The GDPR requires us to detect and respond to personal data breaches – you could be fined up to 20 million Euros if you don't have a process for that – but those processes must themselves be lawful. How does that work?

Here we'll look at three security activities: logging, scanning and responding.

Measures to **prevent** breaches are important, but they can never be perfect.

Measures to detect, investigate, mitigate and learn are also essential. One of the key tools in doing that is logfiles, which let us spot unusual patterns of activity and, once we have spotted something, work back to determine the cause and any consequences.

To be useful, we need logs at many different levels – for example network flows; email traffic (but not usually content); authentication successes and failures; web, database and application requests. That's a lot of data, much of it covered by GDPR.

But keeping it is permitted, so long as:

it contributes to detection and response;

there's no less intrusive way to do that;

and if the risk if we keep it is less than the risk if we don't and so cannot deal with breaches.

So we need to ensure we only keep the data our detection and investigation processes can actually use,

we keep it secure,

and we dispose of it when there is no longer significant benefit from any investigation. For example because other changes will have made investigation impossible, or because the intruders have had ample time to do damage and conceal their traces.

It's much better for a trusted security team to have access to **some** personal data, under strict controls, than for an intruder to have access to **all** of it, under none.

Security teams will often carry out proactive checks for insecure machines, to discover them before they become a threat. This usually involves network scanning, something that is regulated by national, rather than European, law.

Pretty much all countries have laws against "unauthorised access", but how they determine whether scanning breaks those laws is very different: Some ask whether the scan constitutes "access" at all; Others whether the activity was "authorised", either explicitly by the system's owner, or implicitly by connecting it to the Internet; Others ask what the scanner's motivation was.

The UK's Computer Misuse Act has one of the broadest definitions of illegal conduct: there's no "motivation" test, and almost anything constitutes "access".

So Authorisation is key. It's best to have explicit authorisation, from a contract or network policy. For example the Janet Security Policy authorises Janet CSIRT to perform limited scans where there is a threat to the network. Relying on implicit authorisation – for example to decide if external systems pose a threat and should be blocked – is legally tricky.

The few reported cases suggest that scans that only use legitimate features of protocols – such as banner grabbing – may be OK, but actually testing overflows probably isn't. Avoid scans that risk crashing the target machines.

So, what to do when you find a problem?

When investigating a breach you should – both from an operational and legal perspective – try to focus on malicious traffic and look at as little legitimate traffic as possible.

Automation is a great help: try to use computers and scripts first to narrow down what human investigators need to look at.

Even at the human investigation stage, try to stick with machine identifiers such as IP addresses for as long as possible;

And only link those to actual humans when you are confident that they are victims in need of help.

Note that identifying attackers should be a legally and technically separate process.

Security teams will often want to share what they find with others: to warn them of compromised machines in their networks, or to tell others about new attacker techniques.

Here the balance between risk and benefit needs particular care: if you can send a simple notification direct to the individual victim or their organisation, that minimises the risk;

if your announcement needs to be wider in content or distribution, then check the benefits justify that increased risk. Sharing should only include information that the recipient **needs**: that may include information about **their** users, but rarely about anybody else's

A good test is whether your sharing is more likely to help attackers or defenders: good incident response practice is usually good data protection practice, too.