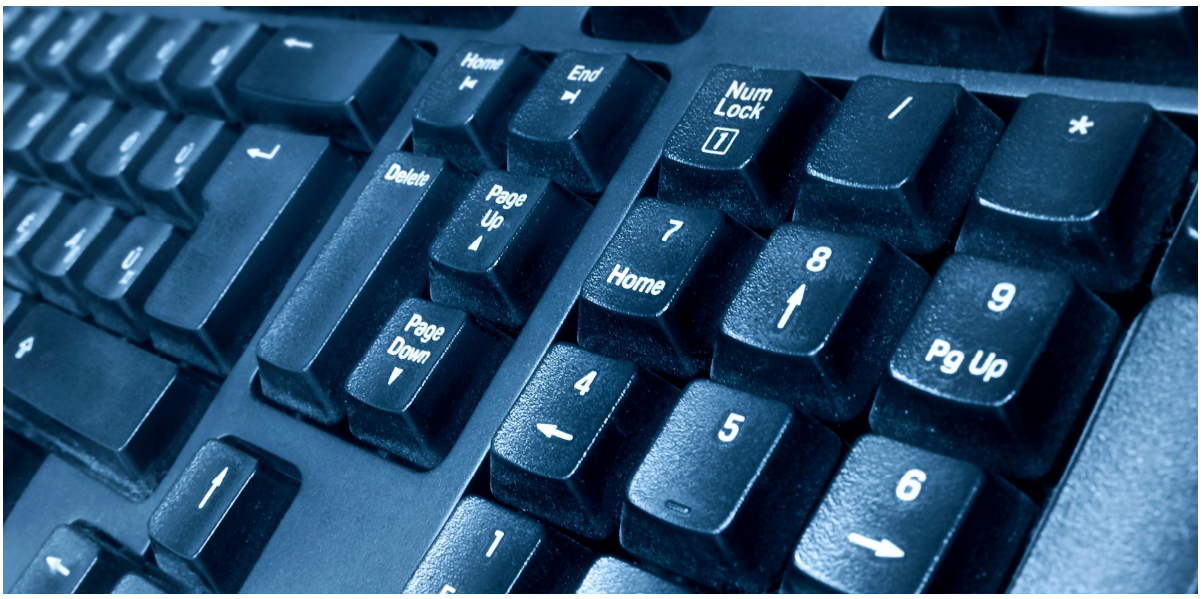

Who is in control of your infrastructure today?

Jamie Bodley-Scott, Account Director – Systems Integrators at Cryptzone explores how the threat landscape is evolving and the resulting need for organisations to review their approach to virus protection.

On the surface nothing much seems to have changed over the last 10 years in relation to virus threats. New vulnerabilities appear regularly, patching has become the norm and anti-virus software is doing its best to keep our networked world working. However take a closer look and there has been a massive evolution.

Step back in time

30 years ago IBM launched the XT5160 – the first hard drive DOS based PC. And in less than 3 decades the 'personal computer' has evolved to form the backbone of the networked world that we all rely on.

The computer virus, nowadays, so seemingly tied to the PC, was in actual existence almost a decade earlier. It took until 1986 for these two threads to come together and the first PC virus 'Brain' was born. By 2000, networks were spreading and so were worms like ILOVEYOU which was considered one of the most damaging.

So where are we today? To a large extent nothing has changed, but the rise of targeted attacks and the involvement of nation states, often linked because of viruses, such as Stuxnet and DuQu, point to separate new and worrying developments.

Targeted attacks can be engineered to seek out a very specific machine, infrastructure or geography. They could be used to target one company, maybe with the intention of stealing trade secrets or discrediting that company. If you want a good example then just look at the infection map for Flame, it is tightly grouped around the Gulf States. The other development is the apparent involvement of the nation state. Now whether this is true or not, the key learning point is the scale of the enterprise behind these attacks. In under 20 years the resources that can be deployed has grown from 'two brothers' to 'nation state' with the ability to shut down critical systems.

Today most organisations still use anti-virus that relies on a snapshot of the signatures of the bad stuff to be kept out. The major disadvantage of this approach is that it does not know what it does not know about! So when a new or adapted threat, not yet in the snapshot, appears then it will be allowed to run. This is the classic mechanism by which some of these targeted attacks have managed to be so successful.

So what of the future? There is some light at the end of the tunnel in the form of Application Whitelisting. This technique has two parts. Firstly a snapshot of the computer is made which will contain signatures for all the programmes, operating system elements, drivers, etc that were originally installed. Second an agent is installed which checks everything just before it runs to make sure it was in the snapshot. Even though this technique still uses signatures, it has the major advantage of being able to block unknown code and prevent what is now known as 'zero day threats'.

Application Whitelisting

So why do we still stick with less effective anti-virus solutions when application whitelisting software now exists? Both techniques use signatures (in part) and signatures need to be generated and managed – so what is the issue?

Back in 1986 when the first PC virus came along there was just that; 'a PC virus'. By then PCs had been around for a few years and programmes already contained thousands of executables. And each PC probably contained different executables because one was in engineering and another in finance. So was it easiest to look for the one piece of static bad stuff that was the same everywhere or a variable amount of good stuff which was different everywhere? I think the answer is obvious!

Two things have changed in the last couple of decades that mean it is now time to reconsider the options. Firstly, the numbers game; the amount of bad stuff grows daily and some anti-virus signature files contain in the region of 20 million signatures. On the other hand the good stuff has not grown as fast and a signature file for a standard operating system, such as Windows XP Professional, will contain something like 50 thousand signatures. Secondly, the rate of change; viruses were a static programme and did not change – but nowadays they are written to self-adapt or operate in a command-control mode where they can be remotely updated. So what do you do now? Look for the 50 thousand relatively static signatures of the good stuff or the growing 20 million adapting signatures of the bad stuff?

Signature Management

By solving the challenge of signature management, we solve the problem of why application whitelisting is not as widely adopted as logic would suggest it should be. Most organisations hope they never see any bad stuff and have no expertise in the dark science of understanding them. So it is sensible that both the generation and updating of anti-virus signatures be 'outsourced' to the experts, and that is how the industry has developed. Application Whitelisting appears to require the opposite approach. The organisation itself has to both generate and update the signatures of the good stuff, because PCs are unique to every organisation. Worse still is the worry that with application whitelisting the signature file of every PC might have to be different! Compare this with anti-virus where the same signature file can be applied to almost every machine. Well worry not as a step change is on the horizon. The 'outsourcing' of signatures is possible with application whitelisting as well, simply by taking advantage of the increasing amounts of signed software available nowadays.

Today the concept of 'signing' software is becoming commonplace and will contain metadata, such as the software author, a checksum to verify that the object has not been altered and versioning information. Signing involves a process using a pair of keys, similar to SSL or SSH sessions. The private key used to sign the code is unique to a developer or organisation. These keys can be self generated or obtained from a trusted certificate authority (CA). When the public key used to authenticate the code signature can be traced back to a trusted root authority CA using secure public key infrastructure (PKI), then you know that the code is genuine. We see this most commonly today in environments where we have not requested a given piece of code, so the source may not be immediately evident - for example a Java Web Start application accessed from your browser.

This same signing process can now be used by application whitelisting solutions, such as Cryptzone's SE46. The agent which checks everything just before it runs clearly trusts the signatures generated for that PC in the first place (especially if they have been signed in a way similar to the above). But the trust model can be extended to include other signing authorities. But this means it would now be possible to have a Windows PC which has the trust model extended to include say Microsoft, Adobe and Cryptzone, so it can now self update without any need to in-house manage the changing signatures. Effectively the management of the signatures of the good stuff has now been outsourced in much the same way as for anti-virus.

Today application whitelisting is being most aggressively adopted for industrial control, medical and manufacturing systems, where there is less configuration variation and the outcomes of infection are potentially extremely severe. Wider PC protection will undoubtedly follow as the trust model is extended to include other signing authorities and it becomes possible to self-update signature files without the need for in-house management.

So coming back to the question, who is in control of your infrastructure today? If you are solely relying on anti-virus protection you can no longer be sure. However if you apply certificate based application whitelisting, you can be certain that just you and any developers you choose to allow will be able to take control – and that will be all!