Executive brief

Commissioned by

Taking Cyber Security to the Next Level

The role of continuous egress monitoring and analytics

September 2015

In a nutshell

As cyber security threats become more sophisticated, targeted and persistent, it's no longer possible to block all attacks. Measures to detect intrusions are now at least as important as those designed to prevent them. With this in mind, the rapidly evolving area of egress monitoring has a lot to offer. While not a magic bullet, implemented as part of a blended strategy, it can significantly reduce your cyber security risks.

About this Document

The insights presented in this document are derived from ongoing independent research, coupled with specific briefings from RedSocks on its 'Malware Threat Defender' solution. While the latter is used to illustrate how the principles discussed translate to practical reality, this should not be taken as a product verification or endorsement.

So much value, so many routes to exploiting it

The 'dark market' is thriving. Whether it's stolen intellectual property, pilfered identities, misappropriated customer data, or access to botnets powered by hijacked computers, there's a lot of money to be made from cyber crime. This in turn means cyber criminals are well-funded and well-resourced. Continuous R&D allows new vulnerabilities to be discovered and shared very quickly, and ever more sophisticated tools and techniques to be developed and rolled out at a pace that most legitimate IT vendors would envy.

In today's hyper-connected business environment, cyber criminals have many routes to get to their objective.

Meanwhile, in today's hyper-connected business environment, cyber criminals have many routes to get to their objective. They can hack internet-connected systems directly, use malware or social engineering to come in via your users, and even exploit B2B access mechanisms to attack your systems via trusted customers, partners and suppliers (or to use your systems and users to attack them).

Furthermore, attacks nowadays are much more likely to be targeted than in the past. The amount of data shared on social networks and via electronic publishing and trading means that approaches can be tuned to the way your organisation and your people work. Add advanced persistent threats (APTs) into the mix, which blend techniques to quietly penetrate your systems undetected over a long period of time, and the modern cyber security challenge is very clear (Figure 1).

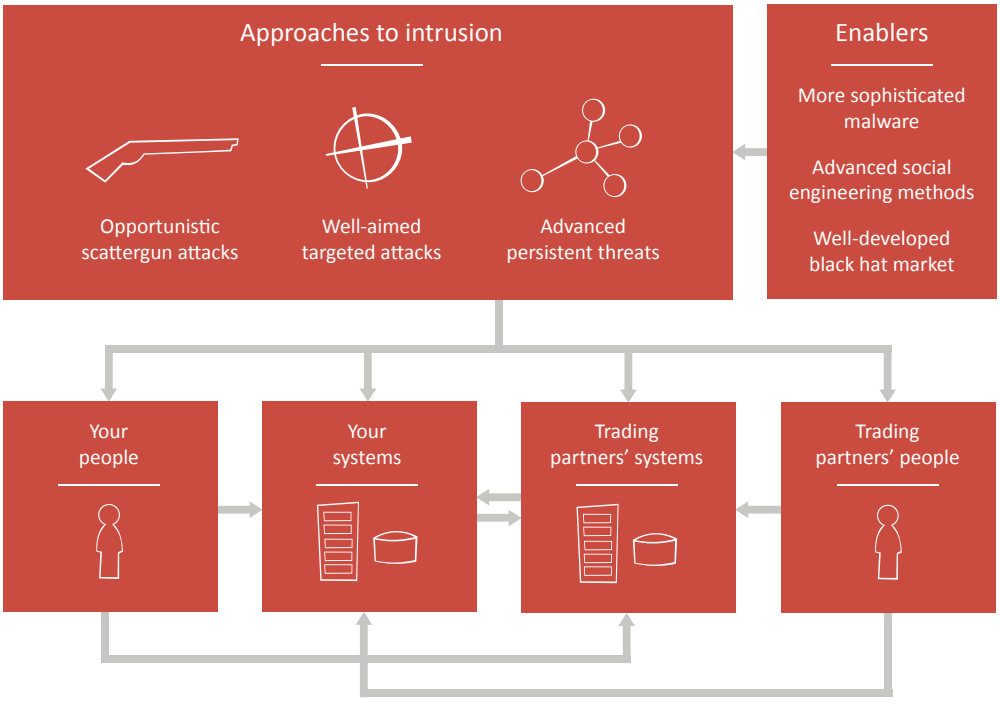

Figure 1

Well-resourced cyber criminals are reaching into your systems and data using sophisticated techniques to exploit available 'attack vectors'

Many security experts argue that the days of being able to prevent intrusion into your network are over.

Against this background, many security experts argue that the days of being able to prevent intrusion into your network are over. There are just so many entry points into your business systems and data that trying to protect them all is futile. Of course this doesn't mean we should all give up and let the cyber criminals run free, but it does mean that we must focus our efforts and investments differently than in the past.

The intrusion detection imperative

The sooner a security breach is discovered, the more quickly it can be dealt with, and the less damage or loss will result.

The sooner a security breach is discovered, the more quickly it can be dealt with, and the less damage or loss will result. It's a simple principle that undoubtedly won't be news to you, but it's one that highlights the importance of intrusion detection given the threat landscape described above.

The trick for controlling risks is to strike the right balance between prevention and detection. The former is clearly important to minimise the number of intrusions and related direct exposure, but also to keep the amount of 'noise' to a minimum, making the job of detection much easier. Effective intrusion detection then allows you to react quickly and minimise damage when a breach occurs or a dormant threat that has previously infiltrated your network becomes active (Figure 2).

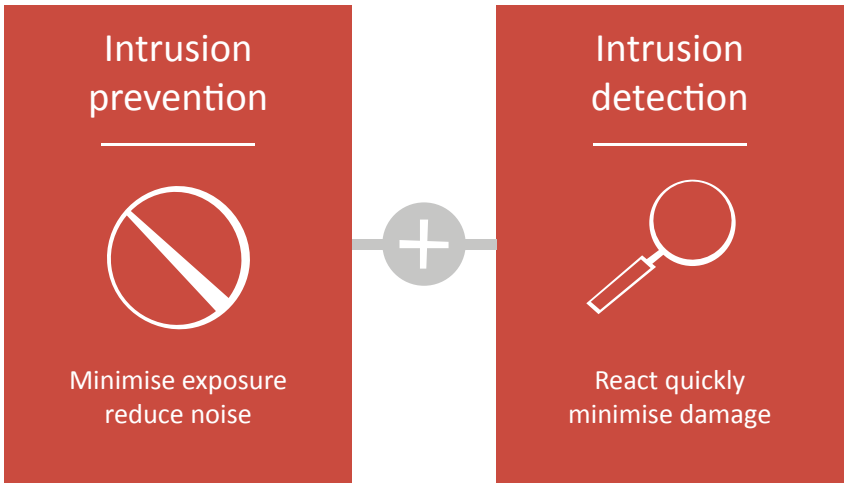

Figure 2

The right blend of intrusion prevention and detection is the key to minimising risks

Most IT and security professionals know they could do a lot better when it comes to cyber security.

Now all of this might sound very obvious, and there's a good chance you will have technology and processes already in place to implement this kind of approach to one degree or another, particularly if you work in a larger organisation. Something that comes through very strongly in our research, however, is that most IT and security professionals know they could do a lot better when it comes to cyber security. Indeed many say that unless they strengthen their approach, they will find themselves increasingly more exposed over time as the abovementioned trends continue.

Assessing your current situation

You can never care too much about your intrusion detection measures.

So how do you know when you have been breached? Sadly for some it's the point at which stolen data surfaces in the outside world, a distributed denial of service (DDOS) attack gets traced back to compromised machines on your network, or some equally disturbing or costly event transpires. And this is despite the fact that organisations falling victim to such cyber criminal activity are making the effort and have a range of protection mechanisms in place. With this in mind, you can never care too much about your intrusion detection measures.

The most immediate objective is to identify malicious activity and initiate appropriate remedial action.

One way to think about effective intrusion detection is that it generally revolves around a central monitoring process that gathers and acts on data and events generated from a range of different sources. The most immediate objective is to identify malicious activity and initiate appropriate remedial action, though an ability to report events, incidents and trends for both compliance purposes and continuous security optimisation is also important (Figure 3).

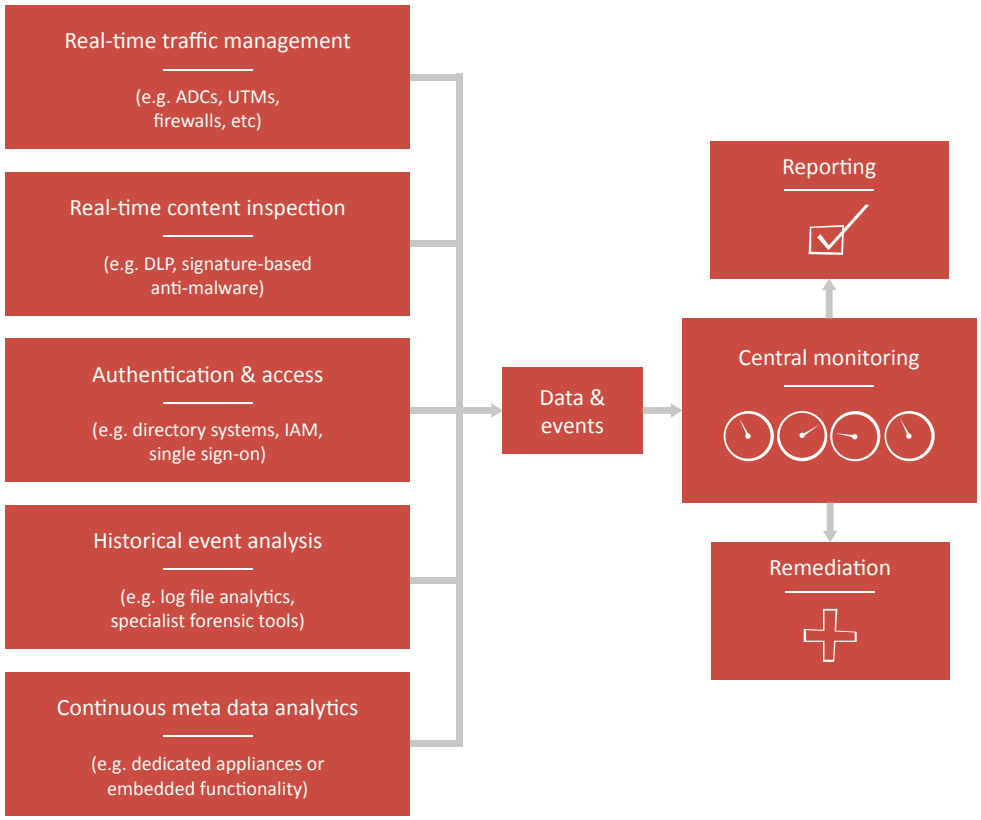

Figure 3

Effective intrusion detection generally revolves around a central monitoring process that gathers and acts on input from a range of sources

With a reliance on pre-set rules and no inherent 'memory' of what's gone before, most network devices cannot detect suspicious patterns over a timeline.

The reason for taking data and events from so many different sources is because they each have their strengths and weaknesses.

Real-time traffic management systems such as firewalls, unified threat managers (UTMs) and application delivery controllers (ADCs) are great for flagging up 'point in time' events. But with a reliance on pre-set rules and no inherent 'memory' of what's gone before, most network devices cannot detect suspicious patterns over a timeline.

Real-time content inspection systems are similarly useful but limited. Much of the functionality of data loss prevention (DLP) solutions, for example, is undermined when data is encrypted. And along with signature-based anti-malware technology, and most authentication and access systems, they are also generally geared up to dealing with known threats and/or scenarios rather than the unknown or unusual.

Tools focused on analysis of history are generally only run on a periodic batch basis.

As a direct result of the limitations mentioned, many IT and security teams have implemented solutions that allow data and events to be analysed retrospectively to yield greater security insights. These range from the very simple, e.g. tools designed to spot trends and anomalies in individual device logs, to complex 'big data' setups capable of cross referencing data-sets and drawing inferences from many detailed sources. The common challenge here, however, apart from the need to manage large amounts of data, is the time it takes to generate useful output. Tools focused on analysis of history are generally only run on a periodic batch basis.

You can't rely on just one input into your intrusion detection and monitoring process – it needs to combine data and events from multiple sources.

Of course the above comments represent a generalisation of the kind of limitations that commonly exist. You may have more sophisticated solutions or hybrid configurations in place that overcome some of the challenges. We would be surprised, however, if you didn't recognise many of the issues highlighted. The point is that you can't rely on just one input into your intrusion detection and monitoring process – it needs to combine data and events from multiple sources to be effective.

On that note, the one potential input source from the above figure we haven't mentioned yet is continuous metadata analytics. This is because such technology is much less common, which is in turn why we'll be spending the remainder of this paper looking at what it is and how the 'egress monitoring' it enables complements the other solutions we have touched on.

Introducing 'egress monitoring'

While trying to detect or block an initial security breach is extremely difficult because of the sheer number of attack vectors, looking out for suspicious traffic leaving your network is actually a very good way of identifying malicious activity.

While the cyber vandalism and purely destructive attacks that defined the early days of malware have not gone away, the bulk of today's threats are motivated by financial gain or political leverage. In order for an attack to deliver on one or both of these objectives, communication needs to be established between your network and the cyber criminal's system. The most obvious form this takes is a connection to facilitate the exfiltration of data, either pulled or pushed as a result of direct hacking or malware infection. In the case of the latter, communication may also manifest itself as signalling traffic, e.g. when an infected machine is being remote controlled as part of a botnet, or is sending out a beacon to signal its readiness for activation.

As it turns out, while trying to detect or block an initial security breach is extremely difficult because of the sheer number of attack vectors, looking out for suspicious traffic leaving your network is actually a very good way of identifying malicious activity. In conceptual terms, it's about creating an 'egress trap' which is able to tell the difference between legitimate and malicious traffic, and generate alerts accordingly (Figure 4).

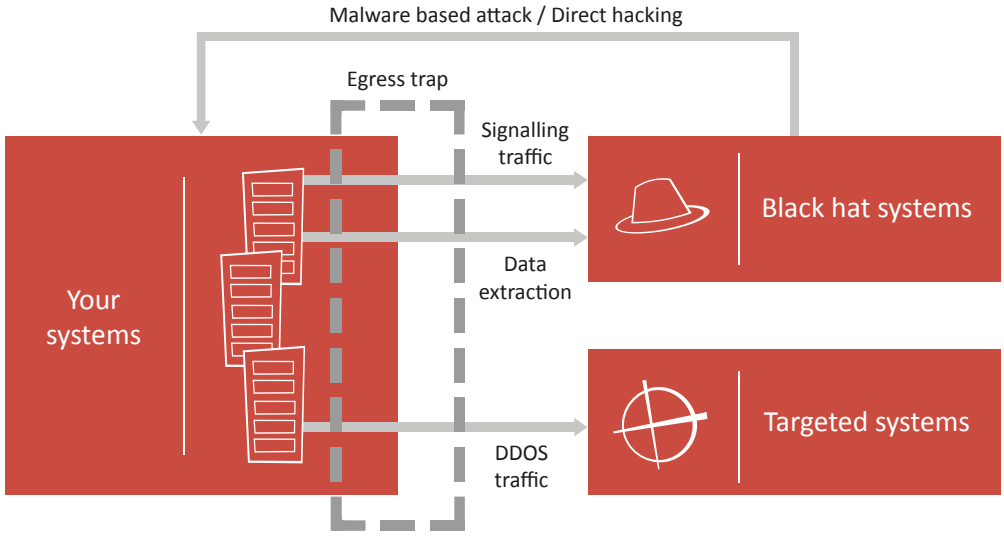

Figure 4

Regardless of the nature of the original breach, an egress trap is designed to detect when malicious activity is taking place

An algorithmic evaluation of traffic legitimacy is performed on the fly.

So what's the difference between this approach and the way in which a next generation firewall or similar solution operates? Well rather than implementing straightforward filtering based on static pre-defined rules, with an egress trap or monitor, an algorithmic evaluation of traffic legitimacy is performed on the fly taking historical activity as well as the immediate event into account.

An egress monitor focuses purely on metadata, as contained in the (typically) 1-2% of network packets that detail the source, destination and nature of the traffic.

Focus on continuous metadata analysis

In practical terms, full and detailed in-flight analysis of all outbound traffic from a sizeable network on an ongoing basis would translate to a huge amount of overhead and expense. An egress monitor therefore focuses purely on metadata, as contained in the (typically) 1-2% of network packets that detail the source, destination and nature of the traffic.

Most mainstream network devices such as switches, routers, firewalls, etc make such 'flow data' available for monitoring and analysis purposes, and with the right technology in place, this can be tapped into on a real-time basis without intercepting or otherwise interfering with the main network traffic flow itself. The end result is that the detail needed to facilitate continuous metadata analysis can be obtained extremely efficiently – it doesn't even require an excessive amount of storage. Furthermore, the approach is not undermined by traffic encryption.

Judgements are made on a near real-time basis on whether threats are present in the flow of traffic out of your network.

But how does it work?

At a high level, each interaction (or egress event) derived from flow data is assessed for legitimacy using algorithms that make reference to both threat intelligence data and monitoring history, and judgements are made on a near real-time basis on whether threats are present in the flow of traffic out of your network (Figure 5).

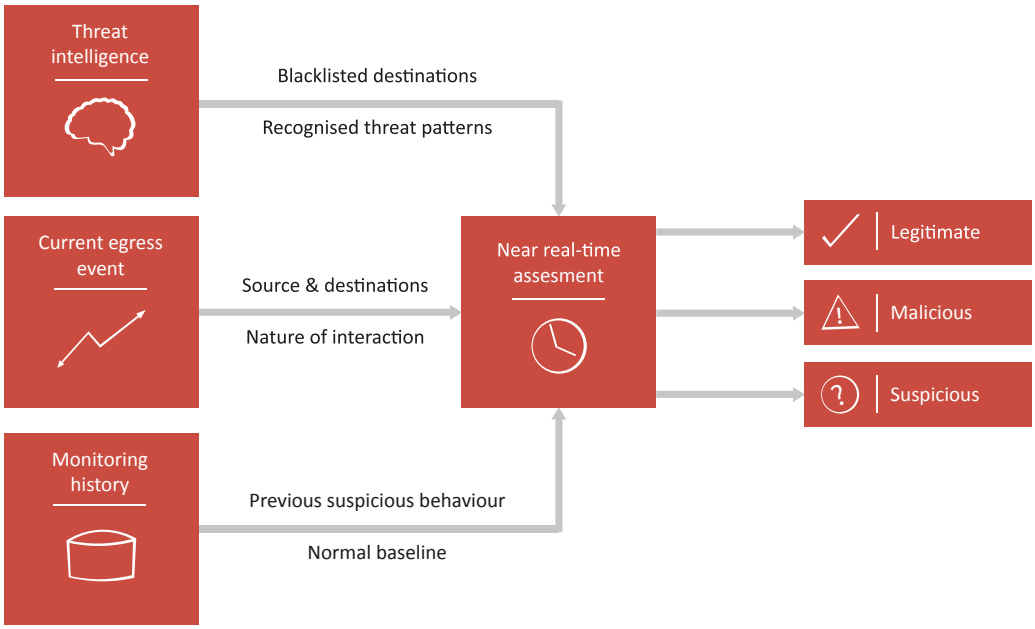

Figure 5

Assessing the legitimacy of traffic leaving your network

While it may not be obvious at first glance, the differentiating element within this approach is the presence of a notional 'memory'. As a simple example, a single ping as part of a beacon signal may not in isolation be enough to indicate infection of a particular device on the network. Put it together with similar pings over the past few hours or days, however, and the pattern observed along a timeline can flag up the presence of an active piece of malware with a high degree of confidence.

The pattern observed along a timeline can flag up the presence of an active piece of malware with a high degree of confidence.

Abnormal traffic patterns associated with data exfiltration, botnet remote control, DDOS hijacking in progress, and so on, can similarly be detected with a good level of accuracy.

The process is not dissimilar to that which takes place in a court of law. Various pieces of evidence are considered to make a judgement.

If you are familiar with the idea of security information and event management (SIEM), then one way of thinking about what's going on here is the egress monitor fulfilling a similar function to an analyst working in a control centre environment. Rather than a human scanning screens for suspicious events, trends and patterns, a similar process takes place automatically behind the scenes on a continuous basis.

But just as a human analyst makes value judgements, if you look more closely at the way in which assessments are made by an egress monitor, the process is not dissimilar to that which takes place in a court of law. Various pieces of evidence are considered to make a judgement of innocence or guilt, or in this case whether an interaction is legitimate or malicious (Figure 6).

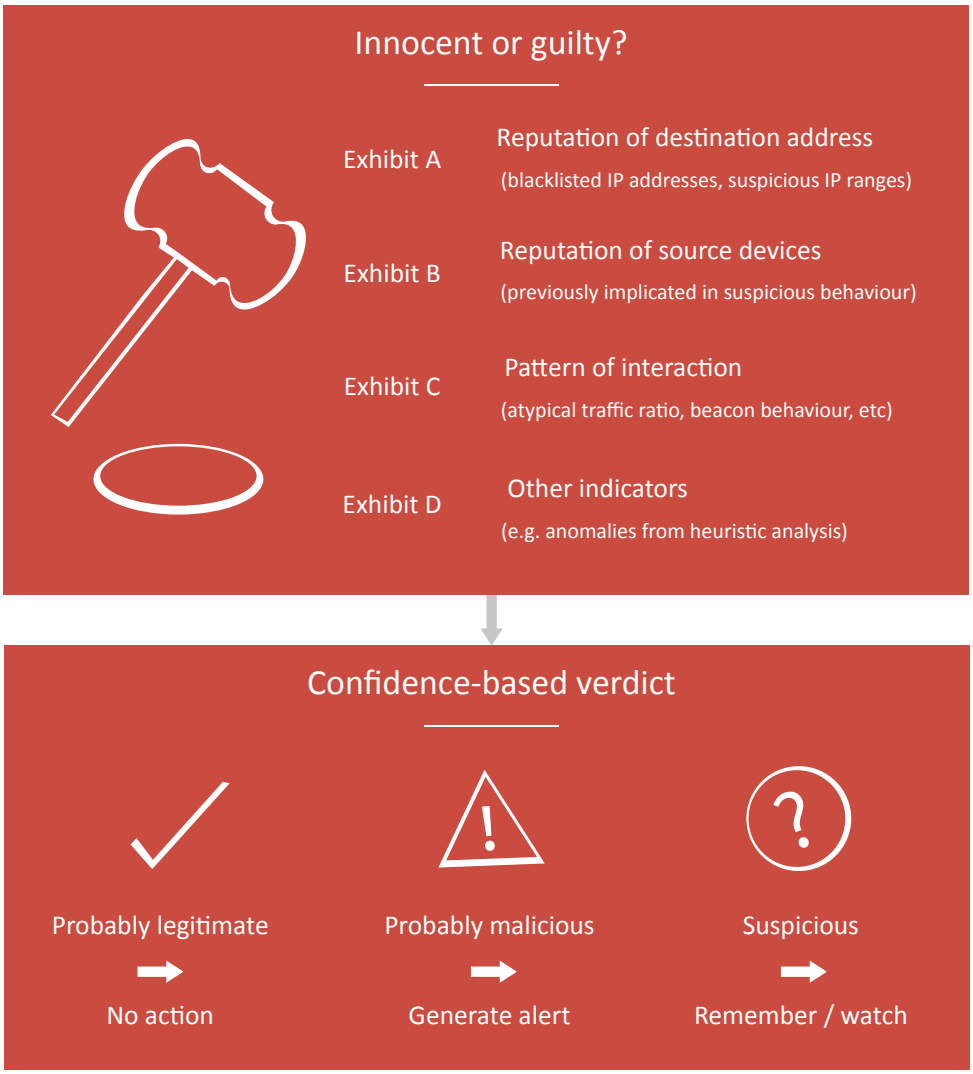

Figure 6

An evidence-based approach to identifying malicious traffic

If some aspects of an interaction are merely suspicious, it may not be immediately flagged, but it will be logged and effectively 'remembered'.

The picture we see here illustrates that egress monitoring is generally based on a probability-centric approach. If the evidence convincingly suggests the presence of malicious activity then a threat alert is raised. If some aspects of an interaction are merely suspicious, it may not be immediately flagged, but it will be logged and effectively 'remembered' when assessing future events, hence could still ultimately contribute to a guilty verdict once more evidence is gathered. While the possible threat is not 'watched' in an explicit sense, the net effect is the same.

While talking about practicalities, it is worth stressing that the continuous metadata analytics technique underpinning egress monitoring is not a total answer to the threat

detection problem in its own right. As per our earlier schematic (Figure 3), egress monitors are designed to sit alongside other solutions that feed into the overall monitoring process.

So much for the generics; let's take a look at an example of a real life egress monitoring product that exploits continuous metadata analytics in the manner described, and how this kind of technology is implemented in practice.

The 'Malware Threat Defender' (MTD), manufactured and supported by the Dutch security specialist, RedSocks, was arguably the first dedicated egress monitoring offering to appear on the market.

Real world implementation

In order to illustrate the practicalities of implementing the egress monitoring approach to cyber security enhancement, it is useful to look at a specific solution in this space. An appliance known as the 'Malware Threat Defender' (MTD), manufactured and supported by the Dutch security specialist, RedSocks, was arguably the first dedicated egress monitoring offering to appear on the market (Figure 7).

Figure 7

Example of a real world egress monitor

RedSocks Malware Threat Defender

Self-contained appliance implementing continuous metadata analytics

19" rack-based form-factor

Solutions of this kind are designed with rapid implementation and low maintenance in mind, but there is obviously some work to be done to get them up and running.

Integration with your existing equipment

Integration with the most commonly used routers, switches and firewalls is easily achieved.

The first job, once the appliance is installed in your datacentre or computer room, is to hook it into the network devices from which flow data will be captured. The RedSocks MTD supports a number of options to facilitate this, and integration with the most commonly used routers, switches and firewalls is easily achieved either natively or via probes appropriately located on the network.

RedSocks provides best practice guidelines to help with this, along with professional/partner services as required to deliver assistance where necessary.

Configuring threat detection settings

RedSocks invests significant time, resources and expertise making sure its devices are always using the most up-to-date assessment algorithms and the latest available threat intelligence.

One of the big advantages of using a solution such as the RedSocks MTD is that you don't need an intimate knowledge of how to configure the equipment to detect different types of threat. RedSocks invests significant time, resources and expertise making sure its devices are always using the most up-to-date assessment algorithms and the latest available threat intelligence. As part of this it brings together the output from an extensive team of security analysts, malware researchers, 'tame' hackers and software developers with feeds from third party providers.

RedSocks MTDs leverage and are kept updated with all of this, and as a result you can make monitoring and alerting decisions based on threat levels and categories that are easy to understand and map onto a business risk view of the world. You can then enable alerting to strike the right balance, e.g. to raise alerts for important events, but not clutter up dashboards and other monitoring tools with low priority activity.

You can pinpoint when in time a given threat emerged, which is invaluable when performing potential damage assessments or investigating possible APTs.

Having said this, the RedSocks MTD will continuously monitor for all types of threat and keep a complete log, backed up with historical metadata. This is important for periodic review to identify and clean up 'nuisance' level threats. It also means you can pinpoint when in time a given threat emerged, which is invaluable when performing potential damage assessments or investigating possible APTs.

In order to get the most from an egress monitor, however, it needs to be implemented as part of your broader intrusion detection and response framework.

Integrating egress monitoring into your overall process

As mentioned earlier, egress monitoring is about identifying threats and intrusions quickly, and raising the relevant alerts, so incidents may be assessed and where necessary remediated in a timely manner to minimise loss and damage. Such assessment and remediation can be achieved through manual processes, and the MTD will surface the data required. This includes events and statistics relating to the types of threats detected, the destination of outbound traffic, the identity of compromised devices, and a range of other relevant details.

The RedSocks MTD publishes standard APIs to allow easy integration into third party environments.

If you work in a larger enterprise environment, the chances are that you will already have higher-level security monitoring and management tools in place, in which case you would want egress monitoring events and alerts to feed into these. In order to facilitate this, the RedSocks MTD publishes standard APIs to allow easy integration into third party environments. The specifics of how integration is achieved will vary from case to case, but the work required need not be onerous. Again, RedSocks (and its partners) can provide guidance on how to get things working together effectively.

In the meantime, for a visual representation of how everything we have discussed fits together in a RedSocks enabled environment, take a look at the infographic presented in Appendix A.

Final thoughts

As the nature of the cyber threat becomes ever more multi-faceted, so too must the nature of your cyber defences.

As the nature of the cyber threat becomes ever more multi-faceted, so too must the nature of your cyber defences. Relying on a small, fixed set of solutions to protect your systems and data is not enough, especially if these are based on traditional technologies that only deal with known specific threats. In today's hyper-connected business environment you need multiple tools and approaches to minimise the chances of intrusion, but you must also have systems and processes in place to identify and deal with breaches when they inevitably occur.

This is where egress monitoring and alerting solutions come into play. Technology such as the RedSocks Malware Threat Defender will increasingly become an essential part of the intrusion detection jigsaw, fitting alongside other important pieces of the puzzle such as traditional blacklisting approaches, heuristics-driven anti-malware solutions, data loss prevention tools, and even business-level fraud detection systems. The key message is that whatever some security vendors might claim, there are no magic bullets in this space.

Technology such as the RedSocks Malware Threat Defender will increasingly become an essential part of the intrusion detection jigsaw.

Overarching all of this is the fundamental principle that risk management is a probability based game. You cannot implement an absolute level of either intrusion prevention or detection, so the objective is to stack the odds in your favour as much as possible. We hope our discussion in this paper has provided some ideas on how egress monitoring might help you to do this.

Further reading

Freeform Dynamics

The following reports are available from the Freeform Dynamics website (www.freeformdynamics.com)

1. Mobile Working without the Tears
   www.freeformdynamics.com/fullarticle.asp?aid=1792
2. User-Centric Mobile Security
   www.freeformdynamics.com/fullarticle.asp?aid=1855
3. Future Proofing Your Network
   www.freeformdynamics.com/fullarticle.asp?aid=1852
4. Controlling Application Access
   www.freeformdynamics.com/fullarticle.asp?aid=1778
5. The End User Security Jigsaw
   www.freeformdynamics.com/fullarticle.asp?aid=1688

RedSocks

The following material is available from the RedSocks website (www.redsocks.nl)

1. RedSocks Malware Threat Defender Endpoint Threat Detection (Product sheet)
   www.redsocks.nl/files/RedSocks_ProductSheetGartner_2015_EN_FINAL_site.pdf
2. How does RedSocks Malware Detection Work? (Infographic)
   www.redsocks.nl/files/Redsocks_HowDoesMTDWork_ENG_FINAL_website.pdf
3. Data Breaches, Smart Companies come Prepared (Solution brief)
   www.redsocks.nl/files/Redsocks_Databreaches_2015_EN_Final_interactive.pdf

Appendix A: Overview of a RedSocks deployment

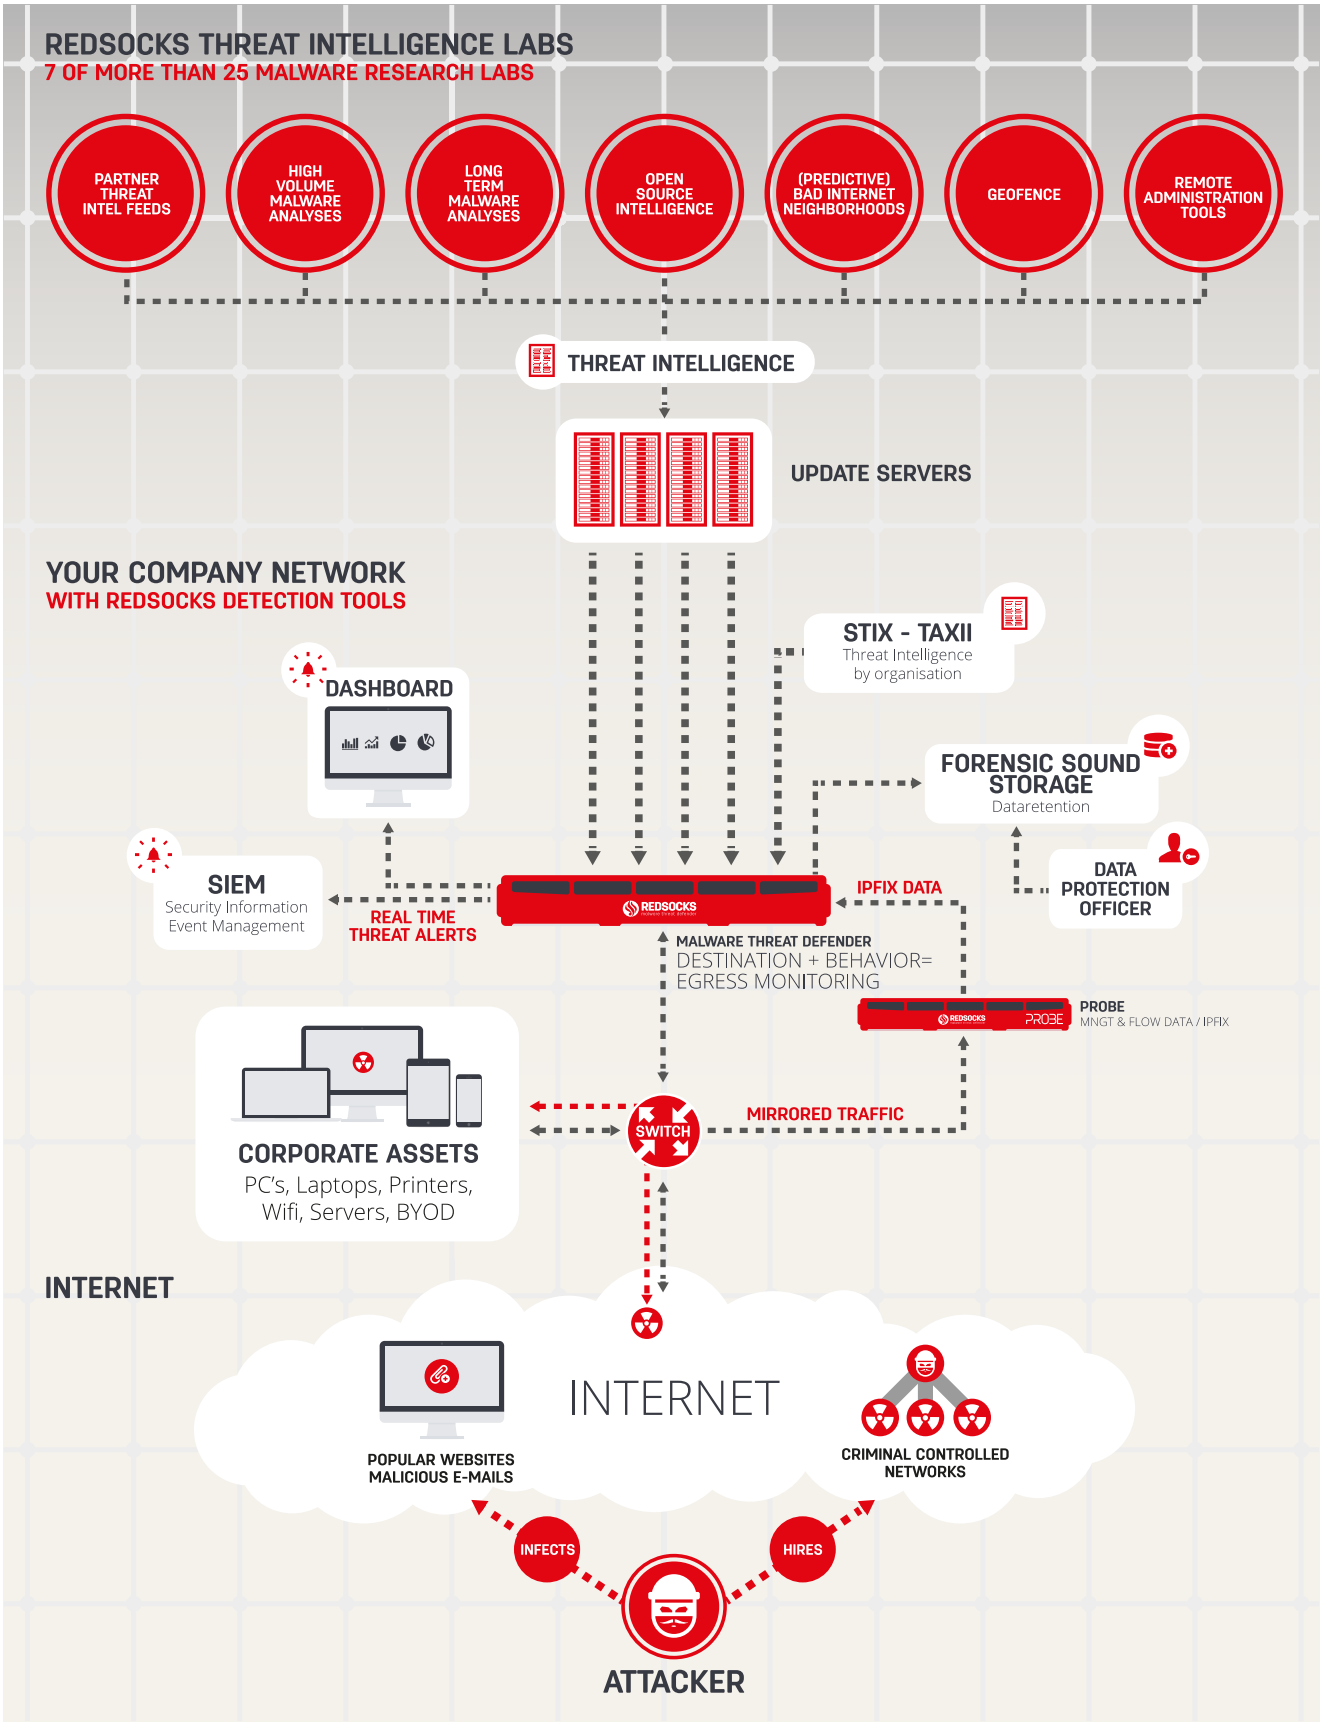

About Freeform Dynamics

Freeform Dynamics is an IT industry analyst firm. Through our research and insights, we aim to help busy IT and business professionals get up to speed on the latest technology developments, and make better-informed investment decisions.

For more information, please visit www.freeformdynamics.com.

About RedSocks

RedSocks is a Dutch IT security company specialising in detecting and defeating malware in the corporate market. RedSocks supplies the RedSocks Malware Threat Defender (MTD) as a network appliance. This innovative application analyses outgoing digital traffic in real-time, also known as 'intelligence-based egress monitoring', based on risk lists and algorithms compiled by the RedSocks Malware Intelligence Team. The team is made up of specialists in spotting new threats on the internet, and translating this into state-of-the-art malware detection. RedSocks is an HSD (Hague Security Delta) partner.

For more information go to www.redsocks.nl or follow us on Twitter @RedSocksMTD.

Terms of Use