about this that can be easy to overlook, because we're wired that way, we deal with single environments, we deal with our company, our organization, our breach, and we really want to use M-Trends to talk about ... when you put all that together, what are we seeing across all of those events. I cannot reinforce what Regina said enough about as we look at dwell time and detection by source, which we assessed years ago, those were important because they were impactful, and we continue that forward, it really comes down to organizations that can find their own intrusion, are typically better prepared to deal with that intrusion, and you see that play out in global median dwell time.

The difference from 24 days, into the external versus internal categories is stark. I mean, internal detection is half of that global dwell time at 12 days. External notification is more than three times that average at 73 days, I think, off top my head. So, that's one of those things, I think, is also easy to get lost, because it looks like a different way to track but it really does speak to all the questions we get all the time across every business unit, "What can clients do to help themselves? What can clients do to be better prepared?" That's why we talk about internal detection, external notification. That is one of those things we think is the most important, is put the effort there, put the resources there, and you can see the dividends on those investments, or conversely, what happens when you're reliant on other organizations.

Luke McNamara:

So this is one of the things in particular, I really wanted to pick both your brains about,

because that internal detection metric, I think you're right, a lot of people look at as a proxy, not just in the United States, but when you're looking at the specific regions that you have broken out as an increase in organizational maturity and the capability of organizations in these different regions to be able to deal with different types of cyber threats. So it's an improvement over last year, the uptick in the number of internal detections, but looking at it for the last several years, it seems like it's kind of flat lined a little bit, hovering around that 40%. It's gone up and down a little bit, but obviously, a big improvement from where we were in 2011.

But do you think that we are starting to see a normalization of the amount of internal detection? Again, kind of viewing that as a soft metric for the capacity of organizations to detect those threats internally, and then conversely, a metric for their maturity and their capability, are we kind of evening off at this point, or do you think we'll continue to see a downward trajectory?

Steve Stone:

Yeah. Great question. So, I would say off the top, I'm not sure. So, everything I can give you is an opinion. Right? I mean, looking forward to the intrusion landscape is always challenging, because every year, we think we know what the next year brings, and then we're writing M-Trends, and we've learned that there's a bunch of stuff that came up that we didn't see, and this year proves no different. So, on the detection by source, I think we're definitively in an area where we don't expect massive jumps one way or the other.

I would be shocked if we're writing M-Trends 2022, and the internal detection is at 80%. I just don't foresee that. We would hope to see that kind of incremental improvement, and it just keeps working and working and working. And again, this is easy to take this discussion around detection by source, and have it turned into what's wrong with a certain organization, but again, I think, there's a whole element here that we openly advocate for as a company. We want other entities to be involved. We openly advocate for industry regulations, we openly advocate for government actions, those entities and those actions then produce external notifications. That's great. That's not a bad news story and of itself, we want governments to improve. We want industry partners to improve, we want our competitors and peers in the vendor landscape to improve. That helps clients, that helps everybody.

So, I think that that's a piece as well. And those are very difficult things to forecast, it's almost impossible to know how, let's just take the government piece, those changes take years to manifest, especially in a holistic way. I mean, look at last week, the week that we rolled out M-Trends, we were dealing with government take downs of web shells, we were dealing with government arrests and significant sentencing for FIN7 personnel. We were dealing with government sanctions, in part, because of breaches that we talked about in M-Trends. Those things, we did not forecast what all happened the week that we rolled out M-Trends, they're just very difficult to forecast and when they'll manifest.

Then again, it's almost impossible to know what other vendors are working on that will then produce these positive impacts. So, that was tough, but what we would hope for and expect

is an incremental improvement year over year. We just don't think there's going to be a big thing that's going to come in and swing everything, we've all been doing this long enough that those silver bullets just don't exist. So we then look for what are these incremental progressions that help on these really, really difficult topics. Oh, by the way, the bad guys get a vote too. They innovate just as hard as the rest of us do.

And ideally, some of that comes through in this M-Trends, you can look at just the ransomware stuff we discuss alone, we used to discuss just a couple years ago, ransomware as a fairly monolithic action. If you look at M-Trends, especially in some of the case studies, you can see we spent a lot of time talking about the evolution of how actors are using that, the evolution of the deployment timeframes alone. So, there's a lot of players involved to get a lot of votes, which make forecasting that pretty difficult.

Luke McNamara:

Yeah. There's obviously a lot of factors, even in the data collection there of where our customers are that we pull this data from, the specific regions they're in even the industries that they're in, that all could have an impact on their organizational maturity. But I think it is an interesting thing to look at year over year, how that specific metric changes, and then what that could be telling us, again, if you view it somewhat as a metric for the overall capacity of the organization to deal with and respond to threats, and is that maturity going up or down in certain regions, I think it's interesting to look at.

We were talking about ransomware, I think the metric in there was about 14% of the incidents that we responded to last year or in 2019 rather involved ransomware of some kind, and that jumped up to 25% last year. It seems to be that no one's expecting that to slow down anytime soon, but was there any noticeable trends from what we saw in terms of further actors entering the space, or obviously, a big shift in the tactics we saw last year was just how creative they got with extortion. But what were some of the big takeaways from ransomware this past year?

Steve Stone:

I think when we look at ransomware, the first trend is the one you just mentioned. I mean, we just see more of it. I mean, that jump from 14 to 25% of all the incident response engagements that we worked in addressing this M-Trends period, that's pretty significant. I mean, we're dealing with one out of four, and again, as Regina mentioned, we're seven months into the next M-Trend, back-of-the-napkin math seems to be that increase is continuing into the next M-Trends. So we'll see, right, fingers crossed, we'll see how that goes.

But the thing I would say with the ransomware as big trends is, there's a reason these numbers are going up. I think when we look at the motivation section, I mean, I really think it's important to pair the ransomware discussion with other parts of what we see in M-Trends and then just obviously in event response. It's working, that's why it's prevalent. I mean, I don't want to be overly simplistic way, but if you look at the section we talk about attacker trends and some of those areas, if you look at just the motivation piece, the highest

single motivation percentage that we assess is financial gain.

If ransomware pays off, actors are going to use it, they're motivated to make money, and that's not new, and they're just finding what they believe is the best, most effective way to do that. What we're seeing is there's more and more of those actors that are assessing that ransomware as the way to do that, it is just that simple at some point. Again, I know that runs the risk of being overly simplistic, but these actors are doing this to get paid, period. Some high majority of them, and if ransomware gets it done, if that generates them more money, or more money per intrusion, or in an easier way, great.

Again, I'll keep stealing Regina thunder, she mentioned that, the point that a ransomware actor can show up and move very fast in environment because they know they're going to get caught because they're going to deploy encryptors, and it will be very obvious, that's got benefits if you're a bad guy compared to an espionage intrusion, you can just do things and not have to do things that traditional actors that I'd be focused on, I think we would be remiss if we didn't recognize those elements have a direct contribution to the uptick in ransomware.

Luke McNamara:

One of the things I thought was actually interesting, and one of the metrics around ransomware was noting a decrease in ransomware in APAC, every year from 2019 to 2020. I think it makes sense when you think about, again, actors gravitating to where the money is, as mentioned, certainly that's going to be in the type of threat activity that they are going to take part in because they're financially motivated, but also makes sense, and I think we've seen this in the past with banking Trojans and other forms of cyber crime, where they also gravitate to the regions where they can maximize the most amount of gains.

I'm curious if anything that you've been seeing so far into this year, suggests that we might see ransomware, it's certainly not alone. I mean, the America is certainly not the only region where we've seen ransomware, EMEA and APAC as well, but do you expect those two regions in particular to see more of that this year?

Regina Elwell:

I mean, I guess what Steve said earlier, it's hard to predict the future. Obviously, we saw a decrease in a APAC last year, but EMEA had an increase as well as the Americas in 2020, for that period. So again, I think for these types of attacks, the attackers are going to go where the money is, and where they're actually going to get paid. So, if they're going to have success and continuing to have success, they're going to go get similar victims for those types of attacks. So, it's hard to say if we'll see increase in a APAC or not aa the numbers did not show that for 2020 though.

Luke McNamara:

Moving on to talking about groupings. I think this is an interesting thing that maybe more so than other previous reports, correct me if I'm wrong, it seems like we're talking a lot about the young groups that we've seen merging specific APT and FIN groups that have been

active this past year. One of the metrics that jumped out to me that I thought was particularly interesting was the note that ... I think we've seen an almost doubling ... yeah 15% to 29% of multiple threat groups active in a victim environment, when we responded to a particular breach.

I think that speaks to the necessity and importance of being very careful and accurate in the cluster of the different threat groups. But can you talk a little about in terms of what we were seeing this past year, with respect to the number of groups that we merge graduated, the number of groups we saw emerge, and new groups on the space, and what that speaks to in your mind for how the threat landscape is changing?

Regina Elwell:

Yeah, so we had 650 plus newly traction groups in 2020, and then 246, that we actually observed during a Mandiant investigation. Then the subset of that, 161 of those were new groups for that year. So the other half being ones that we were already tracking and had seen. We did a big blog earlier this year “DebUNCing Attribution” that really goes into our process. So definitely check that out for the whole merge graduation, like how these groups work.

But we do spend an extensive period of time clustering groups and being diligent of keeping those as separate clusters, and too, we have proof and significant overlap between groups. That merge process really is an extensive process of looking at those overlaps, what TTP do they share, what malware do they share, what timeframe did it happen like similar targets? Looking at the whole spectrum of the overlaps between these groups, and if there are significant overlaps, and we can justify, and in our minds prove that these two groups are the same, we merge them together based on those similarities.

We also did graduate one UNC group into a named APT group to FIN11, and that was our first FIN graduation since 2017, I believe, and that is a really extensive even more so than the merge process like deep dive into this group fully vetting, figuring out as many individual operator details as possible, really going deep dive into that. We're saying, this is a very clearly defined group, and we've tracked it over an extensive period of time, like, this isn't something that just popped up that we think may change, we're very confident that this is an established group, that major group, FIN11, that we did this year.

Luke McNamara:

It was that number you mentioned, of the number of UNC groups that got merged. Was that on par for what we've seen in terms of previous years, or did we do more merging of UNC groups in previous years?

Regina Elwell:

I have the numbers offhand for previous years, but we are trending up, our team in general is growing, I feel like we're having a little bit of extra cycles to be able to spend this extensive time period on these types of overlaps, and we're getting better at it. There's another blog out, Adamicity. So data science process of looking at the overlaps of these

groups and not only like the analysts gut feeling of noticing a couple of things like actually using the data science to say like, "What are the similarities? What group is this most similar to?" Over time, we're definitely spending more time on that, we're getting better at it, and we're doing more of these year over year.

Steve Stone:

I think too, the groups discussion is a great ... like if I had to give one touch point about the show your work mentality of M-Trends, it would be the groups, because I think this is our second year, where we started pulling the curtain back and an M-Trends giving the visibility into how many groups we're tracking, where we are seeing them, what we're seeing them do. But that's not new, we've always done that, that has always been a behind the scenes motion, and a key part of our incident response. For Regina and our team on Advanced Practices, a key part of what we do.

This is directly ... when we're working with incident response consultants, and let's say the overall engagement manager, she comes to us and she's like, "Listen, I need to understand what's going on here, we need to tease out all this activity." Part of our job as Advanced Practices is to come in and say, "Okay, this is one group, or this is two groups, or, hey, there's at least two here, we don't know what's over here on this third one, we need to figure that out." That has always been part of our incident response process as a company, and we've always found great value in that.

I think the irony is that's how it shows up in M-Trends, if you'd have asked me several years ago, would we ever reveal the kind of group trade craft, I would have said, no, just from a company perspective, we just didn't reveal that, and we didn't talk about those things, especially publicly. We were in discussions with the IR leadership, and their perspective is, this matters to us, we spend huge amounts of time on these topics writ large, and also in specific engagements. The entire goal of this is to do the right incident response and remediation based on what you're actually dealing with not, "Hey, what's kind of a cookie-cutter approach?" Because that's proven such a valuable tool for us, now, it's M-Trends.

I think that's really interesting, because this isn't marketing, this isn't something new, because we got to be splashy, we've always done this, and then eventually got to our point where the instant response leadership just believed that it was important to talk about this publicly, because again, our goal here is to really say, "Here's the things that we think are important, we hope this will help inform other organizations." And it became problematic to continue having that discussion, and not discuss threat groups and malware clustering, because we spend so much time on it, and derive so much value out of it.

Luke McNamara:

This is something that's been interesting to see across the industry. There's obviously a lot of different philosophies around this, particularly around the granularity you get with how specific you get with clustering things often to separate entities. I think the whole process detailed around UNCS is a great example of how the intent is to, if we cannot immediately link this to an unknown threat group, then it gets its own UNC. But also you see some

organizations that have very massive threat naming schemas that encompass what we might track as multiple UNCS or even a APT or FIN groups.

I think that at an organizational level, particularly on the business side, you may have some entities that all they care about is that it's North Korea, not that it's specifically APT 37 or 38. But from a very tactical standpoint for folks in a sock that are trying to prioritize which detections should they chase down or knowing the tools, the malware used by specific groups, that seems to make a lot more of a difference. So, I guess, what are some of your thoughts around how granular you should strive to be in terms of separating out specific clusters?

Steve Stone:

Yeah, so I would love to jump on that one, that is a topic very near and dear to my heart, and when we talk about a lot. So, short answer is the goal ... and this is definitely the goal we operate as a company and then again, kind of going back to that core task of Advanced Practices working with, the frontline folks to suss this out, our answer is as granular as absolutely possible. Like you can't be granular enough. I think this is one of the things ... it's always easy to look back in hindsight and think about how you could have landed certain concepts differently or discuss them differently.

One of the things I think we could have done better and I love that our company has gone public with, how we do clusters, how we do groups, I genuinely love that, it's a really important topic. I think one of the things we probably could have done a little bit better, and we deal with this in M-Trends presentations and conversations is, it's important to keep in mind this construct UNCs and the APTs and FINs, it was not designed for ease of use for marketing, it was not designed for ease of explanation. It just wasn't, it was designed for the most detailed, technically driven, incident response possible.

That's why we have these, and that's got pros and cons. Every model has pros and cons. There are other models that are much easier for a typical person to stand up and talk about or summarize whole swaths of activity in a cleaner fashion than typically some of our service delivery people have had to do, we recognize that. But again, we didn't design it to do that. We designed it to guide incident response. There's lots of folks that have been around and contributed to this over the years, and the philosophy that was always built in was, UNCs are cheap, you can always merge an UNC. You can always decide down the road, "Oh, I actually created a net new cluster, we did know about that, let's blow it away."

It is designed for technical accuracy, and then the second part is designed to identify how technical intrusions are evolving, and then get after that. That's why we have UNCs. That's why this entire model exists. It's been interesting to watch that model, which we've used for, I mean, as long as we've been around, I think we've had UNCs for all but the first year of Mandiant as a company, and the use of that and the extrapolation of that, I love we talked about it publicly, I think it's got value. But again, it ultimately comes back to, we're talking about all of these UNCs and where they're at, because that's how serious we take that granular level, you should be responding to a ransomware group totally different than UNC2452.

If your response and remediation plan cannot deal with those two distinct efforts, you're going to have a bad day, because the response is completely different, and that granular tracking down to the specific ... I mean, you name a technical detail, and we track and UNC to that level, because we think it is that important and it's very easy to get attribution wrong, always has been, but as the landscape continues to get more diverse, it's getting more challenging, which we believe the solve for that is to get more granular. We talk about cobalt strike beacon, beacons and M-Trends multiple times. It was the most prevalent malware family we saw using intrusions, even though it's not a malware family in and of itself.

We see UNC2452 use it, one of the most advanced groups we have ever seen. We see FIN11 use it, we see everybody use it. It's not enough to track cobalt strike beacon anymore. You have to track the hundreds and hundreds of variations in iterations, and if you cannot, you don't have the ability to suss out, is this ANC2452, cobalt strike beacon, or is this a pen test? This is a red team that a company hired and they're using cobalt strike beacon. That's why we believe that granular nature is so critically important, because when you're in it, and you've got to make a decision right now on what's going to happen next, and that's going to have profound ripple effects, that technical granular detail is what's going to enable that decision making.

Luke McNamara:

Yeah, that's a perfect segue into actually one of the other things that I wanted to ask you is about some of the malware stats in here. So you mentioned beacon, beacon an empire both showing up here a lot, probably no surprise in terms of what ... people have been tracking the space and the usage of publicly available tools and malware things that we see, pen tester is using but also as you mentioned some of the most advanced threat actors that we track. I thought it was also interesting, and maybe this ties into your point about the variations within beacon. I'm not sure if it's represented in this number. Let me see if I get this quote, right. Of the nearly 300 malware families observed by many experts during intrusions, 144 malware families were malware families which made it begin tracking in 2020.

So about half of those new malware families, we saw for the first time last year. At the same time, we have this continuation of this trend, which has been going on for several years right now, where some of these publicly available tools are being widely utilized by a range of threat actors. Are some of those variations of beacon accounted for in that or how do those two stats play together?

Regina Elwell:

So, I guess the broader point here is adversaries absolutely are going to use what's available to them both in the environment and publicly available tools to make their lives easier. But we're also like half of these malware families are customed to them that they are innovating, they're adapting, they're doing whatever they need to do to be effective in target environments, and also not using public tools.

I guess, if it was specific custom, that it was something beyond beacon itself, it would have been tracked differently, but if it's at its core beacon, it's tracked as beacon. So, those numbers would possibly be slightly inflated for that beacon number, but still at its core, it's still beacon, it's still that publicly available tool, even if it is using some custom campaign code or anything of that nature.

Luke McNamara:

What do you make of that trend? Again, maybe similar to the ransomware one, I'm asking you to be a little bit predictive in terms of what you might expect to see down the road, but that also doesn't seem to be something that we're going to see an end of anytime soon. For a lot of different reasons, I'm sure if there's counter forensics and counter attribution reasons why some of those threat actors will gravitate to those tools, ease of use modularity might be another one. But I mean, you mentioned 2452, that's a threat actor that we've seen utilize beacon alongside some of their own custom tools. What do you make of that trend?

Regina Elwell:

I think that was something that we've seen in the attacker landscape, for as long as we've been tracking is that, threat actors are as sophisticated as they need to be in environments or as they feel they perceive that they need to be in environments. So, if beacon still works, still isn't detected, they aren't kicked out of environments, for all the reasons that you just mentioned, they're going to continue to use it, there's no need to bring out special custom tools that they may blow in this environment and eventually get detected, if they can get through this portion of the compromised using a publicly available tool.

So, beacon gets discovered, it's really not a huge hit to their potential arsenal that they've spent delving this custom tool, if that custom tool is just discovered in an environment, obviously, that's a much greater cost to that threat actor.

Steve Stone:

I think as well, this only speaks to the landscape changes for everybody. I think one thing that we see year over year over year is exactly to Regina's point, the threat actors will do what works, and they'll start at the bottom, and they'll work their way up. We love on the defense side of the house that there are more readily accessible tools that organization can bind implement, everything is much more capable out of the box, that's just a trend that's great for everybody. Same thing on the attacker side, there are more readily available tools that can do more out of the box, the cost of entry can be lower for a certain threat group. So I think that's an important piece.

I think the other piece is, again, to keep going back to Regina's points, these actors are learning too, and they've been doing this forever, they have always sought to use the lowest level of skill while also having to not declare who they are. These kinds of tool mixes the mix of custom and public, the mix of highly evasive, almost single use malware families with cobalt strike beacons to serve that purpose. On 2452 again, is another great example,

we don't know but we assess part of why they use so little malware is, it speaks to the evasive, stealthy nature, that was a driving force for them.

And then if they do have to use a tool, great, use a very good cobalt strike beacon, you blend in like every one of the other hundreds of actors using it, versus a super top shelf thing that now we know like, "Wait a second, this? We've never seen before, we just got our socks blown off, let's now get after this." Why would they do that? It doesn't serve their purposes for what they're trying to do. So the same mix of custom tools, public tools, benefits both the lower end actors and the higher end actors, and again, just to go back to the specificity, this is why we think that's so important.

Then again, one thing that we see if we move off of cobalt strike beacon and even some of the custom UNC2452 tools, we've seen tools that used to be cardinal become very public. SOGU was a perfect example. It was not that long ago in the olden days, when one threat group used SOGU. We knew when we saw SOGU, it was this one thread group. Those days are gone, they're absolutely gone, everybody uses SOGU. Oh, by the way, SOGU has evolved pretty demonstrably over the years. It's a more capable tool than it was the first time we encountered it. So, everyone's evolving, everyone's working, and everyone has different purposes and motivations, and the task is on us to suss that out and make sure we're doing the right thing for the right reason.

Luke McNamara:

Yeah, I think we saw an example of that last year too with APT36, one of their tools started being utilized by entities not believed to be APT36, and that obviously has ramifications, I guess, for how you think about the tool that previously you'd have assessed had only been used by one particular threat actor, or maybe several actors from a particular sponsoring entity. But I think the point you made about 2452 usage of beacon, again, is a great reminder, if it had not already been obvious in the last several years of how sussing threat actors, their capability, solely based on how advanced the malware is where the usage of zero days is not necessarily the best metric for some of the most stealthy threat actors that we look at.

One of the questions I had, I guess, as you mentioned, Regina at the beginning, this is something that is an ongoing focus of your team is looking at these trends as they emerge, not just the period of time where you're collating all this for entrance, but I'm curious, was there anything as you were in the process of putting this all together that surprised you or stood out that maybe wasn't as apparent in your day to day up to that point?

Regina Elwell:

I think the best thing about M-Trends is that it really is that level set moment of what actually occurred through that time period and what the real trends are, because every responder, every intel analyst, every person that has a piece of this puzzle, only sees their piece of the puzzle. So, if you go and ask a bunch of different responders, what's the most prevalent thing or what do you think was the most impactful, they're going to be biased by that set of data that they have available to them. So, one person might think that

ransomware is a lot more prevalent than it is because they've done only ransomware investigations for the past year or examples to the other extreme.

I think that it's really important that when we get back here that we always go into M-Trends with a complete open mind like, "Show me what the trend is." Rather than coming in like, "I want to prove that the x happened in this time period." We definitely get requirements or lists of metrics that, "Oh, we'd like to show this new thing this year." Sometimes we pull the metrics, and we're like, "Well, that actually wasn't a trend, or that really was insignificant in this process." I think to Steve's point from earlier, something that always shocks me a little bit is how prevalent financially motivated actors in general, how that's becoming such a bigger piece of the puzzle. And obviously, ransomware is a component of that.

But to have more than a third be financially motivated, like if we go back to the trends several years ago, like that definitely wasn't the case. So it definitely shows that shift in the total threat actor landscape from over the years.

Luke McNamara:

All right, final question. I'm taking UNC2452 and FIN11 off the table, but favorite threat actor from this past year, if you had to pick one.

Regina Elwell:

UNC1878, probably my favorite of the last year, but FIN7 is, hands down my all time favorite.

Luke McNamara:

Any particular reason why for UNC1817?

Regina Elwell:

I just respect to their trade craft, I guess, and the consistency of it, I guess to the same for FIN7, like I always looked at FIN7 like they're professionals, like this is an actual business to them. And like there's ... I don't know, maybe I shouldn't have respect for them, but on some level, you respect their process, isn't just a random hack. There is intent and true motivation behind that.

Luke McNamara:

That's a good point. Steve?

Steve Stone:

Man, I think for the last year, it'd be hard not to have it be UNC2452, we spent a significant amount of time working against that group, learning that group. I think what I really grew to appreciate about UNC2452, and we kind of really had to get our heads around is, I think one of the things that has been under appreciated about that group is, they do all the little things right, it's not unusual that we see a group trying to be evasive, and they do certain

things in a certain part of the attack lifecycle to be evasive, and they are, but then we pivot left and right, and we can get back on the trail and we can find them and now we're after it.

As we would go after that, with UNC2452, we would pivot and be like, "Oh, wow, they're just as dedicated here. This is crazy. Okay, let's pivot again. Oh, wow, they're just as dedicated here." It was a group that clearly understands incident response, which is rare. They really understood and in many ways, anticipated some of the things we were going to do, which made that highly challenging. On a purely personal level, I have never worked with so many professionals on a single group, which is probably why I like UNC2452 the most.

We brought in so many experts from so many parts of our own company, and then also some of our partners to really try to get our arms around that and that's unusual. That's not normally as an industry, how we come to understand groups. We all kind of work it on our own, and we deal with people who are directly involved, and then we talk about it. This was totally different, responding to UNC2452 had a totally different feel, which gives them a special place in my heart for this year.

Luke McNamara:

Okay. Well, that's a good answer. So, I'll allow UNC2452 for that one. Steve, Regina, thank you both for your time and all the work that your team and others at FireEye spent on this report. For anyone that has not yet, please go check this out, there's several right talk webinars that Steve and Regina have done as well, check those out. Steve you're on Twitter, where can people find you?

Steve Stone:

I'm [StonePown3000 00:45:33] on Twitter, and that's about it, I'm pretty social media light except for that.

Luke McNamara:

Excellent. And Regina?

Regina Elwell:

Similar. I'm Regina Elwell. So, I am my myself.

Luke McNamara:

Fantastic. Well, great insights as always, great talking to you both.

Regina Elwell:

Thanks for having us [crosstalk 00:45:52]

Steve Stone:

Thanks for having us on, Luke.

Luke McNamara:

Take care