invisiron®

Fortifying Existing Cyber Defence by Introducing an Invisible Layer Before the Firewall Appliance

INVISIRON® WHITE PAPERS
version 1.2

INTRODUCTION

Many of the security solutions on the market today focus on threat detection and do not provide defence mechanisms to keep attackers out of the network. Instead, most legacy solutions focus on log collection and the use of a Security Information and Event Manager (SIEM) software suites, coupled with manpower, to analyse and correlate these logs. This is insufficient at protecting against penetrations, data exfiltration, SQL injection and many of the more sophisticated attack types we see today.

Invisiron®'s Cyber Defence platforms have been designed using state of the art cybersecurity techniques. This paper will describe how to protect a network from the more sophisticated cyber-attacks we see today, identify some of the shortfalls in legacy security solutions, and provide insight into the weaknesses in networks and how they may be better protected.

MODERN DEFENCE STRATEGIES

Invisiron® believes that, as a primary line of defence, an organisation's network must be defended from a point as close to their Internet gateway as possible. This typically means installing a capable security hardware device immediately after the Internet router and in front of the organisation's primary firewall, it may also involve installing devices strategically located immediately in front of critical assets or at key gateway points in the network. This architecture allows for preventing intrusions and stopping attacks before they reach inside the network or asset. Stopping attacks before they take a foothold in the network is the foundation of modern cyber threat defence. Most people will agree that keeping an intruder out of the network should be the primary goal of any security solution. After-the-fact detection does serve a purpose since even the best cybersecurity platforms cannot guarantee that all attacks will be prevented. A properly defended network comprises of several layers of defence including Security by Design, Browser Isolation & Sandboxing, Endpoint Protection, Monitoring & Prevention, Policy & User Awareness, Cyber Threat Intelligence (CTI), and a modern hardware solution such as Invisiron® to stop threats in their tracks.

Invisiron®'s Cyber Defence platforms have been designed in a unique way that allows them to be installed inline between an Internet router and a firewall, in front of critical assets, or at key gateway points in the network. Packets can be inspected in real-time as they pass through the platform and in both directions (inbound and outbound protection). Malicious packets can be dropped instantly before they have a chance to enter into the network or infect a key asset and before vital information leaves the network.

It is crucial to monitor packets flowing in both directions at your organisation's Internet gateway. Often it is believed that network security is about monitoring incoming packets from the Internet. This is only

half the truth: it is equally important to monitor outgoing traffic as well. For example. Should a malware successfully infiltrate the network i.e. via a thumb drive or other media, there is then a chance it can be detected by monitoring the outgoing traffic it generates.

The current state of the art in defence strategies requires in-line hardware-based protection which can not only analyse all traffic in both directions (many DPI systems today only analyse header data which equates to about 10% of actual traffic) but can do so at or near line rate (no impact on system latency). The system must as a minimum provide detection engines for malicious data signatures, tracking and blocking of TOR exit nodes, targeted rate limiting, as well as traditional IP & Domain reputation detection. More advanced methods should include the ability to expand the detection of packet signatures to more complex rules to provide detection and protection against protocol-based attacks or pre-indicators of a potential attack such as scan attempts. Importantly all Invisiron® Cyber Defence platforms provide these functionalities and more while being capable of monitoring and protecting traffic in both directions (bi-directionally) at any point in the network.

LEGACY DEFENCE PRODUCTS

Before going into more detail on how Invisiron®'s Cyber Defence solution can protect your systems, this paper will provide an outline and comparison with legacy security technologies, endpoint protection systems, and architectures which many organisations still choose to utilise today.

SOFTWARE ENDPOINT PROTECTION

Software-based security technologies provide limited protection, in that many devices in an organisation's information technology ecosystem cannot be protected by software alone. Examples are network-attached storage, printers, webcams and other Internet-Of-Things (IoT) devices. Additionally, software protection solutions are often rendered ineffective by more sophisticated hacks before the attack or new threats find ways to exploit vulnerabilities in lower system layers which the software has no control over e.g. Stuxnet, Meltdown, and Spectre. An ever more common attack technique is to penetrate a network using an IOT device, such as a network-attached printer, and use it as a base for further penetrations; lateral movement inside the network, and even data exfiltration.

In other examples, we have seen Internet-connected webcams being integrated with extremely poor security awareness. Many of these devices use fixed factory set login credentials or password reset backdoors. On a global scale, these devices are regularly being used in distributed denial of service (DDoS) attacks, crypto mining hijacking, as well as new emerging threats created by the permeation of IoT solutions throughout the worlds IT systems. Software layer protection alone is not enough to stop a

sophisticated attacker from infiltrating and taking control of such devices nor is it effective against attacks such as an attack on unpatched OS vulnerabilities, firmware bugs or backdoors etc. The bottom line is, a network must be defended immediately at the Internet connection, as well as at critical points, to stand a chance of blocking attacks.

LOG COLLECTION & SECURITY INFORMATION EVENT MANAGEMENT (SIEM)

Another popular approach is the use of log collection and SIEM software tools. The premise is to collect log information from as many devices as possible and analyse this data to detect malicious activities in the network. This offers some protection against penetrations and other types of attacks, however without an inline device that can analyse and drop or block malicious packets the protection is limited. Most of these systems focus on providing visibility into what has taken place in the network while requiring costly and manpower intensive manual response. The delay alone in this approach renders many attacks over before they can be stopped or mitigated.

Some more sophisticated systems rely on sending TCP reset packets to delay an attack, however, even these can be ignored or countered by a sophisticated adversary. There is some value in understanding what has taken place in an organisation's network in the past, this is better than knowing nothing at all, however knowing about a past event does not necessarily protect against all types of attacks or malicious activity.

Log collection and SIEM software alone cannot protect against particularly the more serious and damaging cyber-attacks unless combined with a proper hardware-based physical Cyber Defence system. Without such a Cyber Defence system organisations can gain various statistics and analytics regarding past activities. This information could be valuable, but alone do not provide the required protection against today's more sophisticated cyber-attacks.

INTRUSION DETECTION

Many older legacy solutions claim to provide intrusion detection. They are typically marketed as Intrusion Detection Systems (IDS) or Network Intrusion Detection Systems (NIDS) and are still being deployed in many networks today. These devices are normally installed via a tap port on a switch. A tap port is a special port on a switch that can be configured so that it sends out mirrored copies of all network traffic.

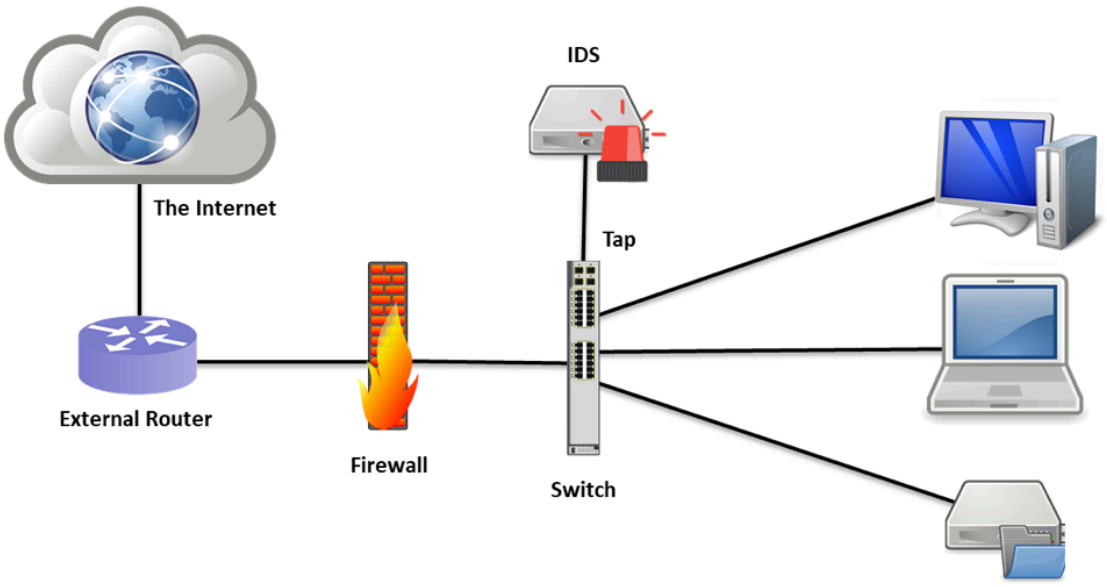

Typical Network Monitoring Installation from a Local Area Network (LAN) Tap

There is limited value in knowing that the network has already been breached. Invisiron®'s approach is to instead disrupt the attacks before they have a chance to enter into the network. Most IDS/NIDS systems on the market today do not protect a network, as they are not capable of blocking malicious network traffic. They aim to detect an attack, determining that a breach has occurred, the expected mitigation then requires human intervention to address the event. Unless someone is monitoring these systems around the clock, it might be hours or even days before a network administrator is made aware that their network has been breached. Additionally a serious drawback of these systems is that as a rule tap ports on switches have a lower priority so as not to cause network disruption, the result, in modern networks with their ever more increasing data rates, is that many of the packets which pass through the switch are not mirrored and as such never seen by the IDS/NIDS or connected SIEM system. Passively listening to network traffic alone is inadequate in countering today's threat landscape.

Active solutions may provide a better level of protection by employing disruptive measures such as sending TCP reset packets to a potential threat endpoint or an infected device. A major drawback of this approach is that a sophisticated attacker can simply set the device to ignore these TCP reset packets in effect rendering ineffective this method of protection.

SIGNATURE BASED MALWARE PROTECTION

Another popular category of security technology is software utilities which are installed directly on the host machine which base their protection on file signatures. If, for example, a user downloads a file from the Internet, a checksum (signature or fingerprint) of the file will be calculated by the software. This signature is then sent to a centralised server owned by the company that provides the security software. It will be compared against other signatures from both known good files and against signatures of

known malicious files. This type of protection might have some value, provided that the vendor's database is updated fast enough with new threat signatures and contains a large enough database of signatures from the past. It is however very easy to bypass this defence. All the hacker has to do is slightly rearrange their malicious file so that it generates a different file signature, which will then easily bypass any signature-based protection solution.

Additionally, this type of signature-based protection is not capable of detecting attacks that are not delivered through a file. It is also not capable of detecting threats that do not use persistence, meaning that they do not store information on the system's hard disk drive. Furthermore, network penetration or breaches cannot be detected by these solutions.

Another category of signature-based defence products sends complete files from the owner's computer up to a server for offline pattern matching. There have been many reports in the news media about this, where high profile companies have had important payroll files and other confidential data exfiltrated out of their network and sent to the security software vendor's signature servers for analysis. In the modern privacy-conscious landscape this method is not viable. What's more, while not only counter too many privacy regulations, in some cases this approach can actually increase the risk of certain types of attacks. If the network contains sensitive data and files, one must be aware that this data might be exfiltrated out to be analysed externally by some systems.

In summary, signature-based malware protection has little value when used alone. It should rather complement properly placed hardware systems, and be part of an overall defence strategy.

MODERN CYBER SECURITY SOLUTIONS

In today's sophisticated attack environment protecting an organisation's networks is no simple task. The first and most important step is endeavouring to keep attackers out of the organisation's networks. If an attacker is unable to penetrate a network from the outside, we greatly reduce the risk factor to the networked systems. A hacker that can penetrate the network from the outside may be able to install what's known as a 'reverse shell', an attack method which gives the attacker full control of the target machine. Keeping attackers out of a network requires the use of Cyber Defence platforms installed inline with the Internet traffic and as close to the Internet gateway as possible. Only when installed inline can the device inspect and block malicious traffic in real-time. It is of the highest importance that Cyber Defence protection mechanisms include inline devices positioned as close to the Internet gateway as possible to allow them to detect and prevent attacks before they reach inside a network. It is also of very high importance that these platforms are capable of blocking malicious traffic by stopping the network packets as soon as they are detected as being malicious. If such packets can enter into a

network, the damage may already be done. These packets must be stopped and dropped as soon as possible before they have a chance to infiltrate the network.

Software solutions alone running on computers inside the network offer a limited defence, in that a software-based solution does not monitor and stop packets inline and directly at the point where traffic enters the network thus leaving each and every computer or device within the network vulnerable to attack.

Some inline installed defence systems suffer from performance degradation, and/or high failure risks. These weaknesses have been overcome by Invisiron®'s unique design architecture, further details on these important points can be found in the section 'Inline Performance' of this paper.

MINIMUM NETWORK DEFENCE CAPABILITIES

A proper defence implementation for any network must include at minimum three different types of advanced cybersecurity capabilities.

- A deep packet inspection engine, or DPI engine.
- A reputation detection engine.
- An evidence collecting and logging engine.

The DPI engine is responsible for detecting and blocking break-in attempts, disrupting malware implants, remote control attempts, preventing data exfiltration, operating system fingerprinting (determining the OS type and vulnerabilities of a potential victim without having an account or logging in directly to the machine), and more.

The reputation detection engine is responsible for checking IP addresses, domain names and URLs against lists of known malicious actors. These two engines complement each other and offer very strong protection.

One of the common complaints by security auditors or researchers after an attack has occurred has been the lack of evidence available to determine exactly what has transpired and improve future defences. Invisiron®'s systems include full packet logging capabilities where the entire network traffic for an attack is kept in a standard pcap format which can be interpreted by network analysis tools. This is vital in improving future protections and determining exactly what has transpired particularly in cases of more sophisticated attacks.

Cyber threat intelligence (CTI) is another key part in a comprehensive Cyber Defence, CTI is the act of keeping the aforementioned engines updated with new threat types. More details on these engines and systems follow.

DEEP PACKET INSPECTION ENGINE

The DPI engine is responsible for performing pattern detection on every network packet. The engine is rule-driven which means it is controlled by a set of externally provided inspection rules that have been compiled down into machine-executable code. The rules contain instructions for what to look for and where to look inside a packet. Rules can be complex to write and understand, especially if they need to be written in a low-level syntax. Writing low-level rules is difficult and requires an in-depth understanding of the rule language syntax.

To counter this complexity Invisiron®'s security platforms incorporate an easy to use Graphical User Interface (GUI) based tool that allows for rules to be written in a much simpler way, without requiring detailed knowledge about the underlying rule command syntax. Invisiron® includes what is called a system rule set through a cyber threat intelligence (CTI) feed. The system rule set includes a wide range of rules that protect against common attacks and penetration attempts as well as some of the more sophisticated attacks like protocol exploits. Examples of such rules are those which block port scans or OS fingerprinting.

The system rule set is updated with new rules as new attack types are detected. Examples of such attacks are the widely publicised 'Eternal Blue' type attacks, where highly advanced government level hacking tools were leaked and then used widely to attack the private sector and critical infrastructure. Invisiron® provided rules for these attacks as soon as they became public and immediately distributed these new rules through the cyber threat intelligence feed. As a result, all Invisiron® cyber protection devices were immunised and provided client networks protection within hours of the first attack being detected, all without requiring any intervention by the end-user.

The configuration and rule writing GUI tools allow for rules to be exported out from and imported into Invisiron®'s Cyber Defence platforms. This means custom rules can be shared between platforms. Rule files are generated in JSON format which makes them easy to read and understand.

REPUTATION DETECTION ENGINE

While the DPI engine protects against the contents of a packet's payload, the reputation engine provides an additional layer of protection by validating the packet source and destination addresses against lists of known malicious IP addresses. It also checks all domain names and Universal Resource Locators (URLs) against similar lists of malicious domains and URLs. Invisiron®'s Cyber Defence platforms contain a database of known malicious domains, URLs, and IP addresses. The database

typically includes several hundreds of thousands of IP addresses and several tens of thousands of malicious domains and URLs. This database is provided to the security platforms through Invisiron®'s cyber intelligence feed located in the cloud. The cloud server harvests real-time threat information from various sources, this information is then vetted and converted into a unified format that is suitable for sending to the platforms. Each active platform will poll the cloud server and within the hour update its database with the latest reputation lists to be protected against new and emerging threats.

It is important to note that the reputation engine is only as good as the intelligence that informs this engine about the reputation of domains, IP's, and URL's. Invisiron®'s reputation engine is informed by high-quality cyber indications and warning (I&W) data from validated and trusted independent sources (e.g. research firms, associations, trusted agents, government contacts, etc.), together with Invisiron®'s own original cyber threat intelligence. This data is fully reviewed, validated, and integrated into the platforms.

The reputation detection engine informed by Invisiron®'s cyber threat intelligence feed checks each incoming and outgoing packet against this database. If a hit is detected, meaning a packet is very likely to be malicious and unwanted, the engine will cause this packet to be dropped. Malicious packets must be dropped, this is if prime importance as this way they will not be allowed to cause any harm in the network.

EVIDENCE COLLECTION AND LOGGING ENGINE

This engine is responsible for collecting information relating to threats that the Cyber Defence platform detects. It generates various log files, both in text format and in Packet Capture (PCAP) format. This makes it easy to open and study the files with third party tools such as Wireshark. Log files are easily downloaded from the device. All log files are rotated to make sure they don't grow too big and also to allow for identifying which date each log file was generated. Log files can also be archived enabling the keeping records as far back as the customer requires. Additionally Invisiron®'s systems have the ability to perform a full packet capture regardless of threat such as during an advanced penetration test or network audit.

INLINE PERFORMANCE

One of the issues when installing a device inline is possible performance degradation of the Internet traffic to and from the internal organisation's network. This is the case with many legacy inline products. Invisiron®'s modern Cyber Defence platforms have been architected and designed in a unique way to allow them to operate inline without causing any noticeable performance degradation at all. To accomplish this goal, Invisiron®'s platforms eliminate the use of the traditional TCP/IP network stacks in the critical path. The high-speed ethernet ports of our Cyber Defence platforms do not require an IP address to be assigned or a Media Access Control (MAC) address to be exposed. Instead, network

packets flow through the platform in stealth mode, while the defence platform performs its advanced security operations directly on each packet inside the transmit and receive buffers of the high-speed network Interface controller card (NIC). The platform remains completely invisible in the network and even with sophisticated detection tools such as NMAP, it is not possible to detect its presence.

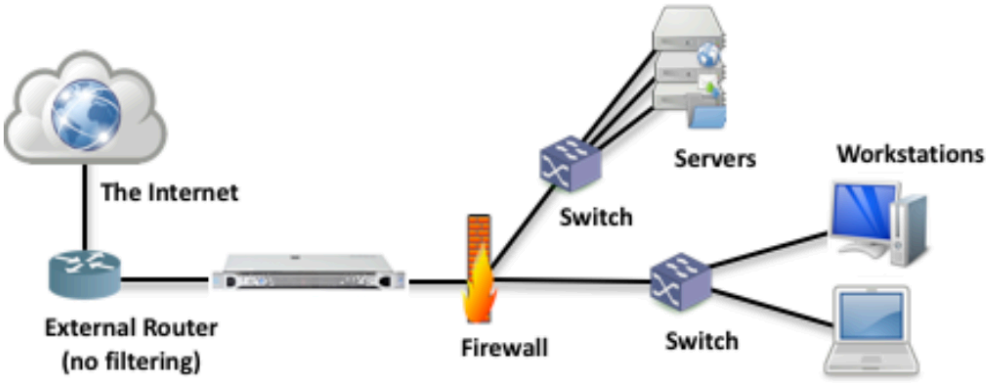

Typical Invisiron® Cyber Defence platform deployment is installed behind the Internet gateway

Instead of using traditional network stacks together with an operating system (OS) such as Linux, Invisiron®'s devices use proprietary software that operates directly on top of the server platform hardware. This software owns and controls the high-speed NIC that handles the protected network traffic. Each packet that arrives at the NIC is delivered to Invisiron®'s security software directly from the receive buffers on the NIC. The software will inspect the packet and perform security checks on the packet before it is let back out on the transmit buffer.

The security analytics are performed by a few independent software processes inside the platform. The most important processes, as aforementioned, are:

- A deep packet inspection engine, or DPI engine.
- A reputation detection engine.
- An evidence collecting and logging engine.

These are independent processes that perform their tasks in parallel with the other processes.

FAULT TOLERANCE

It is a common misconception that installing a security device inline and directly after the Internet router increases the risk of failures and can lead to being cut off from Internet access. Another common misconception is that a security device installed inline must also include bypass circuitry that makes it possible for Internet traffic to pass through the device should it fails, or even if the power is turned off.

It should be noted that the likelihood of being cut off from Internet access is present when a security device is installed between the Internet router and the firewall or main switch. However, this issue exists with all other key components of network operations such as the router, firewall, and core switches. Unless all of these devices are duplicated in a mirrored high availability way, there is always a risk of a single point of failure. This is not a new risk introduced by the security device, it already exists with the other components of the network.

For high availability networks where this is a concern Invisiron®'s Cyber Defence platforms can be mirrored, for this approach to make sense all other devices in the critical path should also be duplicated and mirrored.

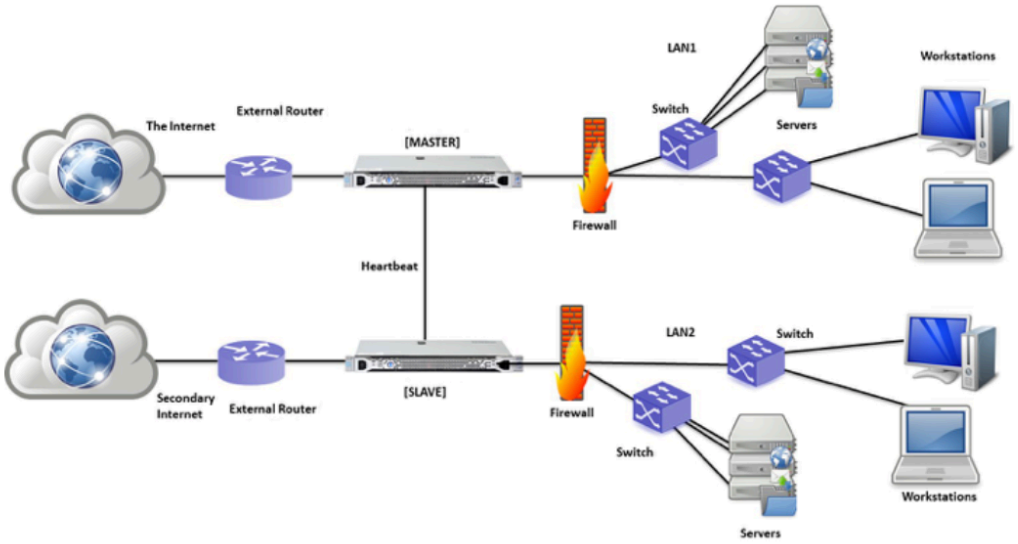

Invisiron® High Availability Configuration using OB-4000 Cyber Defence Platforms

Some Invisiron®'s models include built-in high availability (HA) support. When configured in HA mode, one platform becomes the master platform and the other becomes the slave platform. The two platforms communicate with each other to determine the pair's functional status. Should the master fail, the slave immediately takes over. This happens transparently and does not require user intervention. Once the failed platform is brought back online, the system becomes the master again, provided it was the master when it failed.

Some regions and countries struggle with less than ideal Internet services. In these situations, it might be a good compromise to add fault tolerance to the Internet connection only. This will reduce cost but still provide good protection. To support this type of fault-tolerant internet connection two independent Internet connections can be handled by a single Invisiron® platform.

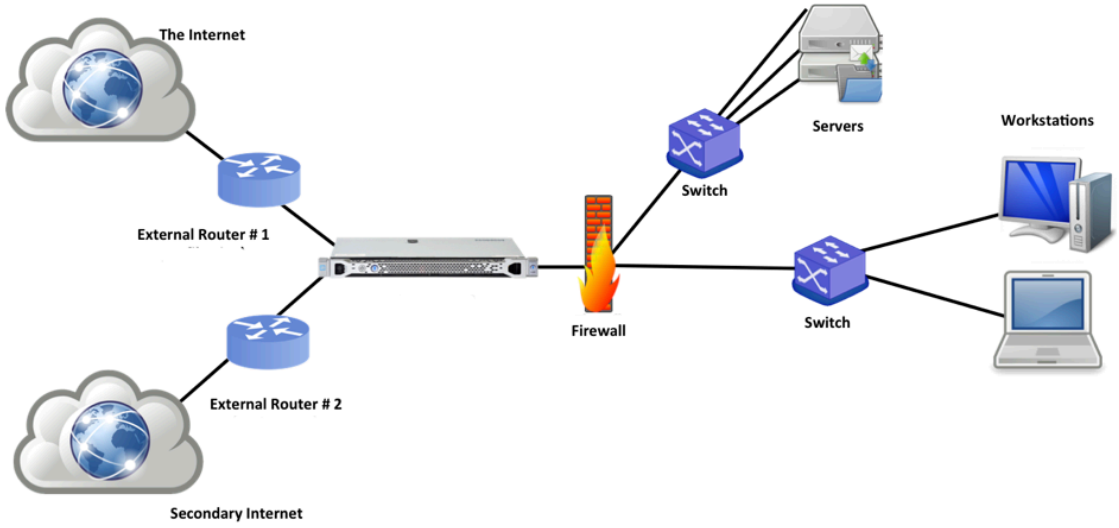

Dual redundant Internet connections configurations using one OB-4000

If the primary Internet connection goes down, the Invisiron® platform will automatically detect this and switch over to the secondary connection. Once the primary connection resumes functionality, it will take over as the primary source for Internet traffic again.

PORT BRIDGING OR BYPASS FUNCTIONALITY

Some solutions include an automatic port bridging functionality. If the device fails or loses power, the bridging kicks in and it will still be possible to reach the Internet, no traffic will be interrupted. In modern defence systems, a security device should never be allowed to perform such bridging as this would open up the attack path towards the network and an attacker is free to penetrate the network with little effort.

Tests in Invisiron®'s security lab, as well as by independent third-parties, have demonstrated that, in the current threat landscape, if a network is left unprotected with various network ports open, it takes less than a minute before scripted attack scanners detect this. Hackers run sophisticated scripts that constantly scan the Internet for vulnerable systems. When found, an unprotected system may be immediately attacked. Therefore today, it is extremely high risk to use products that allow such

bypassing if they fail. Sophisticated vendors such as service providers never allow such a setup in their networks. The correct way to protect against device failures is to use the high availability techniques described earlier in this paper.

REMOTE MONITORING

Invisiron® offers a new and modern remote management tool 'Threat Commander'. This software tool is designed to perform remote monitoring of multiple Invisiron® platforms in real-time. It can replace legacy Security Information and Event Management (SIEM) style tools and focuses exclusively on real-time monitoring of security events and other events in the platforms.

Threat Commander is typically deployed on a virtual server on-premise or in the cloud. Each security platform that needs to be monitored is configured to point towards the virtual server. Once the remote monitoring is activated, the platform will start pushing up all its security-related data to Threat Commander.

Threat commander provides the ability to manage multiple devices and drill down into the details of a platform. This way it is possible to get detailed information about all security events in the security platforms for both current and historical events. The user can also download various log files from the platform such as the system log, event logs, dropped packet logs, and more.

Threat Commander also offers several forensic or analytical functions to generate statistics and analyse security events in detail. This provides users with the ability to make informed decisions and take action, particularly against the more serious and sophisticated attackers. Statistics & data can be exported into comma-separated files which can be imported into other applications, for example, Excel for further processing or report generation.

Invisiron® Cyber Defence platforms also have support for integrating with 3rd party SIEM software tools though it is strongly recommended to use Threat Commander as a targeted solution to the protection provided by our platforms.

COMPLEMENTING AN INVISIRON® TECHNOLOGY BASED DEFENCE

Invisiron® Cyber Defence platforms will, at best, be able to detect and stop approximately 80-90% of attack attempts faced by organisations today. The remaining 10-20% of attacks may result in gaining a foothold in the network unless the defence platform is complemented by other solutions. One example is an employee that brings in their laptop or personally owned smart-phone to the organisation and connects it to the network. This device might have been infected before it was introduced into the

network. This could allow malware to enter the network and cause harm without being detected and blocked by the endpoint device. USB sticks are another source of infection, cases have been found in the past where even brand new USB sticks straight from the manufacturer have been infected by malware.

Depending on the level of security an organisation requires complementary endpoint or cyber protection solutions should be installed.

invisiron®

Invisiron Pte Ltd

1 Pemimpin Drive #08-03
One Pemimpin
Singapore 576151

info@invisiron.com