dspanz. digital service providers australia new zealand

Introduction

Host - Belinda Stewart (DSPANZ): Good morning everyone, thank you for joining us. Good afternoon, for most of you who are on the east coast of Australia.

Just a little bit of housekeeping as we get started today. We have Jordan from ForensicIT, who is presenting for us today the anatomy of a cyber attack. During our webinar, you will all be on mute for the duration. If you have any questions, which we do encourage, please put those through the chat function, and we will host a Q&A session at the end of the presentation. Following our presentation, the recording and slides from this webinar will be made available on the website, we will email and let you all know once they're available

Slide change

We'll kick off by doing our acknowledgment to country. In the spirit of reconciliation, DSPANZ acknowledges the traditional owners of country throughout Australia and recognises their continuing connection to land, water, and culture. We pay our respects to the elders past, present, and emerging. In recognition that we are a binational association, I take this opportunity to acknowledge Maori as Tangata Whenua and Treaty of Waitangi Partners in Aara New Zealand.

Slide change

So, welcome this morning. My name is Belinda Stewart. I'm a director here at DSPANZ, and I also co-chair our security committee with Charles, who is also here. So if we'll go to Charles, we'll do a little bit of an introduction.

Slide change

Host - Charles Gillman (DSPANZ): As we said, I co-chair the security committee. I have over two decades of expertise in information security, specialising in financial services. I have led security teams at major institutions, including CISO roles in diverse settings, offering a comprehensive perspective from a background as a penetration tester, providing practical solutions for current and emerging cyber threats. So, thank you Belinda. I will be hosting our Q&A session following the presentation today.

Host - Belinda Stewart (DSPANZ): Thanks, Charles.

Slide Change

The person we're all here to hear from is Jordan, who is the director of cyber and incident response at ForensicIT. Jordan draws on extensive expertise in cyber incident response and forensic data analysis. He is proficient in threat detection, digital forensics, and mobile device investigations. He brings a distinctive approach to cybersecurity challenges, adept at managing intricate cyber incidents and formulating effective strategies. He is celebrated for his capability to improve cybersecurity defences and establish resilient incident response frameworks.

I know every discussion I've had with Jordan has always been so interesting, and I've learned so much from him in the few times I've met with him. So, I really hope you will get something out of today. We'll hand over to Jordan so he can begin his presentation.

Slide Change

Guest Speaker - Jordan (ForensicIT): Thank you. It was weird hearing someone say lots of things about me, but thank you all for taking the time out of your day to come to this presentation.

Slide Change

Guest Speaker - Jordan (ForensicIT): In a nutshell, we'll cover a little bit of context of the current threat environment and then some steps to take in order to prepare for a cyber incident, the process we would walk you through if you were experiencing one, but also in the context of an actual breach that I think is relevant to your industry and the development environments that you support.

Slide Change

Guest Speaker - Jordan (ForensicIT): Something we like to try and explain is that you may hear me speak between the two concepts of IT security and cybersecurity. You can think of the difference as being one is more akin to your health, looking after your body. You're eating lots of fruits and veggies, you go to the gym, you don't smoke, you don't drink. It's not necessarily going to help you if you get punched in the face. So, it's useful in recovery, but it doesn't help with the attack scenario.

Cybersecurity, you can think of it more along the lines of martial arts training, weapons training, having situational awareness, knowing not to be on King Street in Melbourne at 3 a.m. on a Saturday. So, you know there's a heightened risk of attack. When you apply that to our discussion today, IT security is your preparation, asset change, patch management, having information security policies, staff being aware of them – good hygiene. Whereas cybersecurity is more about threat management, incident response, knowing what to do because taking action is what reduces the risk. In IT security, you're applying controls to reduce the risk. That's the differentiator when we're running through this today.

Slide Change

Guest Speaker - Jordan (ForensicIT): Some context for attack paths. A misconception in a lot of jobs we do is that threat actors just appear in the environment. There's actually a reasonably long process they've gone through to get there. A fair amount of it can be automated. It's not that they're necessarily targeting you. In their reconnaissance phase, they may be targeting out-of-date Citrix NetScaler appliances or out-of-date FortiGate firewalls. They're looking for that in their reconnaissance, then they look at weaponisation, testing the vulnerability, knowing how to deliver it. They deliver it through social engineering, phishing emails, USBs not so much anymore, drive-by compromise of websites, etc. Then they exploit something in your system. Once they're in, they install their applications and malware, establish command and control, looking for out-of-band communications with your network that you can't see.

Recently, this was done through hosted Azure infrastructure in Microsoft's cloud. If you looked at it from a security operations centre perspective, your environment speaking to another Microsoft IP address wouldn't be seen as malicious. It's a good way of doing command and control, getting anonymity in Microsoft's infrastructure to hide. Once they've got command and control, we move into the actions, and that's essentially game over. If you see attack activity only at this point, they've already got access, exfiltrating data, committing fraud, taking intellectual property, etc. A lot of the time it's kicking off ransomware, encrypting the entire environment.

Slide Change

Guest Speaker - Jordan (ForensicIT): So the environment we're looking at, following Lockheed Martin's kill chain is just that process to give you some context of the steps that they take. Each year there's a report produced by IBM and the Ponemon Institute that covers often the cost of a data breach. So, it is severely skewed in the costs by the US and Middle East regions from a few jobs in the last three or four years, it's been a constant thing, we've done some work to average this out.

In Australia, we're looking at around $2 million to recover from an average cyber incident. Just to give you some context from an Australian perspective, ForensicIT has been involved in jobs where the cost ended up, including business interruption costs, up to the high $30 million mark. So, it does get up there in Australia, but on average, it's around the $2 million mark.

Slide Change

Guest Speaker - Jordan (ForensicIT): Some context of how threat actors get in: primarily phishing and stolen or compromised credentials, which have been the two highest for the last five or six years. This Ponemon Institute report has been happening now for about 15 years. They also go hand in hand – phishing is a great way of compromising credentials, but not the only way. The other one we see quite a lot in the financial sector is cloud misconfiguration, people accidentally configuring environment variables for a site or AWS infrastructure to be public, leading to repeated attacks on the organisation.

Impact by industry, the highest cost for a data breach per industry, the financial sector is second to healthcare. Healthcare skews the statistics due to large ransomware attacks on hospitals in the US in recent years. The financial sector is a valuable target to attack, highly regulated, so regulatory fines increase the cost when a breach occurs. In Australian context, the financial services sector is middle of the road. It accounts for a smaller amount of attack activity in Australia, primarily because the sector is mature in its incident response and uptake of cybersecurity measures due to heavy regulation. Organisations we've dealt with have been good at dealing with and knowing what to do in the event of an incident.

Slide Change

Guest Speaker - Jordan (ForensicIT): In Australia, there is noise around cyber incident reporting, with stats from the federal government saying there's an attack every six minutes. This includes incidents from Facebook account breaches to critical infrastructure attacks. When you take into account incidents that meet a severity rating needing to be dealt with in Australia, you're dealing with over 1000 incidents of reported incidents. There is a lack of transparency because not all incidents are reported to the ACSC. This gives better context of the actual attack environment. The work we did for a critical infrastructure asset around the likelihood of an attack in Australia, taking into account all factors, means each business has roughly a 5% chance of being attacked in one year.

Slide Change

Guest Speaker - Jordan (ForensicIT): The things you can do to reduce the cost of a data breach, as far as, I guess, investing in preventative cybersecurity and then trying to reduce the cost of an incident when it happens, the other types of things you're doing is the training and secure development, secure DevOps environment and training, incident response planning, and testing.

So we're doing that now, going through incident response, planning and testing. We're gonna discuss that now as well. So the types of things that you've got at the top are actually quite effective in reducing the cost. But there's something that is not a huge outlay for you to go through and implement. So the first stage of incident response when we start working with clients and the best one to start on is the preparation side of things. So start preparing to bring yourself up on this model.

So at the moment, a lot of organisations may actually be lacking in an inventory of what systems exist and what their subcomponents are to deliver the capabilities that they need to for their software or for their services. Once you've got a solid understanding of the base, start getting telemetry collected from it, look at detection. So you're starting to look at endpoint detection response products and other things, moving them into being able to triage incidents and actually respond to things.

So maybe you're looking at more of a SOC based approach up until actually starting to do threat hunting. Moving up into the top aspects of that pyramid. I don't think I've ever seen an organisation which you wouldn't expect me to, because if you're able to track and act on threats live in your environment, then I wouldn't expect you're gonna fall prey to a full-scale cyber incident and need the resources from Forensic IT to assist.

So, moving people up that I guess preparation stack is what the focus is for the initial aspect. Some of the things we'd advise you to do in that space is essentially profiling networks so that you understand what is expected in your environment. So the simple things you can do there around knowing, say, your Internet traffic inbound and outbound, because if a couple of jobs at the moment that would have been really useful, you would have seen a spike of 300 gig of data leave your environment overnight on a weekend. If that's not normal, then that needs to be triggered. So that's about that sort of situational awareness action that you take understanding normal. So if you're not aware that that is not normal, then it makes it much harder. And it makes it a much longer process to get through the initial identification phase of an incident, which will cover soon.

Slide Change

Guest Speaker - Jordan (ForensicIT): Enabling PowerShell. So if you're a Windows environment, making sure that anything that is run in PowerShell is logged is something that's incredibly important because there are lots of ways to try and use PowerShell and inbuilt Windows tools in an attack. And most threat actors use it because it's in every single Windows operating system. It lets you do things that most really high-scale attack tools that used to be built could do. Now you don't even need to bring them. You just use PowerShell Empire or various other things that you can use to run through and commit pretty much every single phase of the attack. Log retention. So knowing what you've got, how long it's kept for, and where it is and what questions it can answer for you is something that is really valuable to do upfront.

So something we work with organisations to do. I guess you can think of it similar to a tabletop exercise is that we will actually step through a scenario of a likely event and then go and collect the logs so you become almost like a functional exercise. We're collecting logs, we're doing analysis to see, can they prove or answer the questions that they would have? Like if someone logged into Office 365 from an unknown country that doesn't normally come in and then they browsed through your environment, can you answer questions around what they touched? There's a really good test for your environment to see whether your logs are adequate, whether they last long enough because in one job I finished Tuesday.

This week, it was the firewall logs that were the issue. The VPN was used to come into the organisation on the 26th of April, but the logs cycle every 24 hours, so they didn't know that until we asked. But given that's a perimeter, net device entry into the network is coming through it, I would want those kept for at least 90 days so that you've got time to see the actions on objectives and then work back. If you could get them to last for longer, that's even better. But that's more of a business decision than spending money on keeping those logs. Keeping host clocks synchronised sounds really simple. Causes us a lot of issues when we find out that there's time differences between different log sources. So keeping them synchronised, knowing what format they're in, whether they're recorded in UTC, they're in Australian time, local time is really good to know and maintaining a knowledge base ability of your systems in the event you do need to know how it worked and how the architecture of it was in its current state. That becomes really helpful, and it's something that incident response teams will ask you for.

Slide Change

Guest Speaker - Jordan (ForensicIT): These two are very Windows specific. So for Windows devices making sure that the volume shadow copies are enabled and that there's enough space there for them to keep operating because they will be critical in your recovery, especially from a ransomware incident or more lower level data integrity issues because you can use them to restore and also alerting on the use of tools that can interact with volume shadow copies. Because if you see someone deleting them, that is essentially an indicator that ransomware is about to be deployed because there's no legitimate reason for volume shadow copies to be deleted by an admin. And if that's not part of your change process, then it's immediately a risk because you're about to have a bad day, because that is normally the very last thing triggered by ransomware threat actors. The NTFS journal. So we're kind of getting very nerdy there, but there's a journal file on all NT file system volumes, which is like you think, think of your C drive or network drives in your computers. It's recording metadata about the files getting written to it, read from it, moved, copied. So evidence-wise, it is amazing because if threat actors come in and load a lot of stuff in your environment, and then they delete it all to clean up, they won't clean out the journal because it's a fairly unknown Windows artefact, which then gives you the ability to look back in time to see what they did, what folders they interacted with, which gives you the answers of what data did they take. So it's something that is valuable as far as preparation goes.

Slide Change

Guest Speaker - Jordan (ForensicIT): So, I guess getting into a bit of a case study. This was on an organisation who does development and hosts applications that work within Cryptocurrency. So they're global. And in September last year, we get a phone call on a Saturday night.

For some reason, every time I get phone calls, it's Friday or Saturday night, usually around 10 p.m., and there's an incident. So at that point in time, we knew that there'd been 36 unauthorised transactions within one of the hot wallets that supports their infrastructure and that those transactions had actually emptied the wallet. So there was a little bit of money left in it, but it was because the service crashed just prior to them emptying the wallet completely. So the initial starting point was that $62 million had been transferred out of their control by their own application. So it wasn't that they saw, I guess, malicious activity or malicious transactions, cuss coming from an external source. It was their own application that was transferring the money to the wrong location. So it was an interesting starting point and sort of what to look at and go to investigate. So at this point, the start of the incident response process. We're looking at identification and the scoping. So think of it that we want to identify what known, I guess, indicators of compromise we have, and then we want to know the full scope of it. So we want to know everywhere that this is impacting.

So in this initial job, it was AWS, it was their GitHub environment where we expected there was potentially access from, and Docker Hub, which is part of their development process in order to then create the containerised images to go into AWS. So the initial scope was those three solutions. The other part of that identification phase was we started conducting interviews with developers who have access to that environment and worked on the application required to commit this attack. So I guess that goes to your documentation, and your preparation is, do you have the names and details of everyone who worked on that during the time period it was being created? Because in this instance, we needed it quite quickly. So discussions with the development team identified that there'd been a social engineering campaign. So one of the guys had said he'd been contacted by a recruiter, and we believe that was potentially what was going to be the initial entry point, because when you did some research at that time, you actually came up with five other incidents that were identical to what we were dealing with. They were being done by the Lazarus Group, which is a North Korean government-backed organisation who have a habit of stealing Cryptocurrency.

Guest Speaker - Jordan (ForensicIT): At that point, I think in that September time period they stole about $380 million in Cryptocurrency. So it's a good earner. You can see why they would keep doing it, because without leaving your house, if that's the kind of money you can bring in, it's quite lucrative. So then the next part of that, once we've got that level of understanding, we'll move into the containment and intelligence development so we don't want them in there anymore. So the containment considerations are something that has to happen with a good team of people who understand the environment, how best to contain it, and the potential interruption to operations. In this job, there were already some containment measures done in that we identified through the discussions that the way that container was updated was through Docker Hub.

So we ceased all updates and then rolled it back to a previous version. That was about a week or two prior to the event. When we did that, we noticed that one of the so it wasn't the most updated version was set to production, so you can see there. There had been some additional work by the threat actors to complicate the recovery in that they'd modified multiple versions of the Docker Hub containers. So there was a bit of work there. As far as intelligence development goes to build up an intelligence picture to know where they've been feeds back into the scope side of things.

So these two act as a cycle, and it's the best way to make sure that you can move to eradication is constantly move between containment, build intelligence, know the accounts, they’ve got the IP addresses they're coming from and all of the other intel about them, and then feed it back into the identification phase to go. Where else have they been? Where else do I need to look in order to make sure that I know everything they've done and the full scope of the incident.

Slide Change:

Guest Speaker - Jordan Hunt (Forensic IT): So once we've gone through some more analysis, we realised that the part of the recruitment process that the developer went through was that he was contacted first through LinkedIn and then subsequently moved off to Telegram. So it’s very common they don't want it in LinkedIn because that's traceable. We can get information from LinkedIn themselves through information requests, whereas once you're on the Telegram you're on an encrypted channel, we're not gonna be able to see it. So then they asked him to download, work and do some additional work and challenges and tests within a different GitHub repo.

The problem was, when he undertook the tests, it was actually collecting legitimate credentials and everything else from his Chrome browser and everywhere else that he'd saved. And then during discussions he said he had actually been using, he was working from home on the weekend this occurred and he'd been using his personal laptop to log in and do the development work that he would have for the organisation. Raises some questions around bring your own device policies and things like that and whether that's allowed in your organisation. But in this instance, it caused some complications because we then didn't have the power to compel him to give us the laptop to announce to do the forensic analysis. He cooperated and was really good about it, but that could have gone the other way quite easily. Once they were there, they harvested what turned out to be AWS keys with admin rights. And they used them to go through and do the remainder of the attack. So they were in the Docker Hub and they were in AWS. Again, this comes back to your logging. It was quite interesting to see where they were coming in from. The IP addresses were all AWS or Microsoft.

So that was an issue because you're looking at, well, which ones were him because it was very similar to their own working. But once we realised you could tell from the user agent. So the logins to GitHub, Docker Hub and AWS all record the user agent, which is essentially a string telling you the type of device and type of web browser that they're using. So we found that all of the malicious activity had the same user agent. So that was a way for us to quickly try and identify what was the additional attack activity we had to consider. And that again, is that containment and intelligence development moving back to the identification phase. Because as soon as we knew of the user agents, we can then say, OK, now we need to find everywhere they've been, which gives them the full picture, which means you can move to the next phase of instant response, which is eradication and recovery.

Slide Change:

Guest Speaker - Jordan Hunt (Forensic IT): So in this instance, eradication became we burnt all of those Docker Hub images that were interacted with them. They were all deleted. His accounts were deleted. There were new ones created for him to use. The AWS keys and everything were cycled so that we knew that there's nothing left that they're gonna be able to use to come back into the environment.

There's also that concept there that, well, I guess, would you come back if you had already successfully taken $62 million? Is there really a reason for you to come back and try and get more? Turned out, yes. They did start approaching another staff member using a similar social engineering method not long after, which we then impersonated that staff member and tried to get ourselves groomed and hacked. But it turned out to be a different approach. If you were with Intel's kind of suggested that if there were, if you were approached by Jane Wang, it was through LinkedIn and it would be like a recruitment approach. Whereas then there was an Anna Wang, which that was more of a romance scam social engineering type approach, which we actually found other victims online for Anna Wang scams talking about. One guy from the US had given her $800,000 to invest in her Cryptocurrency organisation. And that was sort of the play is that she would be, she'll be your friend and she's a Cryptocurrency investor so that you can give her your money and move through that.

So the Lazarus Group have quite a few different methods for trying to get that type of information. So then these guys moved into more long-term recovery, so that became a lot of the remaining parts of these slides, which we'll go through quickly. But I've deliberately sort of put it in there that you guys can consider this after this event and sort of pay more attention to them in detail because the lessons learned in post-incident activity that we did during those jobs and every other investigation that we've done very much moves to having a log management and retention policy, whether that's for your internal businesses or for software you're delivering to clients and knowing so the jurisdiction they are in will then dictate how long you keep those records for or the regulator. Regulators will also say how long you need them for. Then the types of things that you should keep for an operating system.

These are the critical things that we would want to see if for every operating system you had and that was impacted in that, say, we use that scenario to be Linux operating systems in an AWS environment, we wanna know how the logs can tell us. Use of special privileges, elevated privileges, root accounts. What are the systems? Start-up and shut-down time so that we can see roughly what they've been doing. If we were to move then to a web application, which this was, we'd wanna know authenticated users. So anyone who's been pushed out so failed authentication attempts, any crashes or error messages.

Guest Speaker - Jordan Hunt (Forensic IT): So if you see something crash constantly, it can be an indicator that something has been changed. So it's a good way of doing it. IP address and port. So the port is an interesting one. The reason we put that there is because of an actual job where the connecting port was a Telstra or Optus mobile phone. So you get the idea of network address translation where there's thousands of people coming from that one IP address, you have to know the port number in order to then attribute it to an individual. So it's not something that's ever kept. In my time with law enforcement, it wasn't something we ever had. I think maybe once we were given that level of detail from a breach, but it is really important. If you ever want to be able to attribute attack activity to someone as the, I guess you're developing the tool. You have the option to put it in there.

Things for databases, and this is something that does cause issues quite a bit in that data expelled from within a database can happen in a lot of different ways, and in some instances, not having the execution of commands and the execution of queries causes issues.

So if you've got the ability to turn on this level of logging within the databases that support your infrastructure, so sort of coming to a conclusion there that the essentially from the kill chain side of things, the things that you can actually have an impact on are the ones where you can stop delivery. So educating people around the risks of social engineering and phishing emails and knowing how well it can be done and then looking at the exploitation side of things.

So keeping things patched and up to date back to that IT security side of things is that your actions that you take if you receive a phishing approach or a social engineering, even if you have actually done it, being open and honest about it is the best way because you are actually gonna end up in a good position, being able to deal with it. Whereas not telling someone for a week and then bringing it up once the attack has occurred is a far worse position to be in.

So they're the parts that you can actually change the context in which the offenders operate, so you will have a good chance of moving through and recovering from an incident quickly. I will stop sharing now and we can move through to the question side of things.

Slide Change

Host - Charles Gillman (DSPANZ): Thank you, Jordan. That was really insightful. If you wanted to start putting your questions into the Q&A, we can present them to Jordan. I've got a couple already for you.

Q&A Section

Host - Charles Gillman (DSPANZ): *What does a typical timeline of a cyber attack look like from the initial access to your involvement? Are we talking hours, days, weeks?*

Guest Speaker - Jordan Hunt (ForensicIT): All of the above. The attack from the time the attackers get in there normally goes weeks to, if not, a month prior to when the client notices.

The very large one we're finishing at the moment saw the attack activity on the 11th of May, but we traced the initial actions of the threat actor back to the 22nd of April. They'd been there well before. The initial attack activity, when you look at it in the context of being ransomed, starts to make sense to the organisation because even they brought up that, "Oh, we should have acted then." Yes, but you would have also then have to have been looking. So, it depends on the level of visibility you've got in the environment, but the notification to us and the actions you take relate back to that cost of a data breach.

The faster you take action and kick them out has a direct correlation with the cost of the breach as it goes forward. You want them out of the environment as quickly as possible. Normally, we're talking from the day they get access to them taking actions within a month, but ransomware threat actors are skewing that quite significantly because they will come in very quickly, do their work, and then encrypt the environment because they want to get money and have an impact as quickly as possible.

Host - Charles Gillman (DSPANZ): Great. *Can you highlight the top three vulnerabilities that attackers are using to get into the environment and how we could protect ourselves? For example, is something like MFA a good control?*

Guest Speaker - Jordan Hunt (ForensicIT): At the moment, the main ones we've seen are insecure remote access processes and poor passwords. They all come in hand in hand with password management or hygiene. One of them was using a generic account that was used for tender submission for the organisation, because it was a group and a shared account, they decided they couldn't have MFA on it because multiple people needed access to it. I would have then said, "Don't make it have interactive login. Make it a shared mailbox," because you're reducing that risk. Unfortunately for them, the password was also the name of the company with a one at the end. The password policy and complexity is something that we still see as quite an issue. If you can brute force it, it's a problem. There's been quite a few recently, like a username called Temp with a rubbish password that is part of the VPN coming into the environment. The main things we see as far as the initial entry point go are being able to access and brute force remote access methods and then poor password management.

The MFA side of things is interesting because we've been telling people for years to do it because it stops this type of attack. Then someone built Evilginx, a pen-testing tool used for phishing emails and things like that, but it will bypass MFA because it collects the authentication token. It's not that they've broken MFA, they're exploiting the process where it creates a token that is authorised to impersonate you. The process of you clicking on the phishing email establishes a link to Microsoft through their Evilginx server. You enter your password, it collects it. The MFA allow it collects the authorisation token on the way back. So, Microsoft's 90 days that token will last. Essentially, they have a way of impersonating you as an MFA authenticated user for the next 90 days without any issue. That has caused an issue recently. The concept of creating authorised tokens is the issue, not necessarily MFA. Having MFA is still a good idea, but it means you have more of an onus on education for users to not click on phishing links than you did prior. MFA was the saviour, whereas now it's not.

Host - Charles Gillman (DSPANZ): Not withstanding that new technique, it sounds like they're not getting the basics right. We're getting some questions coming through and we've got one here.

"Do you have experience of attacks occurring in OT environments? If so, what differences have you observed?"

For those who don't know, OT is an operational technology.

Guest Speaker - Jordan Hunt (ForensicIT): We do a lot of work with critical infrastructure and operational technology in the transport, gas, electricity, and water sectors. The main difference in the response approach is the consideration of the impacts on operations. This outweighs a lot of the containment strategies. You can't just disconnect the network because it may cause more of an event than the cyber incident was. We spend a lot more time doing preparation with them. I'm actually doing it tomorrow with a client because getting an understanding of the incident response process in the operational technology side of their plant has already been done. Now, we're working into the impacts that can potentially be seen from the IT network. We're looking at the fact that their ordering and delivery systems are in the IT environment. If you disrupt those, you have to stop production because you no longer know where to deliver the product and when. That overlap is the other thing we're considering a lot with the OT space.

The other more technical differences are the need for more network forensics in OT environments because there's usually a lot more smaller, lower-spec equipment and PLCs and things like that. You're not going to be able to collect really detailed logs and forensic artefacts from them. Identifying choke points for network communications and then figuring out what network telemetry info you can get to understand what malicious activity is occurring on the network would be the primary differences. The preparation in an OT network is very different from what we would normally do for a corporate environment.

Host - Charles Gillman (DSPANZ): Great. Thank you.

"Have you seen any incidents where AI has been used by the threat actors to facilitate the attack?" And I'll add to that, *"How did you know it was AI facilitation?"*

Guest Speaker - Jordan Hunt (ForensicIT): Not really. Brute forcing passwords and things can be done through AI-driven tools. There are now pen-testing frameworks where you can test a network and find every vulnerability in it through AI. We haven't seen it showing up well from our perspective because we're looking at it from the evidence point of view. If it's not leaving logs behind that show us where it came from, it becomes less noticeable. The one strange one we've had recently was a developer testing code in ChatGPT and similar AI tools, but he was uploading code with keys in it and all sorts of stuff.

So, you're getting this thing where it's now learning from your application and your secret keys and all the passwords and credentials. If you want to, you can now go to ChatGPT and say, "Show me the codes and AWS keys for this organisation," and it will come back and give it to you because it learned them. It's not really used for the attack, but it would be used as an initial entry point. How that plays out in the next few years will be interesting to watch.

Host - Charles Gillman (DSPANZ): I think it's already playing out. I saw a report that since the launch of ChatGPT, global phishing volumes went up by 1200%. So, they're certainly being used to craft very realistic-looking emails.

Another question here: *"The study case was intentionally targeting the crypto company. Can you give an example of a non-target attack? Some less appealing company getting attacked."* As lot of companies say, "Oh, we're not going to get hacked, we've got nothing interesting."

Guest Speaker - Jordan Hunt (ForensicIT): We get that a lot. "Why did they target us?" They didn't. They found you because you had some insecure application or open remote access point.

One we just finished was a small Australian company. They had IT people not following good hygiene. They had open RDP directly to the internet, not even a firewall. Someone found that, came straight through, and encrypted the place. They're quite mean. They deleted the backups, reinitialised the NAS storage device, reset all the admin passwords, encrypted the data, then the virtual machine disks as a second layer of encryption, and deleted the passwords. Then they saw a remote desktop icon and used it to go to another organisation and encrypt them as well. That was all through insecure administrative practices. It's not because they were targeted specifically. You could argue that Australia is a good target because we're relatively wealthy, rely heavily on IT infrastructure, and some organisations have a relaxed attitude toward security. The primary reason for those attacks is insecure administrative practices.

We've done jobs for very large financial software services, multinational chemical organisations, and down to a one-person optometrist in regional Victoria. The threat actors don't care. They asked that one-man optometrist for three and a half million in ransom. His business turnover was struggling to get 50k, so clearly, he's not going to pay 3 million in ransom. They don't know who they're attacking until they've already done it. Hopefully, that clarifies things.

Host - Charles Gillman (DSPANZ): Great. Thank you.

Are there some tips that you can give us at the employee level to prevent some of these types of attacks? How can people look out for telltale signs of phishing, for example?

Guest Speaker - Jordan Hunt (ForensicIT): It's hard to say, "Don't click on any links," because there are reasons for you to do it. If we share reports with you at the end of an incident, we'll likely use Citrix ShareFile and say, "Click on this and put your email address in it." If you're expecting it from the person and know it's from their legitimate email address, that's fine. There are technical things you can put in around that, like Safe Links and Defender and other email security platforms. The social engineering side is harder because there's a lot in that that's quite nuanced. A reasonably young bloke at his first job as a developer, being approached by a company called Orby, another cryptocurrency firm. It's relevant to his organisation; he's got development skills. A government-backed hacking group has a lot of resources to impersonate and manipulate him well. In that scenario, education is key. Be sceptical of things if you see them offering you work and asking you to move to an uncontrolled environment, like from LinkedIn to Telegram. They don't want it on LinkedIn because that's traceable. Moving to WhatsApp or Telegram is a red flag. Understanding that a legitimate recruitment company wouldn't take you off to an encrypted messaging channel. Education around these signs is critical to prevent phishing and social engineering attacks. The amount of things now that will bypass security measures is amazing. We've had six jobs from one phishing email around November to December last year about Christmas bonuses with a QR code. Everyone's expecting a bonus due to the cost of living crisis and Christmas. Sixteen people clicked on it and chatted among themselves about it. There's a bit of organisational education needed to address this.

Host - Charles Gillman (DSPANZ): Great. Thank you. Those poor people and their Christmas bonuses.

Another question from the audience: *"If we're using online messaging and meeting platforms like WhatsApp and Zoom, is it a good practice to close all other applications such as Gmail and Internet banking? Does that reduce the risk or make no difference at all?"*

Guest Speaker - Jordan Hunt (ForensicIT): Not a huge amount from the threat actor's point of view. Not saving passwords in your browser is a good start. Use a protected password manager like 1Password or Bitwarden for additional security. There are many known information stealers that grab cached information from browsers. Keeping your device up to date and not using very old devices is also important. Privacy risks around having applications open come more from WhatsApp and Google recording everything on the device. It's more about privacy and invasion of privacy that you agree to in their terms and conditions. Apple's and Microsoft's new AI features also record a lot of activity. The use of apps on mobile phones is generally safer because it's harder to get malware on mobile devices.

Host - Charles Gillman (DSPANZ): Great. Thank you for answering those questions. It was really insightful and slightly depressing.

Guest Speaker - Jordan Hunt (ForensicIT): Sorry to leave you all on a sad note. Thank you for having me.

Host - Belinda Stewart (DSPANZ): Thank you so much, Jordan, for such an insightful presentation this morning. I think we have lots of takeaways and big thank you to Charles for facilitating our Q&A session. And to everybody else who has jumped on today, thank you so much for your time. We hope you have taken some good learnings away.

Host - Belinda Stewart (DSPANZ): We will make sure that this presentation and the slides are provided in the coming weeks. We will also send out a survey to everyone. It would be great if you could fill that out and let us know the types of things you'd like us to do in the future and how you found the presentation.

Thank you so much, everyone. Have a wonderful afternoon. Thank you.