Cybersecurity Incident Response in Large Enterprises

Taz Wake | Halkyn Consulting Ltd | Originally published 2024-01-22

Background

Responding to cybersecurity incidents in a large enterprise setting is often challenging, even for experienced responders. We often focus on smaller environments because these are easier to explain and easier to learn about. The problem is that scaling up to global enterprises with more than 200,000 devices, multiple cloud environments and a wide range of operating systems is not a simple matter of "**it is just bigger**."

Yes. It is *bigger*, but that is a tiny part of the picture.

In practice, dealing with incidents in very large enterprises requires a change in how we do incident response.

The good news is that it is very much an evolutionary change - we do our normal IR work and just add in some adaptations for the new environment. The bad news is that if we don't make these changes, and we just run our normal IR processes then things are much more likely to end badly.

Ultimately, we can look at "enterprise" DFIR as being a *step up* from traditional IR. The enterprise responder needs to build on the knowledge, skills and capabilities that underpin traditional IR.

In this article, I've tried to summarise some of the key areas where new thinking is needed. This is not an exhaustive list, it is just the main areas where I see organisations struggling as they evolve from dealing with intrusions in 1000s of devices to 100,000s of devices.

Enterprise DFIR Tactics and Requirements

Enterprise, and all large-scale IR, shares a lot of things with smaller-scale IR. In general, we do the same things, we examine evidence, we form conclusions and we try to unravel the adversary's hidden activities.

However, three areas hold special importance.

Active Defence

For the absolute avoidance of any doubt, the concept of "hacking back" should be avoided like the plague. Please don't think this makes sense.

However, Active Defence, in this concept is about forcing the attacker to slow down and be noisier. When trying to find an intrusion somewhere in 200,000 endpoints (and maybe another 20,000 servers and 200,000 cloud systems), it is critical that we find ways to force the intruders to give themselves away.

Active defence is essential for enterprise IR.

With active defence, we are only really limited by our imagination. It is important that our senior incident responders understand attacker behaviour and then we can turn this to our advantage.

Some examples of things that you can deploy include (but are not limited to):

- Honeypots
- Honey tokens
- Zip bombs
- Traffic shaping

All of these are excellent ways to force the attacker to fumble and give your incident response team some hope of catching up with the adversary.

Enterprise Visibility

You can only protect the assets you know about. This simple truth is often overlooked. You need to push for the maximum possible visibility into every endpoint because even 99% coverage becomes a problem when you have 400,000 systems. In that example, leaving 4000 devices without a full security stack gives you an incredibly large attack surface, bigger than the whole asset list for some medium-sized enterprises.

We should also look at visibility in two ways.

- Can we see every device in our organisation? This can be thought of as *horizontal* visibility.
- Can we see everything that happens on a given device in our organisation? This is *vertical* visibility.

Both types will directly influence your ability to do incident response at scale.

Enterprise Documentation

The bigger the environment, the more important it is that things are written down. In a very small business, say with 10 people, it is likely that everyone knows what every device does, what is important, what isn't etc.

That really doesn't scale. It isn't even close.

As we look at bigger and bigger environments, it becomes crucial that accurate records are kept **and** made available during an incident. Remember, we can't pick the time an intrusion will take place and we normally only detect them at really inconvenient times. You need to plan for this.

Good documentation will help you. Consider the following:

- Can an incident responder identify the role, criticality and business owners of a server at 0200hrs?
- Are there current, and accurate, data-flow diagrams showing how assets should communicate with each other?
- Can you quickly identify what an application should be doing, how it is configured and where it writes its event logs at 0200hrs?
- When was the last time a given device was patched, what patches were applied, what ones were held back and why?

These are just a tiny example of the questions an incident responder might need to answer.

Managing People during IR

While technology matters, incident response is really *people* (the responders) dealing with other *people* (stakeholders) and trying to prevent a different group of *people* (the attackers) from causing more problems.

IR is absolutely a people thing. Managing people is important.

People management matters

The incident response team

Burnout is real. Responders will often work without breaks. None of this is good. It is important that anyone running, or leading, an incident makes sure that the people involved are able to keep going.

Some elements are trivial, such as making sure they have food and drink. Some elements are more challenging, such as ensuring people dont work too long, have time off and get time to recover between incidents.

Security Stakeholders

Dealing with CISOs, Security Directors and the like requires time and effort. Security leaders will have valid questions and a genuine need to get answers, but responding to these will take time away from the incident response.

This is a difficult balancing act that has to be managed. Managing it takes effort and skill. Never assume it will just work.

Ideally, the incident response team will have a lead who can field all the requests and questions, to allow the responders the time and breathing room to run the response.

IT teams

Most cybersecurity incidents will involve IT teams. You have to spend time and effort managing the relationships here.

During an incident, lots of IT staff might feel at fault (even if they aren't) and might even feel out of their depth.

A good enterprise incident responder can identify this and take measures to minimise any impact it might have on the response work.

Legal, Compliance and Comms

Here it gets complex. During an incident, we have to be able to work with these departments and their input into our process will be essential. However, we frequently face a language barrier in that the terms we use will make no sense to them, and their terminology will make no sense to us.

Spending time understanding the needs of these teams, and setting expectations will make your IR process much smoother and reduce confusion, distrust and even mistakes.

External stakeholders

- Law Enforcement. Although cyber intrusions are often a low priority for law enforcement (certainly in the private sector), they can do things most IR teams cant. Examples include domain takedowns or working with financial services organisations to see about refunding fake payments.
- The public, customers and media. Here it gets challenging. As a general rule of thumb, your responders shouldn't be directly communicating with the public (we are often a bit more negative than we need to be). A much better approach is to work with your comms team, with your legal team and security stakeholders to build an effective public communications strategy. Try to avoid screaming "*A nation-state attacked us*" if it is going to turn out to be a 15-year-old launching SQLmap from their bedroom.

Scale and Time

If Spiderman had been written about incident responders, Peter Parker's uncle would have said "*With great size comes great complexity*", because that is definitely something that can catch responders off guard.

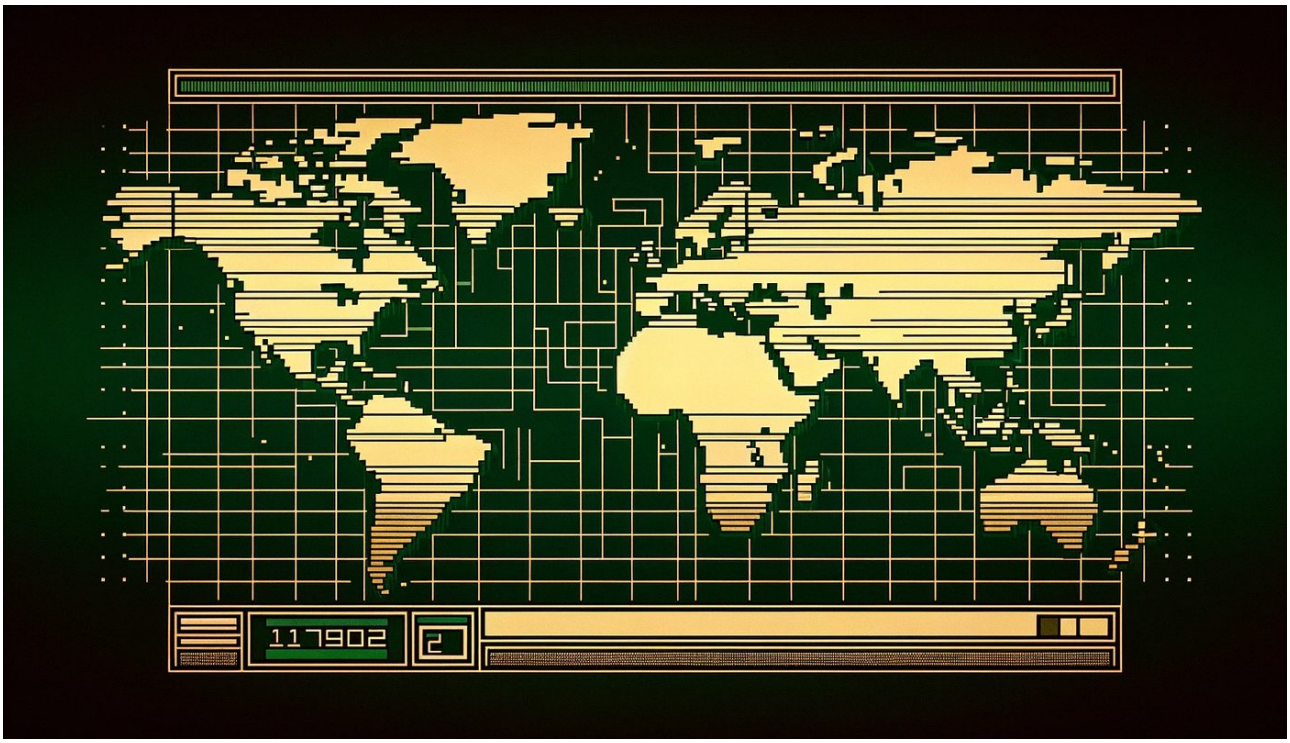

Global networks create complex problems.

With a global network, an intrusion can quickly pivot across continents and timezones before the incident responders can even realise the attack has started. This is often a situation that takes away the "home field" advantage we have in IR, as the attackers are less likely to care about timezones, user behaviour, laws etc. However, we have to care.

As a result of this, the scale of an organisation creates some new challenges to consider.

Geography and Networks

This matters in a couple of ways. First, the geography means we need to consider where our responders are located. If we want physical "hands and eyes" on a device, we need to make sure we have the right people in the right place. Thirty years ago it might have been an option to fly an incident responder halfway around the world to capture evidence, but it is vanishingly rare today.

Instead, we rely much more on remote connections or having local teams who can help out. Remote connections are always at the mercy of network speeds, and if you are responding to an intrusion in a rural location where the network gives 2MBPS uploads, getting a disk image is simply *never* going to happen.

In those situations, we tend to fall back on having local support teams. This makes it critical for our Incident Responders to provide detailed instructions or training. If we get this wrong, the evidence collection will fail and that can derail an entire incident.

Geography also presents a legal problem. When we collect data in IR, there is always a significant risk of collateral intrusion (where we collect unrelated data, such as the user's legitimate web browsing history). This creates a new question - is it legal to take the data from the user's device in Country A for analysis in Country B?

As IR professionals we absolutely shouldn't try to guess the answer here. Speak to your legal team and make sure you have the correct approvals.

Timezones

A related problem is caused by timezones. Again, this can trip us up in a couple of ways.

Time is not on MY side.

The most annoying is that some Operating Systems (*cough, Linux, cough*) and some event sources (*cough, most security products, cough*) try to helpfully capture data in local time. This means when we look at an evidence item, we need to work through some occasionally complex mathematics to work out when the event happened in relation to the timestamps on other systems or even on other evidence items on the same system. Remember, in DFIR, the only acceptable timezone is UTC.

Another issue arises when the responders are in a very different timezone from the victims. If you have an IR team in Seattle USA and they are dealing with an intrusion in Papua New Guinea, there is a 19-hour time difference. This can cause significant issues and delays so you need to make sure solutions are built into your incident response plan.

And no. It is not as simple as saying someone has to suck it up and work through the night/weekend etc. While that might feel like the right answer in the heat of an incident, you need to consider contractual obligations, working time directives etc. *Try to avoid thinking that declaring an incident means you get to do whatever you want, whenever you want*.

Technology

The last challenge I want to talk about here, is the problem with technology.

Now I get that it sounds strange to say Technology is a *problem* for IR, because we kind of *need it to do IR*, but it is all about **complexity**.

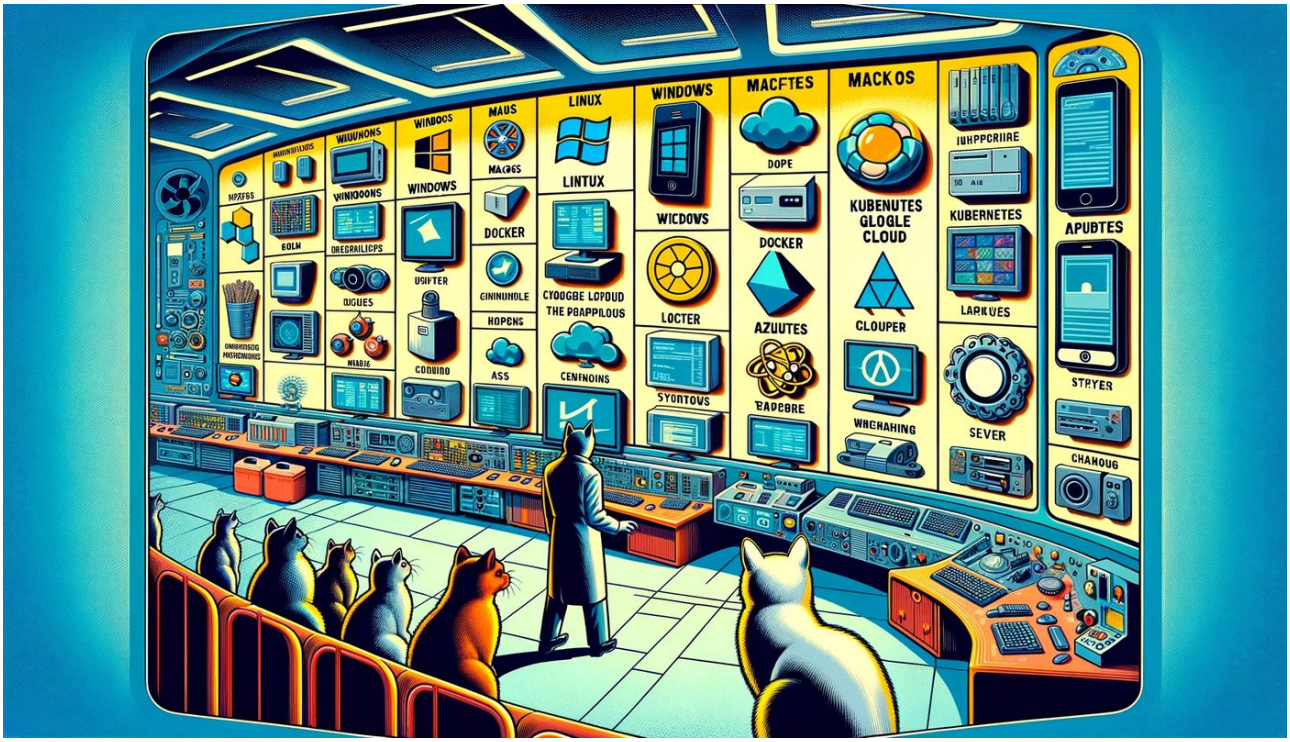

Cats really aren't very good at spelling.

In a small business, we can be pretty confident that we can identify every user's device. It is likely that everyone will have the same type of device. They are all probably managed by one or two people.

In a global enterprise, every bit of that goes away.

We need to deal with intrusions where an attacker has compromised an Ubuntu web server in the DMZ, pivoted to a RHEL database server (exfiltrated data via TFTP), and then compromised a Microsoft Exchange Server. From there they stole a set of credentials and used RDP to go through the DMZ firewall into the corporate network. Once out of the DMZ, they pivot mercilessly through our enterprise compromising the long-forgotten Windows 2003 Server we have running a "business critical application", the obscure NetBSD devices a few devs left running in 2010 and everyone forgot about, several hundred Windows machines ranging from Win 7/Server 2012 through to Windows 11/Server 2019. After they've caused carnage here, they find a set of credentials to a developer's MacBook and the attack moves in a new direction. They take advantage of sloppy sharing getting into multiple macOS devices and even the company GitHub. As icing on the cake, in here the attacker finds credentials for the cloud platforms and now the AWS and Azure environments are compromised.

Does that sound far-fetched? Sadly, it is the reality of lots of intrusions today. It also means that the poor unfortunate incident responders have to be able to understand evidence from dozens of different platforms.

See what I mean about complexity? This hasn't even considered what happens if the intrusion gets into an ICS/OT or IOT platform. Do you have an IR plan for what happens if your HVAC system gets compromised? Do you even know what the evidence would look like?

To be effective in enterprise-scale DFIR, you need to have at least a basic understanding of the forensic requirements of most platforms. It is not enough to say "I use Linux every day" or "I use macOS every day", we need to understand it from an intrusion analysis perspective. Assuming daily use gives us this detailed view is often mistaken.

Would you like to know more?

Would you like to know more?

Clearly, there are limits to what we can cover in a LinkedIn article. There is no way I could even start to provide solutions to all the issues here and still have it readable. The best approach to this is to ensure that your incident responders are well-trained, well-motivated and well-looked after. They are an absolutely essential resource.

If you want to spend six days learning about ways you can become a better, more efficient, more skilled enterprise incident responder then have a look at the SANS Institute course FOR608 "Enterprise-Class Incident Response & Threat Hunting", you can find out more about it at https://sans.org/for608.

In this course, you will practice responding to a complex intrusion, where the highly skilled threat actor has used a range of modern intrusion techniques to compromise a well-defended organisation. You will deal with evidence from varied platforms and the constant issue of having to work through large volumes of often confusing data to quickly find answers for your senior executives.

About Halkyn Consulting

Halkyn Consulting Ltd is a UK-based security and risk management consultancy providing 24/7 incident response, digital forensics, threat hunting, virtual security leadership (vCISO) and security compliance support to organisations worldwide. We deliver clear, practical, vendor-independent advice based on real-world threats — no hype, no upsell.

More free security resources are available at www.halkynconsulting.co.uk/security-resources.

Web: www.halkynconsulting.co.uk | **Email:** info@halkynconsulting.co.uk | **Tel:** +44 1244 940858