Evaluating enterprise firewalls

What IT professionals should look for
in their next purchasing cycle

wanstor

Contents

Introduction

Networks are more complex than ever before, with large amounts of data going across them, employees accessing applications they want and need (even if not approved), end users using work or personal devices to access corporate data and of course the ever present IT security threat from outsiders.

To combat security threats many IT teams have simply been adding layers to their IT security infrastructure and buying security solutions which promise to combat different threats. The reality is that for many businesses their IT security infrastructure is now overly complex.

In many cases the existing IT security infrastructure is actually limiting or slowing down the IT team's ability to respond to cybersecurity challenges.

When things become overly complex Wanstor has one piece of advice – Go back to basics and define what you actually need for safe business and IT operations to take place.

Remember IT should enable the business not hinder it. When increasing complexity it usually limits or slows the business or the IT department in their day to day operations.

Wanstor believes it is always helpful to focus on the basics especially when it comes to IT security and firewalls.

IT teams need to remember that although firewalls are not the most exciting of IT topics to address, they are a must-have, because:

+ The firewall is a strategic part of your IT infrastructure – without it your IT infrastructure is vulnerable to a wide range of security threats and will fail to function

+ It acts as the access control point for all traffic – allowing or denying traffic into the network-based on policy

+ It helps to eliminate risks of the unknown by using a positive control model, which allows what you want and nothing else

Over time, firewalls fundamental functions have been nullified by the very traffic they were meant to control. In summary, applications have evolved at a faster pace than the humble firewall, and as a result, existing firewalls have trouble exerting the control they require in order to protect digital assets.

Additionally employees are using these applications to get their jobs done. Some examples of applications and threats found on many businesses networks include:

- **End-user apps:** These applications include social media, file sharing, video, instant messaging and email. Employees may use some of them for work purposes; others will be for personal use. These applications are often highly extensible and include features that introduce unwarranted risk. Usually they represent both business and security risks and the challenge as the IT team is how to balance what applications they allow vs what they block.

- **Core business apps:** These are the applications that run your business. They include databases, directories and ERP applications in your data centre, and applications in the cloud. This group of applications is your most valuable and is quite often the focus of attack from cyberattackers. The IT team need to work out how to protect these applications from stealth and evasion attacks through hackers who try to evade company firewalls.

- **Infrastructure and custom apps:** This group of applications represents core infrastructure applications, such as SSL, SSH and DNS, as well as internally developed, custom or unknown applications. These applications are commonly used to mask command and control traffic generated by bots and other types of malware.

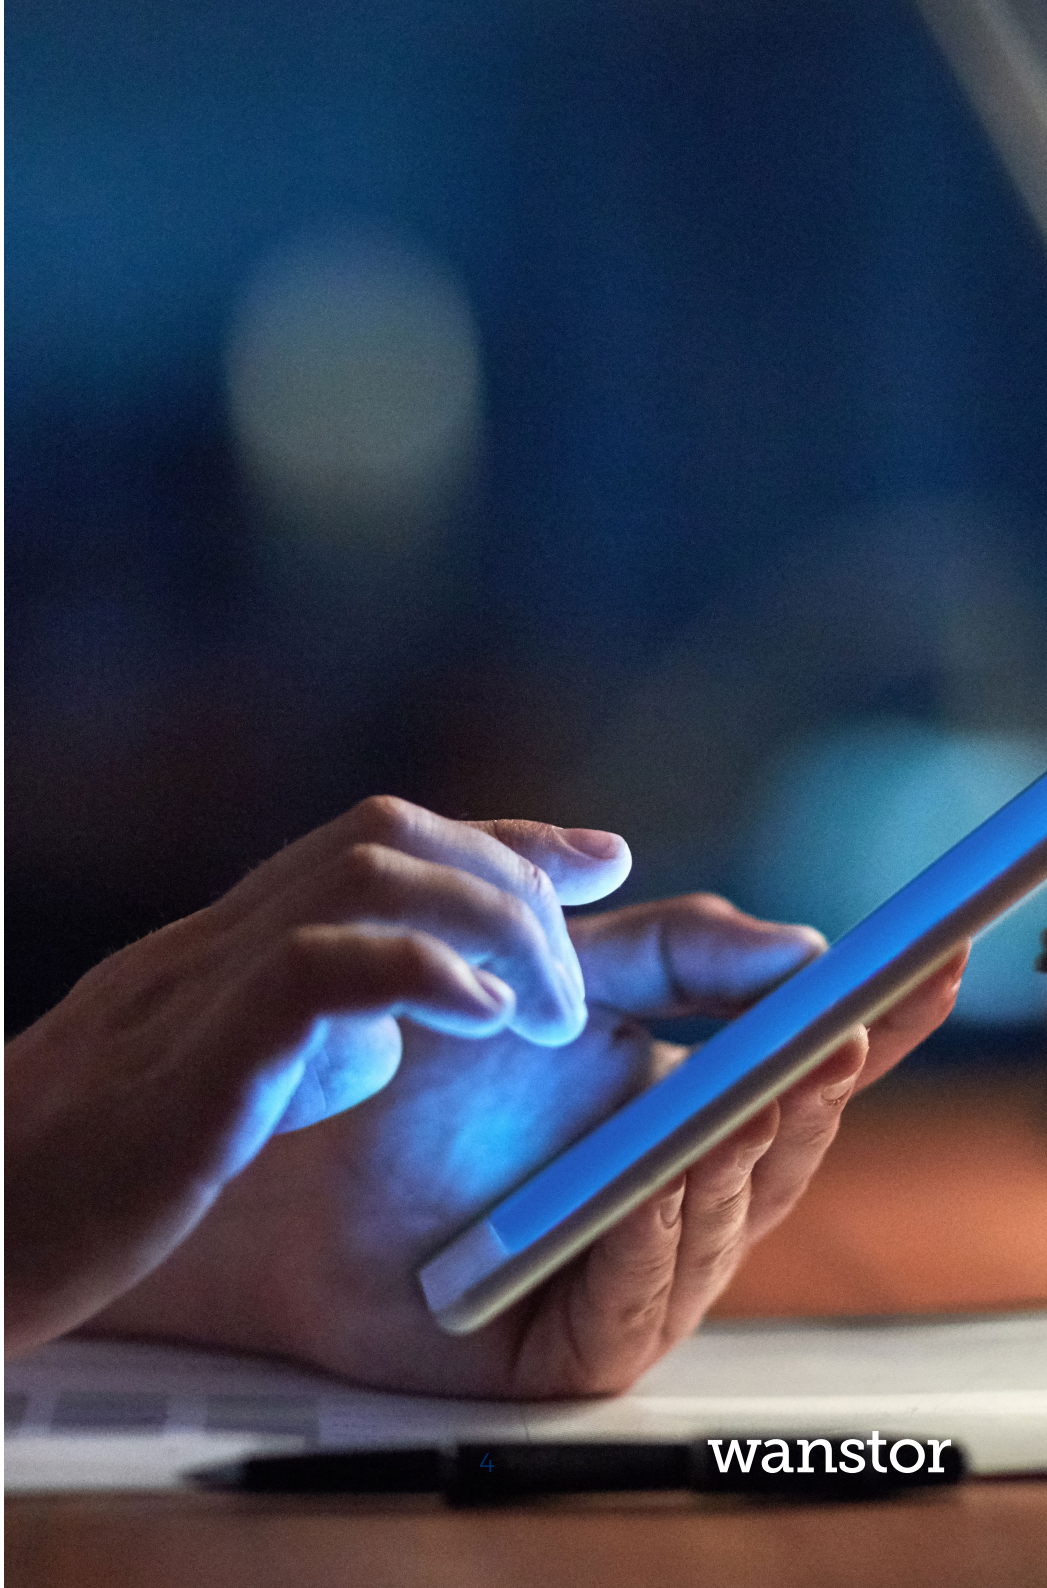

Many of these applications are using a wide range of non-standard ports. To address these challenges, IT teams need to focus on the fundamentals of firewalls.

Over the past two years, almost every network firewall vendor has been re-thinking how they identify and control traffic based on the application itself, instead of just the port and protocol.

This has led to many vendors calling their latest offerings '*next-generation*' with virtually every firewall vendor acknowledging that application control is an increasingly critical part of network security.

At Wanstor we believe there are two key reasons for this renewed focus on the fundamentals by firewall vendors. Firstly, applications and the associated threats can easily bypass port-based firewalls.

Secondly, the firewall is the only place where all the traffic flowing across your network is seen, and it is still the most logical location to enforce access-control policies.

By focussing on the fundamentals of firewalls, Wanstor believes IT teams can significantly improve their security posture, whilst the administrative effort associated with firewall management and incident response should become less of a burden.

Characteristics of next generation firewalls

- *Identify applications regardless of port, protocol, evasive tactic or decryption*
- *Identify users regardless of device or IP address*
- *Decrypt encrypted traffic*
- *Protect in real-time against known and unknown threats embedded across applications*
- *Deliver predictable, multi-gigabit inline deployment*

Purchasing selection criteria for next generation firewalls

At Wanstor we believe in keeping things as simple as possible when it comes to purchasing solutions for our customer's IT environments. In terms of firewalls we believe firewall selection criteria should be across three areas: security functions, operations and performance. Whilst each business will have varied requirements and priorities within the three selection criteria, fundamentally all IT departments must consider the following when thinking about their next firewall solution purchase:

The firewall must identify and control applications and application functions on all ports, at all times

Application developers no longer follow standard port/protocol/application development methods. More and more apps are capable of operating on non-standard ports or can hop ports *(e.g. instant messaging applications, peer-to-peer file sharing, VoIP).*

Additionally, users are increasingly technology conscious enough to force applications to run over non-standard ports. In order to enforce application specific firewall policies where ports are increasingly irrelevant, IT departments must make sure that their next firewall must assume that any application can run on any port.

The concept of any application on any port is one of the fundamental changes in the application landscape that is driving the migration from port-based firewalls to next-generation firewalls. Any application on any port also highlights why a negative control model cannot solve the problem.

If an application can move to any port, a product based on negative control would require previous knowledge or have to run all signatures on all ports, all the time.

At Wanstor we recommend IT teams assume that any application can run on any port, and their next firewall purchase must classify traffic by application on all ports all the time, by default.

Otherwise, port-based controls will continue to be outwitted by the same techniques that have plagued them for years.

Identify and control

A small number of the applications on a business network may be used to purposely evade the security policies that have been put in place to protect a business IT infrastructure assets. Two classes of applications fall into the security evasion tools:

+ External proxies and non-VPN related encrypted tunnel applications are specifically used to evade the in-place security controls using a range of circumvention techniques. These applications have no business value to a network as they are designed to evade security, introducing unseen business and security risks.

+ Remote server / desktop management tools, such as RDP and TeamViewer, are typically used by support and IT professionals to work more efficiently. They are also frequently used by employees to bypass the firewall, establishing connections to their home or other computer outside of the network. Cyber attackers know these applications are commonly used.

+ As with everything in IT security, no two attacks are the same and therefore, not all of these applications carry the same risks for example; remote access applications have legitimate uses, as do many encrypted tunnel applications.

However, these same tools are increasingly being adopted by attackers as part of their ongoing persistent attacks. Without the ability to control these security evasion tools, businesses cannot enforce their security policies, exposing themselves to risks they thought their controls mitigated.

At Wanstor we believe there are different types of circumvention applications, each using slightly different techniques. There are both public and private external proxies that can use both HTTP and HTTPS. Private proxies are often set up on unclassified IP addresses with applications such as PHProxy or CGIProxy.

Remote access applications such as RDP, TeamViewer or GoToMyPC have legitimate uses, but due to the associated risk, should be managed more closely.

Most other circumventors have no business use case to be on a corporate network. Regardless of your businesses existing security policy stance, the next firewall needs to have specific techniques to identify and control all of these applications, regardless of port, protocol, encryption, or other evasive tactic. Additionally applications that enable circumvention are regularly updated to make them harder to detect and control.

The IT team need to understand it is important that the next firewall technology solution they purchase should identify these circumvention applications; it is also important to know how often that firewall's application intelligence is updated and maintained.

Decrypt and inspect SSL and control SSH

Given the increasing adoption of HTTPS for many high-risk, high-reward applications that end-users employ *(e.g., Gmail, Facebook)*, and the user's ability to force SSL on many websites, many IT security teams have a large and growing area that they do not know about without the ability to decrypt, classify, control, and scan SSL-encrypted traffic.

A next-generation firewall must be flexible enough that certain types of SSL-encrypted traffic can get through, while other types can be decrypted via policy. SSH is used universally and can be easily configured by end users for non-work purposes in the same way that a remote desktop tool is used.

The fact that SSH is encrypted also makes it a useful tool to hide non-work related activity.The ability to decrypt SSL is a basic element and not just because it's an increasingly significant percentage of business traffic, but also because it enables other key features that would end up ineffective without the ability to decrypt SSL.

Key elements to look for include recognition and decryption of SSL on any port, inbound and outbound; policy control over decryption, and the necessary hardware and software elements to perform SSL decryption across tens of thousands of simultaneous SSL connections with predictable performance.

IT Managers should also consider the ability to identify and control the use of SSH. Specifically, SSH control should include the ability to determine if it is being used for port forwarding (local, remote, X11) or native use (SCP, SFTP and shell access). Knowledge of how SSH is being used can then be translated into appropriate security policies.

Application function control

Application platform developers such as Facebook, Google and Microsoft provide users with a rich set of features and functions that make users loyal to their applications but expose them to a range of risks. E.g. Allowing WebEx is a valuable business tool, but using WebEx Desktop Sharing to take over your employees' desktop from an external source may be an internal or regulatory compliance violation.

Another example may be Gmail and Gtalk. Once a user is signed into Gmail, which may be allowed by policy, they can easily switch to Gtalk, which may not be allowed through security policies.

A *"next gen"* firewall must be able to recognise and define individual features and functions so that an appropriate policy response can be implemented.

At Wanstor we believe a firewall must continually classify each application, monitoring for changes that may indicate when a different function is being used.

The concept of *"once and done"* traffic classification is no longer an option as it ignores the fact that these commonly used applications share sessions and support multiple functions.

If a different function or feature is introduced in the session, the firewall must make a note of it within the state tables and perform a policy check.

Continual state tracking to understand the different functions that each application may support, and the different associated risks, is a critical requirement for the next firewall purchase.

Systematically manage unknown traffic

Unknown traffic exists in small amounts on every network, yet to your business, it represents a range of significant risks. There are several important elements to consider with unknown traffic.

Is it categorised?

Can you minimise it through policy control? Can your firewall characterise custom applications so they are “known” within your security policy? Does your firewall help you determine if the unknown traffic is a threat?

Unknown traffic is also strongly tied to threats in the network. Attackers are often forced to modify a protocol in order to exploit a target application. For example, to attack a web server, an attacker may need to modify the HTTP header so much that the resulting traffic is no longer identified as web traffic.

Such an anomaly can be an early indicator of an attack. Also, malware will often use customised protocols as part of its command and control model, enabling security teams to identify and take action against any unknown malware infections.

By default, a firewall must classify all traffic on all ports. Positive models classify everything; negative models classify only what they’re told to classify. Classifying everything is only a small part of the challenge that unknown traffic introduces.

A firewall must give IT Managers the ability to see all unknown traffic, on all ports, in one management location and quickly analyse the traffic to determine if it is:

+ An internal or custom application
+ A commercial application without a signature
+ A security threat

Additionally, a firewall must provide IT Managers with the necessary tools to not only see the unknown traffic but to systematically manage it by controlling it via relevant security policies, creating a custom signature, submitting a commercial application PCAP for further analysis, or performing a forensic investigation to determine if it is a threat.

Protect your network from known and unknown threats across all apps and ports

Businesses continue to adopt a wide range of applications to enable their employees to do their jobs and quite often these applications are hosted in a mixture of places inside and outside of the business's physical location.

Whether it's hosted by SharePoint, Google Docs, or Office 365, or an extranet application hosted by a partner, many businesses require the use of an application that may use non-standard ports, SSL or can share files.

In summary, these applications may enable the business, but they can also act as a cyber threat path. Some of these applications (e.g. SharePoint) rely on supporting technologies that are regular targets for exploits *(e.g. IIS, SQL Server)*.

Blocking the application isn't appropriate, but neither is blindly allowing the applications and the associated business and cybersecurity risks.

This tendency to use non-standard ports is highly emphasised in the world of malware.

Since malware resides in the network, and most communication involves a malicious client *(malware)* communicating to a malicious server *(command and control)*, then the attacker has full freedom to use any port and protocol combination they choose.

Part of safe enablement is allowing an application to come into the business security perimeter and scanning it for threats. These applications can communicate over a combination of protocols, and requires a more sophisticated firewall policy than “block the application.”

The first step is to identify the application, regardless of port or encryption, then determine the functions IT can allow or deny, then scan the allowed components for any of the appropriate threats e.g. malware, spyware, or even confidential, regulated, or sensitive information.

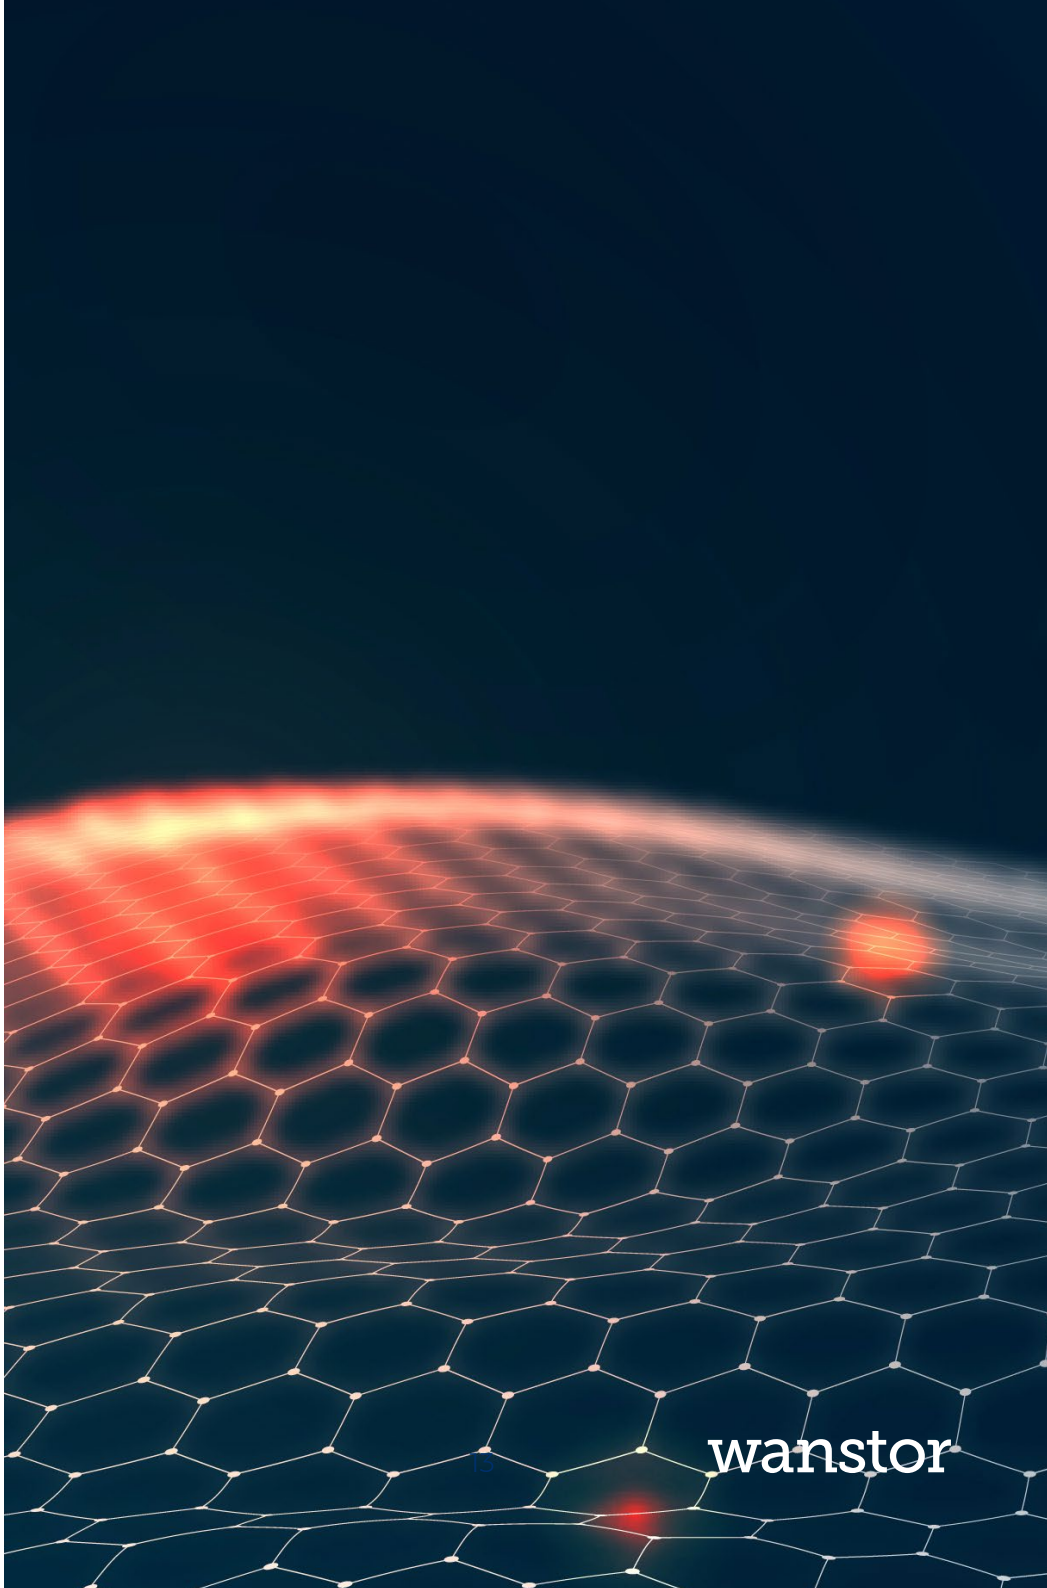

Deliver consistent controls to all user groups, regardless of location or device type

Business users are increasingly working from locations which are not a fixed company location. For example you regularly see people working in coffee shops, at train stations, motorway service stations and airports all the time.

More often than not users who are representing the company and are mobile in their workstyle need to access the corporate network on smartphones, tablets or laptops.

Whether working from a coffee shop, home or a customer site, your users expect to connect to their applications via Wi-Fi.

Regardless of where the user is, or even where the application being employed might be, the same standard of firewall control should apply.

A firewall enables application visibility and control over traffic inside the business, but not always outside and this leaves users vulnerable to accessing risky websites and thus in turn exposing the business to risky traffic.

As a concept to get over this challenge it is simple, a firewall must have consistent visibility and control over traffic, regardless of where the user is.

This is not to say that your business will have the same policy for both; for example, some businesses might want employees to use Skype or Jabber when on the road, but not inside headquarters, where others might have a policy that states users may not download Salesforce.com attachments unless they have hard-disk encryption turned on.

This should be achievable on a “next gen” firewall without introducing significant latency for the end user, undue operational hassle for the administrator, or significant cost for the business.

The firewall must simplify network security with the addition of application control

Many businesses struggle with incorporating more information feeds, more policies, and more management into already overworked security processes and people.

If your team can't manage what it's already got, adding more devices, managing interfaces along with associated policies and information will not help reduce your team's administrative effort nor will it help reduce incident response time.

Given that typical port-based firewall installations have thousands of rules, adding thousands of application signatures across tens of thousands of ports is going to increase complexity by several orders of magnitude.

Your business is based on applications, users and content, and the next firewall purchased must allow the IT team to build policies that directly support business initiatives.

Shared context across the application, user, and content in all aspects – visibility, policy control, logging and reporting – will help the IT team to simplify the businesses security infrastructure significantly.

A firewall policy based on port and IP address, followed by separate policies for application control, IPS and anti-malware will only complicate policy management processes and will in many cases end up inhibiting the business.

Deliver the same throughput and performance with application control

Many businesses struggle with the forced compromise between IT performance and security requirements. In many cases, turning up security features on the firewall means accepting significantly lower throughput and performance from IT infrastructure, devices and users. Designing and building the next firewall you purchase is therefore crucial to success.

At Wanstor we recommend IT Managers always start with the objectives of the business then build a security solution that enables the right levels of user access to applications not the other way round.

Otherwise IT teams will discover they have a firewall which may protect IT infrastructure but does not actually enable the business to achieve its objectives. Therefore the importance of architecture is clear. Putting together a port-based firewall and other security functions from different technology origins usually means there are redundant networking layers, scanning engines and policies resulting in poor IT performance.

From a software perspective, the firewall must be designed to do this from the beginning. Additionally, given the requirement for computationally intensive tasks performed on high traffic volumes and with the low tolerance for latency associated with critical infrastructure, a firewall must have hardware designed for the task as well, meaning dedicated, and specific processing for networking, security and content scanning.

Delivering the same firewall functions in both a hardware and virtualized form factor

The unmitigated growth of virtualization and cloud computing introduces new security challenges that are difficult or impossible for legacy firewalls to manage effectively. This is due to inconsistent functionality, disparate management, and a lack of integration points with the virtualization environment.

In order to protect traffic flowing in and out of the data centre within virtualized environments and in the public cloud, firewalls must support the same functionality in both a hardware and virtualized form factor. The dynamic setup and tear down of applications within a virtualized data centre exacerbates the challenges of identifying and controlling applications using a port - and IP address - centric approach.

In addition to delivering the features already described in both hardware and virtualized form factors, it is imperative that the next firewall solution the IT team purchases provides in depth integration with the virtualization environment to streamline the creation of application-centric policies as new virtual machines and applications are established and taken down.

This is the only way to ensure the IT team can support evolving data centre architectures with operational flexibility while addressing risk and compliance requirements.

Final Thoughts

Your users continue to adopt new applications and technologies in order to be productive in their jobs, but with little regard to the associated business and security risks. In some cases, if the IT team blocks these applications, it may hinder the business they work for.

Applications help employees be productive in their job roles and maintain productivity in the face of competing priorities. Because of this, safe application enablement is increasingly the correct policy stance.

To safely enable applications and technologies on your network and the business that rides atop them, network security teams need to put in place the appropriate policies governing use, and also the controls capable of enforcing them.

At one time, the concept of allowing an employee to use an external or personal application for work-related purposes was unheard of.

Today, employees are always online and are continually using the latest applications, often mixing personal and work-related usage. Summarily blocking these applications is equivalent to blocking the business.

The areas we have discussed in this whitepaper about what your next firewall should do validates the fact that the best location to execute secure application enablement is at the firewall.

This can be done by using the application identity and traditional positive control model (firewall) policies that allow IT teams to define, based on the business, which applications are enabled and which are denied.

For more information about Wanstor's IT Security Solutions, please contact us on **0333 123 0360**, email us at **info@wanstor.com** or visit us at **www.wanstor.com**.

Wanstor | 124-126 Borough High Street | London | SE1 1LB | info@wanstor.com

wanstor