Mesh Security

The Next Evolution of Data Security

A White Paper by Carl Gottlieb, 18th July 2011

Contents

Version Control

Date	Version	Description	Author	Comments
18th July 2011	1.1	Official Release	Carl Gottlieb	Official Public Release

Glossary of Terms

AV	Anti-Virus
Firewall	Network based security device designed to control network traffic based on pre-determined rules based on elements such as source, destination and service
HA	High Availability
IPS	Network based security device designed to control network traffic based on pre-determined rules based on malicious activity
LAN	Local Area Network
Protected Assets	Assets such as data, applications, physical and logical servers, appliances and devices that the organisation chooses to protect via the Mesh Security Policy
UTM	Unified Threat Management
WAN	Wide Area Network

Overview

Data Security strategy is currently stuck within the confines of the network. Organisations continue to protect their data the way they always have done, with layers of network devices to segment their architecture into secure enclaves of assets. Over the last five years, this strategy has become less sufficient with perimeters dissolving; internal users no longer trusted as they once were and cost becoming king. Perimeter Security was the first phase in providing data security. The next evolutionary step is Mesh Security.

Mesh Security is the concept of layering security over the application or asset itself, rather than on a network boundary, providing point-to-point security between any two entities. This concept allows for the flattening of network architectures where network firewalls and ACLs (Access Control Lists) have been previously used to cordon off networks from each other. Mesh Security is focussed on protecting the asset and not the network.

Background

Since its first conception in 1988, firewall technology has been the primary method of providing data security within a network. Dedicated firewall devices now reside in almost every network environment and are still an accepted method for providing data security through inspection of source, destination service and content. Controls such as network IPS (Intrusion Prevention System) and application proxies focus more on content security than network arrangement, and more recently have dovetailed with firewalls to form UTM (Unified Threat Management) devices. Critically these systems are solely focussed on the network traffic traversing them and not the actual application or data the business is trying to protect. These are boundary systems providing Perimeter Security.

As data security awareness has grown in the last twenty years, organisations have initially focussed on data integrity, keeping networks separated with strong barriers and perimeters. Attention to data availability then started to grow as live systems became more critical and SLAs became the norm. This necessitated a shift towards greater resiliency across the network with HA (High Availability) pairs or clusters of devices, networks and data centres now the norm. Technically this resiliency requires additional hardware and necessarily additional costs to the organisation, to essentially provide the same functional service but in a more robust fashion.

Confidentiality of data has always been at the forefront of the minds of security personnel, but traditionally has been an easy problem to solve. Simply locking the data in a huge mainframe, accessed by only one authorised administrator negated the problem, but unfortunately this is no longer an option. Today's systems host data in numerous geographic locations, data is constantly in motion and use; and is accessed by a multitude of people. The vast majority of network architectures have grown but not necessarily evolved to handle this change and still focus on simple perimeter controls at each location where the data is accessed and hosted.

The Data Security Challenge

Today's enterprise environments have similar challenges to that of a modern prison. Prisons are fairly simple organisations, existing to manage its prisoners, all under the watchful eye of the Prison Governor. The Governor's responsibilities include securely detaining the prisoners whilst allowing their controlled movement within designated areas and protecting them from harm in what can be a hostile environment.

Whilst the organisation prohibits it, and controls exist to prevent it, contraband does enter the prison via different vectors (visitors, guards, deliveries) and occasionally prisoners do escape. The prisons know this goes on, yet seem powerless to completely eradicate these issues altogether despite spending more and more money on security each year.

On the surface the prison appears to be layer after layer of security to provide an incredibly secure structure. Crucially, all of these controls are actually about securing the prisoners and not the corridors, entrances and gateways. If the Governor could control the prisoners themselves (e.g. through some kind of mind control) then he could strip away all those layers of bars and concrete and let the prisoners operate in a much more open environment.

All organisations are looking to control their data. Be it a database of customer information or a single image file, the same requirements over access and movement of data exist, seeking to control its confidentiality, availability and integrity. As environments grow, so do the threats of malicious attacks and data leakage along with the ever increasing cost of additional security measures. Ultimately the challenge facing security personnel is often not being able to technically control the data, since files themselves are often just raw information, and instead have to wrap security layers around the data.

Currently this securing of data is predominantly achieved through network perimeters together with point solutions aiming to provide security closer to the data itself where possible. Organisations often only utilise host based security products, such as AV, as a secondary line of defence with out-of-the-box blanket policies, finding more granular controls almost impossible to manage with existing toolsets. Figure 1 demonstrates different organisational assets and the security controls that are used to protect them.

Figure 1 – Assets and security layer locations

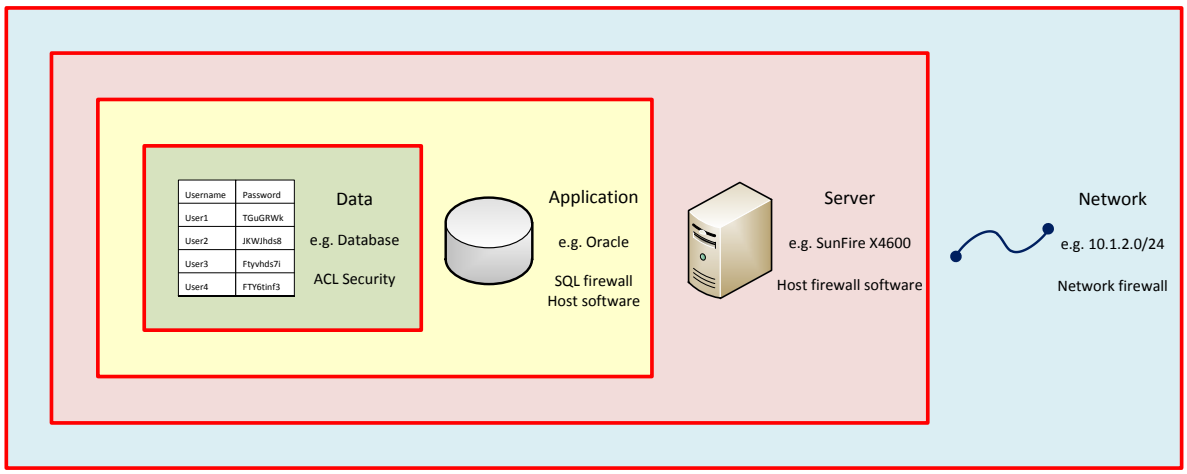

The modern enterprise environment is a relatively secure place but its layers of perimeter security present many challenges:

1. **Access Points Everywhere** – Organisations increasingly need to provide access from multiple locations through a multitude of WAN connections to data resident on virtually any corporate system. The notion of the perimeter is dissolving rendering perimeter security of limited value.
2. **Management of Disparate Security Functions** - Point security products and controls throughout the organisation will commonly be managed by different management solutions resulting in inconsistent security policies with additional cost and complexity of management.
3. **Data Mobility** – Traditionally servers rarely change their IP addresses, allowing for firewall rule sets to remain relatively static. As data moves to more dynamic environments with changeable IP addresses (Cloud hosting, VMware vMotion, DHCP on mobile devices) these controls are no longer effective, resulting in looser controls and increased Change process.
4. **Auditability** – With security controls placed away from the data there is no simple method for demonstrating what data controls exist, who can access what data and how does data securely move from point A to point B.
5. **Change Control of Shared Services** – Network located controls such as firewalls apply protection to numerous servers, applications and data. They are in themselves a form of shared security service. In turn, applying changes to the firewall often requires approval by all service owners whose traffic traverses the device, even if the change itself only relates to other parties' services. Simple changes therefore become costly and time consuming.
6. **Value for Money** – In many organisations 50% of all network devices sit dormant, acting as hot standby solutions in resilient architectures. These devices and their associated management provide zero functional benefit for their cost.
7. **Encryption Everywhere** – Organisations are increasingly encrypting traffic over their WAN and LAN links in an attempt to prevent eavesdropping and interception of the data in motion. Whilst this is a technically good practice, it often renders perimeter security devices blind to the traffic that traverses them. The use of encryption can therefore cause a greater security problem than it is trying to solve.
8. **Scalability** – Deploying physical security boundaries at every potential gateway becomes exponentially more expensive as networks grow with more and more security devices required. As traffic volumes increase, black box security appliances are rarely upgradeable and will need to be replaced, further adding to expense.

The Next Evolution of Security - Mesh Security

As business requirements grow and the security threat landscape evolves so does increased reliance on perimeter security. The concept of Mesh Security is by no means the panacea to enterprise security but is the next stepping stone towards securing the data itself and not just the perimeter.

Mesh Security is based around three essential components:

1. Single centralised management of the whole security mesh.
2. Security enforcement resides on the protected asset itself.
3. Security policies are based on unique identifiers of Protected Assets and not specifically IP addresses.

Technically the Mesh Security model is achieved through the use of host based security rules and enforcement on each Protected Asset. This ensures that however and whoever is accessing the data, the required controls will always be enforced. An example would be enforcing the Security Policy on the virtual server or the web application through the use of host installed security software. This Security Policy would extend to every Protected Asset, be it a data file, application, server or black box appliance. The network itself does not need to be protected. In turn Mesh Security allows for the removal of Perimeter Security controls and frees the network to do its job of transporting data with minimal latency.

Figure 2 – Mesh Security Overlay & Point-to-Point Connections between Protected Assets

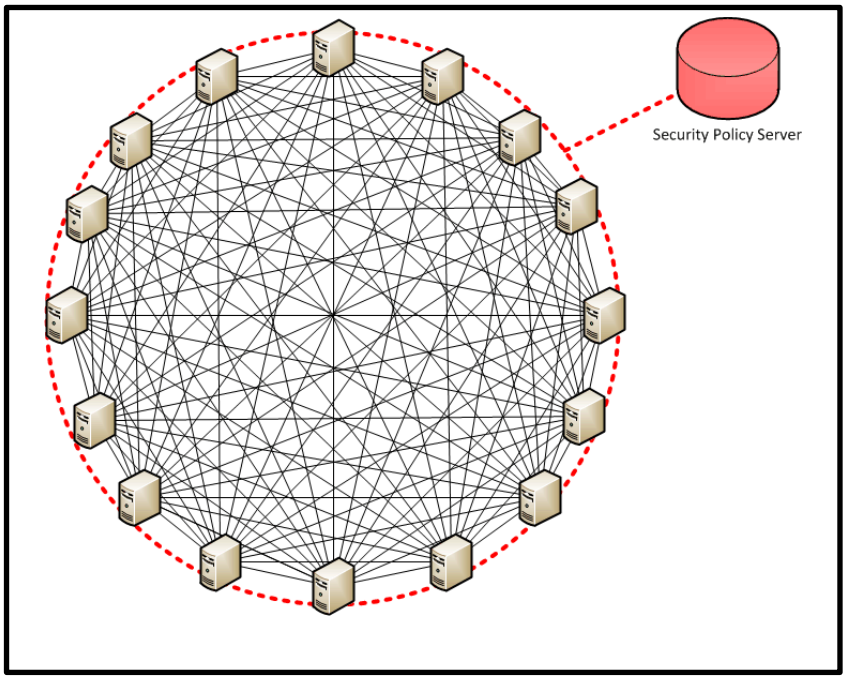

Crucially this mesh model dictates the single enterprise wide security model to be managed by one management entity. This management system contains one organisation Security Policy which is disaggregated into individual policy files when pushed to each of the Protected Assets. Each of these policy files will contain only the rule sets that apply to that specific Asset. Each rule set would specify a point-to-point relationship with another Protected Asset, e.g. WebserverA to DatabaseAppX, with rules within that set specifying regular controls, such as authentication, authorisation, content and allow/deny.

Figure 3 demonstrates the relationship between the organisational Security Policy and individual Policy Files.

Figure 3 – Security Policy & Security Policy File

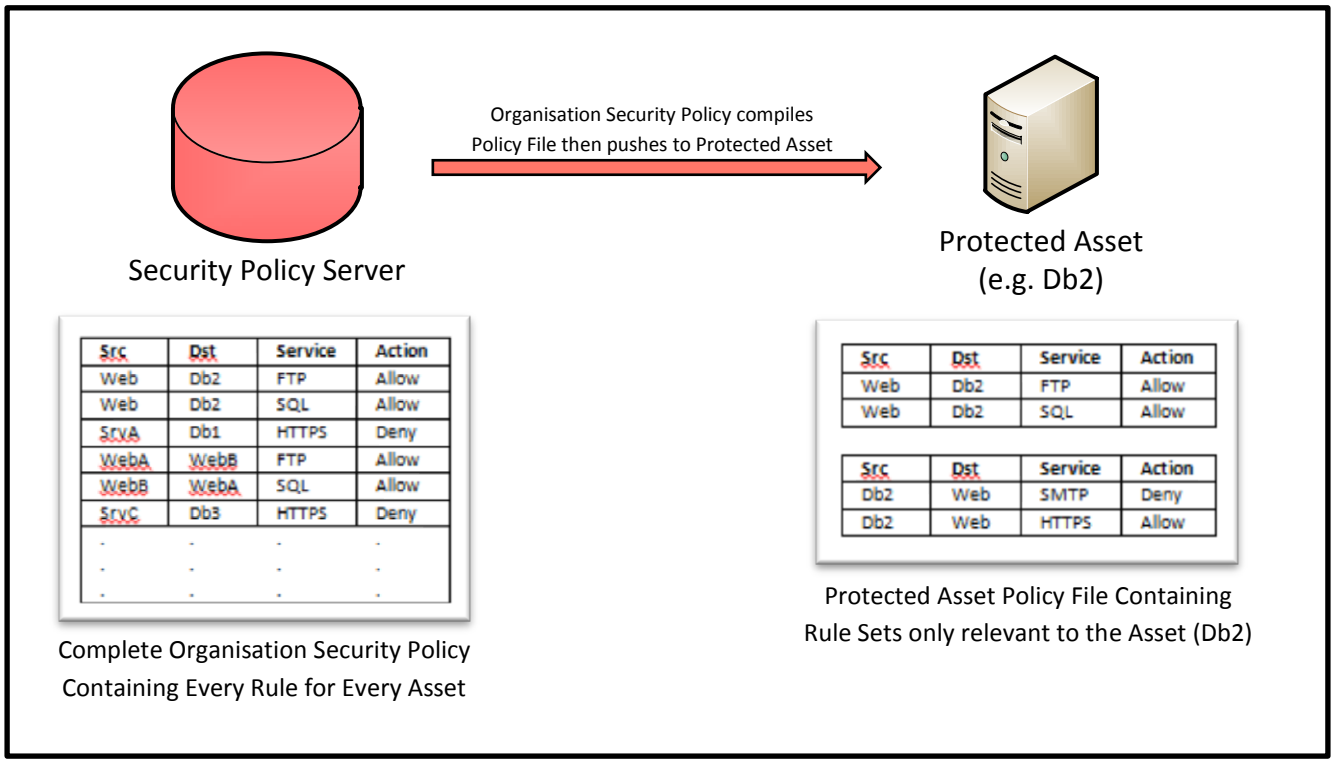

In an IT environment there will only need to be one Security Policy. (Organisations such as service providers could theoretically require more). If for example there were five servers, one single security policy would produce five policy files, one for each of the servers. If each of the five servers communicates with the other four in a two-way fashion then there would be twenty relationships, dictating 20 rule sets. Each of these rule sets could contain numerous rules depending on what security controls are required.

Figure 4 below demonstrates the quantity of Policy entities

Figure 4 – Mesh Security Policy entities

	Formula	Example
# of Protected Assets	n	5
# of Security Policies	1	1
# of Policy Files	=n	5
Maximum # of directional point-to-point relationships (Rule Sets)	$=n^2-n$	20

This paper will now discuss in more detail the three essential components of Mesh Security:

1. Single centralised management of the whole security mesh.
2. Security enforcement resides on the protected asset itself.
3. Security policies are based on unique identifiers of Protected Assets and not specifically IP addresses.

1) **Single centralised management of the whole security mesh.**
   Mesh Security provides one security view of the whole estate of Protected Assets. This provides real time analysis and reporting of the true security posture. E.g. If an audit asks what access is allowed to and from a specific asset, the single security policy will display this. To facilitate the inclusion of assets from various vendors it is ideal for the Mesh Security provider to offer an API or plug-in to connect a third party asset's security functions into the Mesh Security Policy. Clearly the leading server platforms and operating systems such as Linux, Unix and Windows will be the most rapid to incorporate, with platforms such as storage arrays, mainframes and black box appliances requiring more API/plugin type development. Organisations that still require some perimeter controls could simply bolt their management into the Mesh. The goal is for all entities to be manageable and controllable by the single Mesh Security Policy.

2) **Security enforcement resides on the protected asset itself.**
   Mesh Security is focussed on keeping security as close to the data as possible, moving from perimeter/network based controls to host based security. By definition this results in all asset to asset communication traversing two layers of security, one inbound/outbound policy at each side, providing defence in depth through granular least access privilege controls. In turn this point-to-point security does not require any third party security boundary devices such as firewalls or IPS for data confidentiality or integrity.

3) **Security policies are based on unique identifiers of Protected Assets and not specifically IP addresses.**
   Mesh Security is based on the notion of the Protected Asset by name or identifier, which is a fundamental differentiation against the perimeter security reliance on an IP address. This concept is crucial to facilitate more fluid networks in the age of mobile assets, motion of virtual servers and cloud hosting. What defines the asset isn't itself necessarily important, it is the uniqueness and dynamic manageability that is essential, such that if the asset were to appear in a different geographic location (e.g. VMware vMotion to a cloud provider), the security policy would still apply to that asset. Assets can use different identifier types to provide flexibility and granular control, but all can be grouped and templated to streamline the management of assets and their policies.

Benefits of Mesh Security

This section now revisits the challenges of Perimeter Security and analyses them within a Mesh Security framework.

1. **Access Points Everywhere** – Since security is enforced on the Protected Asset itself, the point of access is of no consequence. The controls cannot be bypassed.
2. **Management of Disparate Security Functions** – All security enforcement is managed by one policy in one place providing consistent security policies and simple, cost effective management. Policies based on data/application/server groups and templates allow for even simpler management and stronger control than is currently available.
3. **Data Mobility** – Since policies are based on unique identifiers and not potentially transient IP addresses, assets are free to move around the network whilst maintaining consistent security control.
4. **Auditability** – Single centralised management ensures complete real-time auditability.
5. **Change Control of Shared Services** – Security policies and rules relate only to the Protected Assets which in general will have few service owners, allowing changes to be effected more cheaply and in a more timely fashion.
6. **Value for Money** – Security control is placed on the existing organisational assets, thereby requiring no additional security hardware or resiliency considerations beyond the network fabric. The majority of Perimeter Security controls can be removed.
7. **Encryption Everywhere** – Mesh Security controls are located on the asset behind the encrypting function (pre-encryption, post-decryption or intra-host communication) which therefore allows for complete network encryption without any impact on security visibility.
8. **Scalability** – Host based controls require no additional hardware and have no impact on the estate architecture as traffic requirements grow. Embedding security into hardware and the kernel of the server provides for even greater performance benefits and latency reduction.

Additionally a significant benefit of moving away from perimeter security is the ability to shift to flatter high performance networks such as Juniper's QFabric. All organisations are highly cost conscious and the ability to consolidate and rationalise network infrastructure is highly desirable.

Mesh Security in 2011

At the time of writing this White Paper (July 2011) there is little in the way of complete Mesh Security products in the market place. Microsoft made a good start with Active Directory many years ago producing an architecture of file, folder and group based permissions which is used by almost every organisation globally, but is still only limited to Microsoft files, share, users and products. Trend Micro produces a very interesting product named “Trend Micro Core Protection for Virtual Machines” which protects virtual machines from other assets through the use of host based software controls. Currently the product is limited in its control of just Trend Micro security functions, but its approach to securing a virtual machine as if it were any other physical asset goes some way towards the next generation of security products.

Mesh Security is currently a space that few vendors participate in, commonly due to the difficulty in establishing this framework by oneself, and without a common drive in the marketplace. Notably customers are still cautious over the concept of one solution pushing changes to third party systems, but generally this is due to lack of effective toolsets demonstrating their value. ExaProtect’s SolSoft ChangeManager was such a management tool, utilising the Mesh Security concept of a single organisational Security Policy pushing individually compiled policy files to each managed firewall. Whilst this product had its strengths, it was still managing firewalls and network devices rather than the actual security of data.

The Next Steps

It is evident that organisations face a chicken and egg situation. The tools to manage host-based point-to-point security do not exist and so sticking with network devices is the easy option. Vendors are generally averse to shifting research and development funding away from network security whilst the market is still strong, and so are not developing the new management tools and host based security systems.

Fortunately the current economic climate and consumerisation of IT is pushing organisations towards the adoption of Cloud environments and unmanaged endpoint assets (e.g. iPad), to keep costs down and productivity up. This drive is stretching the capabilities and effectiveness of existing security controls and forcing organisations and vendors to rethink their overall security strategy. EMC have made significant steps in this direction with their VMware VMsafe technology allowing third party vendors to install their security functions into the underlying fabric of the VMware server, allowing stronger and better performing enforcement. Intel’s acquisition of McAfee demonstrated their appetite to embed security controls into the server hardware, again moving the protection towards the data and away from the perimeter.

Technically there are no barriers to all-encompassing security controls being available; it is simply a case of waiting for vendors to invest in the next evolution of data protection, Mesh Security.