



Barbican RNP

Real-time Network Protection
for the Introduction of a new CyberWarfare Defense Layer

Manual & Documentation

Version 1.0



Congratulations on your purchase of Barbican RNP!

You are about to add a new CyberWarfare Defense layer to your infrastructure, thus instantly protecting you from:

- ❑ Penetration Testing probes
- ❑ The effects of a distributed Denial-of-Service (dDoS) Attacks
- ❑ Malware and „Junk“ traffic of other sorts

As soon as you connect Barbican RNP, you will also discover other benefits quickly:

- ❑ Higher Uptime and systems availability
- ❑ Decreased bandwidth costs
- ❑ Reduced processing load on your firewall
- ❑ Reduced log sizes in your firewall and Intrusion Detection Systems (IDS)
- ❑ Detail data on attacks on your infrastructure (if you choose to configure the „syslog“ reporting feature and collect the data sent from each Barbican RNP unit)
- ❑ Fewer compromised PC Desktops in your organization
- ❑ Improved efficiency of your Information Technology staff

The installation of Barbican RNP will not take more than 5 minutes

– **follow the QuickStart Guide.**



Barbican RNP – Manual & Documentation

Page

QuickStart Guide 4 - 8

Setup, Installation & Deployment of Barbican

Deployment Options 9 - 16

Online Hosting 9

Perimeter Defense 12

Internal Networks 14

Background Information 17 - 45

Introduction of the CyberWarfare Defense Layer 17

Is a firewall enough? 22

Barbican RNP Attack Protection 25

Background on dDoS 27

Comparison: previous dDoS Defenses 36

Product Development History 42

TroubleShooting

Hardware Failure? 46

Quick Exchange Procedure

Further information can be found on Melior, Inc.'s web site at www.dDoS.com

For any support questions and to renew your Service & Maintenance contract,
please contact support@dDoS.com or:

Melior, Inc.

CyberWarfare Defense Solutions

1501 Beaumont Street

Columbus A. Langley Building

Dallas, Texas 75215

USA

Telephone: +1-214-421-5975 (or toll-free in the USA: 888-4MELIOR)

Facsimile: +1-214-421-5951 (or toll-free in the USA: 888-TOFAXUS)

QuickStart Guide

Setup, Installation & Deployment of Barbican

Unpack Barbican RNP

Barbican RNP units are shipped in shock-proof packages.

Prior to opening (best at the time of delivery) check the package for damages during transport. If you find any indication of mishandling during shipment, please file a report with the carrier (UPS, FedEx, DHL, etc.).

Carefully open the package, remove the top layers of the packaging material, take the unit out, and remove the protective plastic wrapping.

Check Components

In the package, you should find the following parts:

1x	Barbican RNP GbitE Device v1.0
2x	Rack mount Brackets
4x	Rack mount Screws
1x	Standard Power Cable
1x	Documentation Packet (<i>this Manual</i>)

Please make sure all parts are included. Should you find any part missing, please contact Melior, Inc. at support@dDoS.com


Physical Barbican RNP Setup



If you choose to set up Barbican RNP in a 19" rack, please add the rack-mount brackets to either side of the unit.

Mount the unit in the rack; you may wish to have a second person hold the unit, while you fasten the screws.

Logical Barbican RNP Setup



*CyberWarfare Defense
Port I – Internal Network*

*CyberWarfare Defense
Port II – External Network*

*Admin & Reporting
Management Port*

*Power
Cord*

1. First identify the ideal location in your network for Barbican RNP to be placed. For additional information, see also the section on "Deployment".
2. Use easily identifiable cables labeled RNP (Real-time Network Protection) or similar, to clearly mark the connections – document along your internal guidelines for infrastructure moves, adds, and changes.

(Melior recommends using **red colored** certified "Cat-6"-grade Ethernet cable to connect your Barbican RNP system).

3. Plug in the power cable and network cables in their respective ports, according to the picture above:

Connect the “inside” – the port to the network you wish to protect – to the **CyberWarfare Defense Port I** (left).

Connect the “outside” – the port to the network from which you expect malicious traffic to originate – to the **CyberWarfare Defense Port II** (right).

The two non-addressable CyberWarfare Defense ports support 10 MBit/s, 100 MBit/s, and 1 GigaBit/s connections. Please use the appropriate copper Twisted-Pair Ethernet cables with matching specifications for the connections.

If you are using gigabit-speed bandwidth, your connection media may be via fiber connections. Various standard media converters for TP copper Ethernet to fiber media are available from Melior or from appropriate network media supply sources.

Note:

- ❑ *If Barbican RNP is plugged between network devices (routers, switches), standard, or “straight-through” Ethernet Cat-5 or better cables need to be used.*
 - ❑ *Use a “crossover” cable, if you are connecting Barbican RNP directly into a server.*
4. *Optional for CyberWarfare Defense functionality:* to allow for remote administration, maintenance feature upgrades, and “syslog” reporting, the third Ethernet port (**Management Port**) needs to be connected.

The Management Port supports both 10 Mbit/s and 100 Mbit/s connections. Connect the Management Port to your secure administration network (the same, which you use to administer and collect alerts from your Intrusion Detection Systems).

DO NOT connect the Management Port to your “live” network (i.e. the network you are protecting against with Barbican RNP). The unit cannot be detected or compromised on the CyberWarfare Defense Ports, as those have neither a MAC- nor IP-Address, and is not “finger-printable”.

Note: Connecting the addressable Management Port to the “live” network could potentially compromise this feature.

Until the Command Line Interface (CLI) is released (Feb 2005) and allows for configuration of the Management Port, and for initial setup, this port requires a DHCP lease to function.

5. To permit automatic feature updates, you will need to grant inside-to-outside Internet access to the Management Port, so the unit can check for, and retrieve, updates from Melior.

By default, Barbican RNP checks for updates every night at 2am CST U.S. time, and will update and reload, when new firmware is available. All feature upgrades are extensively tested through a rigorous Quality Assurance process prior to the release to the customer.

If your maintenance window requires a different time to check and/or apply feature updates, or you prefer to schedule the reboot-time to apply the features, you can:

1. Not grant Internet access to Barbican RNP from the Management Port
 2. Configure Barbican with the CLI (upon release) to check at a different time, and set the time for the feature application (reboot)
6. Command Line Interface Configuration
(Configuration and “syslog” functionality are disabled until the CLI is released – an updated manual will be made available at that time)
 7. Barbican RNP does not require a monitor, keyboard or mouse, but they can be used to access the console (a Barbican RNP CyberWarfare Defense logo is displayed as a screensaver).
 8. You are now ready to power on Barbican RNP. Please allow up to five full minutes for Barbican RNP to properly load and start operating.

9. **The system is now operational - you have CyberWarfare Defense!**

**You are protected against Penetration Testing probes,
and safe against distributed Denial-of-Service (dDoS) attacks.**

For details on all attacks Barbican RNP defends against, please see the section on “Attack Protection” in this manual.

10. To verify (you are surely curious so see Barbican RNP at work) Melior recommends:

- ❑ A port scan (i.e. using a free Penetration Testing tool, such as “NMAP”, *available on the Melior F.I.R.E. CD along with other dDoS attacks, or at www.insecure.org/nmap*). You will receive a port scan result indicating the ports are masked (“open” with no further information revealed) and the fingerprint is unsuccessful (“unknown device”).
- ❑ To test for dDoS defense capabilities, use the attacks provided by Melior on the F.I.R.E. CD, or download any tool from the Internet (*for details on attacks Barbican RNP defends against, see the section on “Attack Protection”*).




Deployment Options

The CyberWarfare Defense Layer, provided by the installation of Barbican RNP, can be deployed in multiple environments. Here are a few suggestions:

- ❑ Online Hosting
- ❑ Perimeter Defense
- ❑ Internal Networks

Example 1: Barbican RNP Deployment defending Online Hosting

Barbican RNP CyberWarfare Defense defending your online presence in a data center.



Most online presences for content web sites and e-commerce applications are hosted in data centers, in which the customer rents space in cabinets or cages for their own equipment, and ISPs provide Internet connectivity typically by way of 100 Mbit/s or faster Ethernet connections on copper or fiber wiring.

The customer's infrastructure setup is usually fairly similar: firewalls protect against basic attacks and limiting the number of TCP, UDP, and ICMP ports, on which the outside world can communicate with the systems deployed in the customers' space, and a variety of network- and load balancers, as well as a number of application servers.

Location of any online presence is usually chosen for environmental redundancies, as well as for availability of a multitude of Internet feeds. The customer pays for the space used within the data center (fixed cost), as well as for the amount of traffic delivered to and from the infrastructure by the ISP (dynamic cost). Usually the dynamic bandwidth cost has a fixed minimum price (subscription rate), with the capability to „burst“ from the minimum rate to the full capacity of the medium (the „pipe“).

If the customer comes under „bandwidth flooding“ dDoS attack, the amount of traffic delivered to the pipe typically goes to the maximum, resulting in saturation of the pipe by invalid traffic, thus overwhelming the transport capabilities through the pipe, and resulting in maximum revenue for the ISP. For the revenue aspect, no ISP has a business interest to defend the customer against such Denial-of-Service attacks (and Service Level Agreements –SLAs- typically exclude dDoS attacks), nor does the customer usually have other leverage to request action to be taken by the ISP without compensation.

In order for Barbican RNP to defend against all attacks (see brochure on „Attack Protection“), including the costly flooding of the bandwidth pipe, it must be installed „upstream“, i.e. on the ISP end of the traffic delivery medium, prior to the point of measurements for bandwidth cost.

To provide the ISP with a business incentive, ISPs typically prefer to purchase Barbican RNP units, and offer their deployment and usage as a premium service to the customer, to offset the lost revenue for the additional dDoS, PenTest, and junk traffic.

As a customer, you would want to choose your co-location provider from a list of ISPs, who deploy and offer Barbican RNP CyberWarfare Defense Solutions as an add-on service. If you already have an online presence in a co-location facility, and your ISP is not on the list of Melior resellers or premium providers

(see our web site), you may want to work with your ISP to consider offering this defense as a premium service, or to deploy the Barbican RNP unit you just purchased.

The cost benefit is obvious: compare the rates of ongoing bandwidth cost exceeding your monthly subscription rate, and the cost of damages, outages, and vulnerability exploitation (PenTest - against which your current security measures do not defend) against the cost of Barbican RNP units (or their premium service cost by the ISP).


You will find Barbican RNP amortizes very quickly.

For additional cost / benefit data, please contact your local reseller or Melior's sales department.

Example 2: Barbican RNP Deployment defending the Perimeter

Barbican RNP CyberWarfare Defense defending your entire perimeter edge towards the outside world.

Deployment: Perimeter Defense: On-Premise



Most medium and large networks are set up in a similar architecture design, following a „best practices“ approach established - until now.

One or more Internet connections are shielded by firewalls, protecting against basic attacks and limiting the number of TCP, UDP, and ICMP ports, on which the outside world can communicate with the systems made available for public access, or on which services using the Internet can communicate (out-bound e-mail, web surfing, Voice-over-IP, DNS, etc).

Penetration Tests probes for vulnerability probing are not intercepted, and the infrastructure is open to attacks and compromises. Each firewall is addressable and runs – like any other device accessible from the Internet – on a TCP/IP stack, which in itself is vulnerable.

With freely available tools, any individual, group, or competitor with malicious intent can determine the exact configuration of the customer's infrastructure (with tools such as NMAP or Nessus), and attack it at will to disable individual components or shut down devices and access to and from the entire infrastructure. A simple stack attack on the firewall will disable all communications instantly.

Barbican RNP defends against these attacks (*please see brochure on „Attack Protection“*), including the PenTest probes (to also protect the on-premise router, it must be installed „upstream“, i.e. on the ISP end of the traffic delivery medium – *please consult our brochure on on-line presence defense*).

As a customer, you will want to eliminate the risks of vulnerability exploitation and direct attacks on your infrastructure by shielding your firewall(s), network devices (switches, routers), and servers, now that CyberWarfare Defense Solutions have become available with Barbican RNP v1.0.


The benefits are obvious: compare the cost of cost of damages, outages, IT staff usage, and vulnerability exploitation (against which your current security measures do not defend) against the cost of Barbican RNP units (or their premium service cost by the ISP).

You will find Barbican RNP amortizes very quickly.

For additional cost / benefit data, please contact your local reseller or Melior's sales department.

Example 3: Barbican RNP Deployment protecting Internal Networks

Barbican RNP CyberWarfare Defense defending your entire internal infrastructure.



Most networks in corporations and governments are set up in a similar architecture design, following a „best practices“ approach established - until now.

User desktop systems share local subnets (LANs), interconnected with Ethernet switches and routers, and providing connectivity to commonly shared resources, such as file- and print servers, e-mail servers, intranet web servers, Voice-over-IP (VoIP) telephony services, and many other applications. External sites (manufacturing, remote offices, privileged vendors) are also often connected via low-bandwidth WAN lines, or through the Internet. The entire infrastructure is typically connected to the Internet via one or more ISP links, and a firewall protects against basic attacks and limits the number of TCP, UDP, and ICMP ports, on which services using the Internet can communicate (in- and out-bound e-mail, web surfing, Voice-over-IP, DNS, instant messaging, etc).

Some high-security conscious sites deploy firewalls also internally between LANs and WANs, and in front of internal application farms, to control access to resources within the infrastructure; however, the cost to deploy and maintain firewalls and their policies is too cost-prohibitive in most cases to actually implement.

Penetration Tests probes for vulnerability probing are not intercepted from external or internal sources, and the infrastructure is open to attacks and compromises. Each firewall is addressable and runs – like any other device accessible from the Internet – on a TCP/IP stack, which in itself is vulnerable – and the same is true for network devices (switches, routers, VPN gateways), servers, and desktop PCs.

With freely available tools, any individual, group, or competitor with malicious intent can determine the exact configuration of the customer's infrastructure (with tools such as NMAP or Nessus), and attack it at will from external or internal sites to disable individual components or shut down devices and access to and from the entire infrastructure. A simple stack attack on a firewall, router, or Ethernet switch will disable all communications instantly; attacks on servers and applications will disrupt those equally quickly.

In addition, user desktop PCs are often compromised by virii, worms, and other forms of malware (studies have shown 40% of desktops in US Fortune 100 companies are compromised), and the malware spreads very quickly across the entire infrastructure, contaminating other LANs, remote sites via WANs, and server farms.

Even simple invalid and „junk“ traffic (such as jabber from malfunctioning network devices; i.e. bad traffic without malice) flows freely across the LANs and WANs, eating up resources. While on a 100 Mbit/s or faster local network this extra traffic is not as much noticable, it often easily saturates the smaller wide-area network connections, and leads – with growing networks – to costly upgrades of such connections.

Barbican RNP defends against both invalid, non-malicious traffic, as well as against malicious attacks from internal and external sources (*please see the section on „Attack Protection“*), including internal or external PenTest probes to exploit vulnerabilities in local or remotes sites as preparation for targeted attacks.

As a customer, you will want to eliminate the risks of internal and external vulnerability exploitation and direct attacks on your infrastructure by shielding your firewall(s), network devices (switches, routers), desktops, and servers, now that CyberWarfare Defense Solutions have become available with Barbican RNP.

The benefits are obvious: compare the cost of cost of damages, outages, IT staff usage, and vulnerability exploitation (against which your current security measures do not defend) against the cost of Barbican RNP units (or their premium service cost by the ISP).

Benefits of decreased IT staff costs due to:

- Containment of malware within individual subnets with infected desktop PCs
- Protection of critical LANs and WANs by at least partial Barbican RNP deployment for those networks only
- Drastically reduced log sizes in firewalls and Intrusion Detection Systems (IDS) - less IT staff time spent to review logs efficiently
- Reduced costs to rebuild firewalls, servers, desktops, and other devices from TCP/IP stack attacks
- Reduced costs from compromises from worm/virus attacks
- Reduction/Avoidance of follow-up costs of vulnerability exploits due to elimination of Penetration Testing for yet-uninfected systems (IT staff time spent to keep up with system patches)
- Budget benefits from improved reporting on valid vs. invalid network traffic

You will find Barbican RNP amortizes very quickly.

For additional cost / benefit data, please contact your local reseller or Melior's sales department.

The Introduction of a new CyberWarfare Defense Layer

Since the invention of firewalls, no major technology development has occurred in computer, network, and application security in the last 10 years. Firewalls came as the first measure to provide some security, after local networks and the Internet were set up to simply provide ubiquitous connectivity with no concern for security. The network devices directing the traffic (i.e. routers) could be configured for some basic security (by way of Access Control Lists – ACLs), but due to very limited CPU capabilities, and an enormous cost in performance, proved to be of little use for security.

The introduction of firewalls then allowed a function of a gateway on the perimeter, to authenticate access, and to restrict the total number of 65,535 TCP and 65,535 UDP communication channels to the few ports actually used. However, this approach provides at best a „Swiss Cheese“, as the ‚open‘ ports – i.e. the most popular ones - remain available for both valid traffic (such as web access, e-mail, file transfer, encrypted traffic, etc) as well as invalid traffic, since firewalls have to pass all traffic without discrimination, and do not have the ability to make decisions, which traffic is valid, and which is not.

In addition, firewalls and all other devices communicating with a TCP/IP protocol („stack“) are addressable, and can be directly attacked themselves. By sending malformed data, non-complying with the established protocol rules, or carefully crafted to operate outside established boundaries, every single one of such devices can be very quickly eliminated („crashed“), thus disrupting communications and shutting down the availability of resources.

See also the section „Is A Firewall Enough?“

Penetration Testing: Mapping Vulnerabilities

Additional tools allow for an exploitation of configuration, thus mapping the infrastructure, and providing an exact catalogue of deployed hardware, software, and its configuration, which in turn can be easily used to determine the weak points („vulnerabilities“) and to exploit those in targeted attacks. These tools are freely available and very powerful („Penetration Testing Probes“). Programs such as NMAP or NESSUS will provide such maps of an entire architecture design, and in addition even rate the difficulty for compromise.

Such vulnerabilities can then be exploited even further, with distributed Denial-of-Service (dDoS) attacks.

What is a Denial-of-Service (DoS) attack?

Broadly defined, a DoS attack seeks to overwhelm the capabilities of the targeted site, either from the outside (the Internet) or from the inside (internal systems), with the intention to disrupt (deny) service – by disabling („crashing“) the targets or by sending so much traffic, that the target is no longer able to respond.

These attacks come in many different flavors. They can be very small, sending a short burst of only a few malformed packets (such as in „shrew“ attacks, „synk4“, or „synk5“ attacks) against just one critical system, representing a key, single-point-of-failure component (such as a firewall), and almost instantly taking it down, thus shutting down access to and from all systems behind this device.

The attacks can also be very large. The concept of „bot-networks“ has matured, in which a large number of computers on the public Internet are quietly (without the knowledge of the owner) compromised, and thus listen to a single command from a „bot-net-operator“ to send attack traffic against a particular target simultaneously with thousands (or even millions) of other computers. These attacks are usually carefully crafted, and hide the IP address of each such computer, pretending with each individual packet sent to come from a different IP address (this function is called „sender address spoofing“). Since this form of DoS attack comes from so many distributed systems, it is called distributed Denial-of-Service, or dDoS, attack.

With existing security measures, the attacked site has no means to determine which part of the incoming bandwidth is valid traffic from regular visitors with their correct IP addresses, and which part is invalid, coming from „spoofed“ computers. Consequently, the attack is successful, usually works immediately, and keeps the targeted site offline, until a central command is issued to all the systems in such a „bot-net“ to cease the attack.

This method is virtually risk-free to the attacker, as it is technically impossible to identify all the participating computers in a „bot-net“, and trace back from those to the operator, issuing the single attack and cease commands. Even if that were possible, the attacker is likely located in a different jurisdiction than the targeted site, thus excluding any reasonable possibility to take appropriate action.

Distributed Denial-of-Service attacks have become increasingly popular in the last years. They are used by organized crime groups to extort large sums of money from e-commerce services (for example, virtually all online betting sites in the UK are paying extortion fees after several large attacks), they are used by political groups (news sites have been shut down for extended periods of time),

and the risk of terrorists attacking entire countries with the desire to disrupt that countries' Internet commerce, communication channels, and economy is another great concern to Government officials.

The Introduction of a new CyberWarfare Defense Layer


First major technical advancement in network security since firewall introduction

Universal Threat Management

- through conjunction of traditional and new methods

Intrusion Prevention

- with True Zero-Day Attack Protection



In order to defend against these malware, PenTest, DoS, and dDoS attacks, the existing security measures need an additional layer: a layer, which inspects all traffic prior to reaching critical systems, and eliminates all (sic!) invalid traffic.

This functionality, now for the first time available with Melior's **Barbican RNP**, goes deeper, to Layer 2 and 3 of the OSI model, while inspecting traffic on application Layers 4 and above for validity as well.

- ❑ **Not detectable** on the network – so it cannot be addressed itself, and thus not be compromised (no MAC or IP address)
- ❑ **„Inline“ Appliance**: the only viable approach to inspect **all** traffic first, and not let a single bad packet pass, requires the CyberWarfare Layer to sit „on the wire“. Any other approach, observing traffic flowing through other devices, does not have sufficient time to react, and thus – by design - has to let invalid (or „bad“) traffic pass, while it determines if action needs to be taken.
In addition, the Barbican RNP approach as an „Inline“ device requires no changes to the various existing environments, no additional switch ports,

or any re-engineering of the protected architecture. Since this is a single point of failure, the CyberWarfare Defense layer must be protected against failures, thus have as little moving parts as possible, and come with extensive protection against failure (at least the same level as other critical single-points-of-failure devices)

- ❑ **Real-Time Defense** is mandatory to this layer to function well. Valid traffic must pass the CyberWarfare Defense layer with a minimum of latency, not to exceed 10 milliseconds, while all invalid and attack traffic is discarded instantaneously.
- ❑ **True „Zero-Day“ Attack Defense** is another mandatory requirement. ALL other security devices use a library or „signature“ data base to compare observed traffic against, such as firewalls („policy definition“), Intrusion Detection Systems („signatures“), anti-Virus software („Virus Libraries“), and anti-Spam software (a variety of solutions, from dictionaries, blacklists, heuristic definitions, etc).

This approach can only recognize „known“ – previously observed – forms of attacks, and by design, the growing size of such definitions (which need to be constantly updated) make the comparison process increasingly slow, before any counter measurements can be taken. New and modified forms of attacks, deviating from those pre-defined, known attacks, will pass the security device without challenge.

The CyberWarfare Defense layer must not „know“ any attacks, and look at every single packet with the ability to decide if it is valid or invalid: it must operate „**signature-free**“; thus providing true „Zero-Day“ attack defenses. Barbican RNP is the first such product: it's patented technology provides a „Truth Table“, which inspects all traffic for validity without any comparison process, and acts in real-time to pass on the valid traffic, while discarding (and reporting on) the invalid traffic.

For details on defense capabilities, please see the section „Attack Protection“

- ❑ **Penetration Testing Defense** is another critical component of the CyberWarfare Defense layer. All vulnerability and mapping probes must be intercepted prior to reaching any device on the network: if a firewall were to perform this function (or any other addressable device), it would still minimally disclose the type of this device, thus opening the opportunity to address, attack, and disable the device.

Barbican RNP intercepts and responds to such probes, increasing the difficulty level for compromise by adding partial (probed port is reported as „open“, but no other information is provided) and random (so Barbican RNP itself cannot be catalogued) data to the response. In turn, it allows other, existing security technologies, such as Intrusion Detection Systems (IDSs) to function better, by „tuning“ the device to listen for such Barbican RNP PenTest responses and trace the source of the probe for appropriate counter-measures.

- ❑ **Ease of Deployment** is important to deploy the CyberWarfare Defense layer across a variety of infrastructures, with or without skilled IT Security staff present, as well as to achieve instant protection. If it were to require a lengthy „learning“ period to gain information of „normal“ operations, it would deploy old bandwidth throttling technology, and be inefficient to provide instant protection.
Barbican RNP works in a „Plug & Protect“ mode: as soon as the CyberWarfare Defense device is plugged in, it protects instantaneously.
- ❑ **Affordability** is key for wide-spread deployment of the CyberWarfare Defense Layer, so access to this technology is open to every single organization, whose systems are critical to its infrastructure, regardless of the available security budget.

Barbican RNP is aggressively priced to provide this affordability in a variety of deployment options for the CyberWarfare Defense Layer.

With the availability of Melior's Barbican RNP, a CyberWarfare Defense Layer is available for the first time since firewalls were invented, providing a major step forward in the mix of security measures towards a unified threat management – and offering better options to comply with legal network- and application-protection requirements.

300 million users, who were already successfully protected against dDoS attacks by Melior CyberWarfare Defense technology, attest to the validity of the CyberWarfare Defense layer.

Additional testimonials demonstrate the benefits of PenTest defense, reduced processing loads on firewalls as a result, reduced log sizes, and additional financial benefits in the reduction of IT staff resources.

Is a firewall enough?

The majority of network administrators use some kind of Firewall to protect their network. A Firewall provides policy-based access control to the various devices behind it, generally while allowing access to the Internet from the internal machines. Most people assume that a Firewall can protect the network from an attack.

There are many types of Firewalls. Some operate as filtering gateways with simple port level protection, while others provide detailed connection state tracking. Some Firewalls operate at the application level using a proxy server. There are advantages and disadvantages to each approach, and some more expensive Firewalls use multiple approaches, involving both policy- and proxy-based systems.

All Firewalls share a similar disadvantage: they are addressable.

Firewalls work on the upper layers of the OSI model, taking data from one IP address and sending it to another (or from one group to another). They make an "appearance" on the network as any other device would. Once there, Firewalls always listen on an IP address, and data can be sent to a Firewall directly, since for the most part Firewalls concern themselves with layer 3 and above of the ISO stack. Because of this, they are generally not prepared to handle network anomalies, such as the foundation of hackers' and attackers' tools in use. The systems these "bad guys" use, allows them to hand craft packets, which otherwise would not exist on a network, such as oversize ICMP pings, bad packet fragments, smurf amplification attacks, badly formed packets, land attacks, and SYN floods. It is this departure from the expected network behavior, which has allowed these attacks to have good success exploiting weaknesses inherent within the modern day operating systems and application servers.

The Expected vs. the Unexpected

Because Firewalls work on the expected and apply their policies to that, they are weak in the area of the unexpected. Moreover, since they operate in the upper levels of the OSI model, they simply do not need to understand how the data gets to them - and that can be exploited as well. Hackers know all the details of this information and can craft packets, which break standards or simply bend the rules to their own gain. To a hacker, this is a clever little game to see which is smarter: their expertise or the Firewall? If a hacker knows the TCP/IP stack better than the Firewall does, or exploits a known vulnerability, he or she can bring it down, and your whole network with it. Don't think of a Firewall as the end

of network protection, but as possibly one of the weakest links in your network security solution - and often a single point of failure as well. It is addressable and it does not know much if anything about the lower levels of packets, and relies primarily on administrative policy to protect network assets.

What this means is that it is likely your Firewall is vulnerable to various forms of attacks. Even worse, it may blindly pass the attack through to your internal network or simply crash when it sees certain attacks (such as "synk4"). In addition, the "holes" you have punched into your Firewall to let valid traffic through, are staring hackers in the face. By way of powerful (and free) Penetration Testing tools, such as NMAP or Nessus, they can see exactly which ports are available and will concentrate their attacks there. Your weaknesses are as good as advertised! While Firewall vendors have attempted to address the dDoS problem, the basic problem will always exist. As long as your Firewall is addressable, it is vulnerable!

The solution then is a network protection device that is not addressable or detectable in any way and is capable of blocking anomalous traffic.

Melior, Inc. introduces Barbican CyberWarfare Defense Systems

Barbican - "A fortification at a gateway or drawbridge on the approach to a castle."

Is your Castle's Internet gateway in need of fortification? Are you currently experiencing distributed Denial of Service (dDoS) attacks, or observe several Penetration Testing probes on a daily basis?

Barbican dDoS attack defense, developed by the authorities on CyberWarfare Defense, is an expansion on Melior's proven expertise in dDoS attack defense.

Barbican takes a radically different approach from most Firewalls. Instead of being a server that connects two networks through its assigned IP numbers, the Barbican device can be compared to an Ethernet bridge that functions in-line. The Barbican device does not "listen" to any IP. It is as completely un-addressable as an Ethernet cable would be.

This means that unlike traditional Firewalls, it cannot be addressed directly so it cannot be attacked. More over, good valid data enters one port and exits the other without any IP or MAC address being changed in the packet.

However, Barbican is much more astute than any Firewall, because it has access to every detail of every packet. It knows the real-time details of IP packet assembly and TCP stream sequencing better than the hackers, and can respond far faster than any hacker, generally in the millisecond range! Unlike an Intrusion Prevention System (IDS), it uses no signatures to slow it down or provide false positives, and it does not require network baselines.

Barbican actually handles every connection, and once a packet is verified as being part of a valid connection, it recreates each packet on its protected network interface to be sure it is pristine. Following every TCP/IP rule to the point of being pedantic, Barbican glues IP fragments back together to protect from fragment attacks, deals with port scans (i.e. Penetration Testing probes), floods, hand-crafted packets and similar anonymous behavior. Furthermore, Barbican will detect and report to port scans, and report the IP address of attackers. Forged IP packets are reported on and then simply dropped.

The Barbican device is actually capable of making sure that every packet passing through it is part of a valid connection. It is so efficient, that not a single dDoS packet is ever sent into the protected network. This includes SYN floods! Not even a single SYN of a SYN-flood attack will enter the protected network!

Barbican can sit in front of a traditional Firewall to protect that Firewall from attack, while often lightening its processing load, allowing completely transparent and configuration-free installation into networks, where Firewalls are already deployed and configured. In addition, it can act similar to a policy Firewall and completely block unused ports and services, without anyone on the outside knowing that the port is blocked. Data sent to blocked ports, like any invalid data, is simply dropped.

The Barbican device not only tracks connections like an advanced Firewall, but acts much like an application level proxy-based Firewall to give the benefits of proxy firewalling, but on every TCP port. Furthermore, it does this in such a way as to make it look like all ports are open. Hackers don't know where to focus the attack! It also has a unique TCP fingerprint evasion technique preventing hackers from determining what operating system your servers are running. Simpler scanners will often end up scanning their own reflections!

The combination of intensive connection tracking and network cloaking against Penetration Testing on a completely un-addressable device marks a new approach to network security that can be deployed into any network for immediate dDoS protection.

Attack Protection

As an „inline“ device, Barbican RNP is designed to deliver on two key purposes:

- ❑ Inspect every single packet and packet stream for validity, to identify invalid traffic and threats, and discard and report on such identified „bad“ traffic
- ❑ As little as possible impact and latency on valid traffic due to real-time inspection, so normal operations are guaranteed

In order to deliver the desired benefits, Barbican RNP protects against these attacks:

Connection flood attacks:

The attacker exhausts the target's resources by creating a large number of unused connections to the target.

SYN flood attacks:

The attacker exhausts the target's resources by starting a large number of connection "handshakes" that it never finishes.

Malformed header attacks (TCP, UDP, ICMP):

The attacker attempts to take down, take over or degrade the performance of a target by sending packets whose headers contain invalid data.

Packet fragmentation attacks (TCP, UDP, ICMP):

The attacker attempts to take down the target using packets deliberately crafted to take advantage of faults in certain operating systems or network devices.

Random-source UDP flood attacks:

The attacker consumes the target's network resources by originating a large number of UDP datagrams that appear to come from random origins.

TCP State Machine attacks:

The attacker attempts to take down or degrade the performance of a target by performing protocol actions out of expected order, or in ways that "confuse" the target device's network communication system.

Malformed TCP packet attacks:

The attacker attempts to take down, take over or degrade the performance of a target by sending packets that contain various kinds of invalid data.

These attacks can be geared against specific network devices, operating systems, applications (forinstance: web servers, E-Mail servers, DNS, Voice-over-IP telephony services, infrastructure services such as BGP4) in either broad-scale attacks, overwhelming services or bandwidth pipes, or specifically targeted against known or emerging vulnerabilities.


The determination of vulnerabilities specific to each targeted site and infrastructure is prevented by the Barbican RNP Penetration Testing Defense (PenTest), which „cloaks“ the infrastructure against probing attacks.

In the initial v1.0 release, Barbican RNP is geared to protect against OSI Layer 2 and 3 attacks; however, many attacks targeting applications on OSI Layer 4 and above can be recognized and thus defended against on the lower levels. Melior has several modules in development, addressing CyberWarfare Defenses against OSI Layer 4 and above attacks, which will become available for easy installation on existing Barbican RNP platforms by the customer.


Barbican RNP can be deployed across the entire network infrastructure, to defend against external and internal threats. The Barbican RNP product is purposely priced very aggressively, to make wide-spread CyberWarfare Defense affordable to deploy

- ❑ At data center locations to protect your online presence
- ❑ At the external perimeter to protect against external threats
- ❑ At internal network connection points between LAN subnets and between LANs and WANs, to protect against spreading attacks from contaminated PC desktops by containing contamination within local subnets, and to protect against bandwidth saturation on WAN links


Data Center



Perimeter Defense



Internal Network Defense



See the other sections discussing the details of various deployment scenarios.

Background on distributed Denial-of-Service (DDoS) attacks

Over the last thirty years, most software development targeted delivery of functionality, not security. This holds true for communications software and protocols, operating systems, application software, and code embedded in hardware appliances. While computers and networks have become necessary tools to conduct commercial, government, and educational business, provide the means of global communications linking individuals, institutions, customers to vendors, citizens to governments, students to universities, thus creating a dependency well beyond a point of return to other, traditional means, the entire infrastructure remains highly vulnerable against misuse.

These vulnerabilities affect every level of communications use: from the small Local Area Network, the Personal Computer at a residential home with a connection to the Internet, to enterprise networks, e-commerce businesses, government systems, nations, and the global Internet infrastructure. In essence, the exploitation of vulnerabilities establishes "CyberWarfare" between legitimate users of resources and their abusers, those being individuals of various degrees of sophistication, competing businesses, or unfriendly nations.

With the evolution of the Internet, the TCP/IP network protocol has emerged to be the predominant communications protocol with its IPv4 structure, which in itself is highly insecure. Even without their own inherent insecurity, all network-based software applications and advanced communication protocols are vulnerable from the underlying TCP/IP network protocol. The development of IPv6, initially geared towards providing more network address space to meet the growth of the Internet, will also provide some improved security, but is many years away from practical implementation.

Vulnerabilities are well known and accessible to everyone: a simple Google search in June 2003 for the term "Network Vulnerabilities" returns over 442,000 results, a search for "Application Vulnerabilities" returns over 275,000 results. Compiled in large databases such as BugTraq, this information allows both system administrators as well as malicious attackers to learn of detected risks and utilize the information.

Since computer security firms and attackers are usually the ones finding the vulnerabilities, system administrators and those responsible for the security, integrity, and availability of computer/network infrastructure find themselves in the undesirable situation to keep up with newly detected risks and the search for solutions to mitigate the risks. This has always been a re-active approach, looking for application software patches, and devices to address individual risks.

Commercial vendors and research projects in universities and governments have always tried to overcome the technology hurdles with mixed success to secure the critical computing and networking infrastructure from its own inherent insecurity, while the malicious attackers develop ever more sophisticated tools to compromise such efforts, and freely share these tools. Due to the re-active nature of defenses, the malicious attackers always stay ahead.

Vulnerability Determination

The first step for any attack, both from the inside as well as the outside (commonly the Internet), is to find out which systems are in use at the targeted infrastructure. Powerful and freely available tools are easily accessible to anyone, and return detailed configuration results within minutes (this process is also frequently referred to as "enumeration" by scanning network addresses and ports).

Tools such as NMAP and NESSUS (to name just two) penetrate multiple levels of firewalls, Demilitarized Zones (DMZs), and internal networks, and return exact results, specifying which systems and mechanisms are deployed to defend the infrastructure, which brand, operating systems, release versions, servers, and applications are in use. These tools also provide an automatic assessment of the grade of difficulty to compromise each detected system.

With this information, the malicious attacker can look up how to best target an attack, exploit known vulnerabilities ("known" attack) or go a step further and develop a new form of attack specifically geared towards the mapped target ("unknown" attack).

Currently, without exception, all network devices (such as routers, bridges, and switches), all defense devices (such as firewalls), and all application devices (such as servers, storage systems, workstations, and individual computers) are vulnerable to this information gathering. In other words, all systems connected to the Internet are "sitting ducks" for being mapped from public sources, as well as all systems on internal-only networks from any connected individual computer on that network.

Information gathering on system and network configuration is conducted on an ongoing basis. Logs from IDS systems and firewalls are filled with these probes, which are either automatic sweeps by IP ranges, or probes against specific targets.

The ideal solution to mitigate this first and most important step for every attack is to "hide", or "cloak", each infrastructure and system on internal networks and on the Internet. Such a solution is able to detect information gathering probes, and respond actively in a fashion which does not reveal the systems and configurations, while at the same time not impact the normal traffic.

The technology hurdle to deliver such a solution had not been overcome prior to Melior's iSecure and Barbican products.

TCP/IP Protocol Attacks

TCP/IP has evolved to be the most-used network communications protocol. It is exclusively used for the Internet, and widely used on internal networks, complimented by IPX and other legacy protocols, which are increasingly diminished. The internal structure of the TCP/IP protocol, represented by the 7-layer OSI model, is extensively documented and freely accessible to anyone for analysis of vulnerabilities (in "RFCs", also known as Requests for Comments during the design of every commonly shared function of interaction between computing systems).

All systems using TCP/IP-based network communications are accessible (and thus attackable) via all seven layers of the OSI model, spread over a total 65,535 ports for TCP, ICMP, and UDP communication.

Along with the hardware Layer 1, Layer 2 (data link, addressable via the six-byte unique identifier, or MAC – Media Access Control) and Layer 3 (logical network address, represented by the assigned IP number) are the prime targets to access and identify any system on a network.

The first three bytes of the MAC address is vendor-assigned, allowing to determine which vendor provides the hardware in use.

Within TCP/IP, the packet, by which information is exchanged, is well documented, and allows for a lot of options to be set between two communicating devices. Many of these options are not practically used, but allow for the exploitation of vulnerabilities. Many low-level network (or TCP/IP protocol) attacks work through the manipulation of any of the less or not used fields, or by sending and requesting information outside the protocol specification. Every network device using TCP/IP as their communications protocol is vulnerable to such attacks, and many known forms of such attacks exist. The behavior of responses to low-level communications is used to determine system configuration.

Current security defense technology has not overcome the technical hurdles to eliminate packet-based attacks. Any system deployed on a network (including firewalls, Intrusion Detection Systems, VPN Gateways), which is addressable, is in itself vulnerable to such attacks. Many threats and risks are based on malformed packets (such as most Denial-of-Service -or DoS- attacks). Known tools such as "Synk4" are able to make any TCP/IP-based device unavailable within seconds.

The ideal and effective solution to TCP/IP protocol attacks is an inline-scanner inspecting all traffic flowing through it for validity. This device must not be addressable by either MAC (Layer 2) or IP number (Layer 3) in order not to be vulnerable in itself, and must have the ability to determine valid ("good") packets from invalid ("bad") packets, discarding the invalid packets and allowing the valid packets to follow through. Such a device must work in real time or near real-time to not introduce latency within the communication.

(distributed) Denial-of-Service Attacks (DoS/dDoS)

Distributed Denial of Service Attacks emerged as one of the most newsworthy, if not the greatest, weaknesses of the Internet.

These distributed Denial of Service (dDoS) attacks are not a new development; but have been aggressively and increasingly used in the last years. They first appeared and were first widely discussed starting in the summer of 1999. During the week of February 7th through 11th, 2000, they emerged as a major new category of attack on the Internet. They took out many sites, including Yahoo, Buy.com, eBay, Amazon, Datek, E*Trade, and CNN. The victims were unreachable from several hours to over three days each.

Subsequently, large-scale dDoS attacks almost disabled the entire Internet twice by attacking the DNS root structure in October 2002, and created significant tangible and intangible damages in the "Slammer" attack during Superbowl 2003, in which Bank of America (USA) lost all of its 13,000 Automatic Teller Machines (ATMs) for two days. During the Iraq conflict, media sites suffered dDoS attacks resulting in extended outages rendering the web sites inaccessible. The most widely reported victim during this period was the independent Arabic news source "Al Jazeera", whose English site was unavailable for several weeks. Even the largest providers such as Akamai were unable so withstand dDoS attacks.

Due to the impact of dDoS attacks, businesses have been shut down (such as one of the oldest ISPs in the United Kingdom, "Cloud Nine", in 2002), or severely impacted with revenue loss and intangible damages.

By now, dDoS attacks have become a successful tool for organized crime in extortion attempts, shutting down payment processing providers such as WorldPay, Authorize.net, and many others.

In response to the significance of the threat, the Federal Bureau of Investigation (FBI, USA) classified distributed Denial-of-Service attacks as a threat to national security as early as February 2003, and the risk remains as a high-priority item on the newly formed Department of Homeland Security ("DHS") in the United States, as well as for many other governments and commercial businesses.

A 2001 study by the Supercomputer Institute of the University of San Diego measuring the backscatter traffic found an average of 4,800 dDoS attacks per week on the Internet, with increasing tendency.

Denial-of-Service Attacks also rank amongst the highest risks for damages due to their maturity and visibility, as the Gartner Group reported already in June 2003.

Entire armies of "bots" can be rented for as little as \$150 per attack; according to a 2004 Symantec study, 30,000 computers are taken over every day in extension of these "bot" networks (up from 2,000 / day six months earlier).

dDoS Background

dDoS attacks involve breaking into hundreds or thousands of machines all over the Internet by means of a direct compromise or via worms, virii, and automatic scans and execution of vulnerabilities. The attacker installs dDoS software on the compromised systems, allowing the attacker to control all these compromised systems to launch coordinated attacks on victim sites. These attacks typically either exhaust bandwidth, router processing capacity, network stack resources (see above, TCP/IP protocol attacks), or target applications directly with low-level port traffic. Through breaking network connectivity to the victims, disabling security devices such as firewalls, or exhausting operating system or system application resources.

The attacker starts by breaking into weakly secured computers, using well-known defects in standard network service programs, and common weak configurations in operating systems. On each system, once the attacker breaks in, some additional steps are performed. First, software is installed to conceal the fact of the break-in, and to hide the traces of the subsequent activity. For example, the standard commands for displaying running processes are replaced with versions that fail to display the attacker's processes.

These replacement tools are collectively called a "rootkit", since they are installed once the attacker has "broken root", taken over system administrator privileges, to keep other "root users" from being able to find the attacker. The attacker then installs a special process, used to remote control the compromised system. In the public environment, this process accepts commands via the Internet, and in response to those commands it launches an attack over the Internet against the designated victim site; the same process applies to internal attacks within an enclosed network. In an automated process, the attacker catalogues the addresses of the compromised systems. A cautious intruder will begin by breaking into just a few sites, then using them to break into some more, and repeating this cycle for several steps, to reduce the chance they are caught during the preparation phase. By the time the attacker is ready to mount the kind of attacks as widely reported (gigabytes per second of traffic for the attack on Yahoo in 2000, according to reports in SANS) the attacker has taken over thousands of machines ("dDoS handlers") and assembled them into a dDoS attack network; meaning the attacker has the attack software installed on them, and knows all their addresses (stored in a file on their control system).

To run the actual dDoS attack, the attacker runs a single command, which sends command packets to all the compromised systems, instructing them to launch a particular attack (from a menu of different varieties of dDoS attacks) against a specific victim.

When the attacker decides to stop the attack, another single command is sent. Only the attacker in control of the dDoS handler systems can start and stop a dDoS attack.

While there are variations, (distributed) Denial-of-Service attacks generally take one or a combination of the three common forms:

- Bandwidth-Flooding Attacks, overwhelming the available network capacity
- Network-Level Attacks, exploiting vulnerabilities in the TCP/IP (ICMP) protocol
- Application Attacks, targeting specific ports within TCP (such as port 80 for HTTP/web traffic) or UDP (such as port 53 for DNS lookups, or the port range used for VoIP telephony services) by packet or data payload manipulation

The controlled machines being used to mount the attacks send a stream of packets. For most of the attacks, these packets are directed at the target systems (Network-Level and Application-Level attacks). For the bandwidth-flooding variant (such as a "smurf" attack, named after the first circulated program to perform this attack) the packets are aimed at other networks, indirectly causing the damage by provoking multiple reverse ICMP echoes all aimed at the victim.

To go into further detail, some background description of the Internet:

The Internet consists of hundreds of thousands or millions of small networks (called Local Area Networks, or LANs), all interconnected; attached to these LANs are many millions of separate computers. Any of these computers can communicate with any other computer, using the above described TCP/IP address. The addresses are structured (organized into groups) so that special-purpose traffic-handling computers (routers), can direct them in the right direction to reach their intended destination. A typical connection today may require 15 or more hops, crossing from one LAN to another, before it reaches its final destination. But most of these "LANs" are actually special-purpose links within and between network transport companies (Internet Service Providers, ISPs). The larger ISPs (backbone providers) and interconnect exchanges provide the routing of traffic between networks.

As described above, when one computer wants to send a message to another, it divides it into fixed-size pieces, called "packets". Each of these packets is handled separately while in transit via the Internet, and then the message (if it is larger than a single packet) is reassembled at the remote computer. So the traffic passing between machines consists entirely of packets of data. Each of these packets contains a pair of addresses, called the Source and Destination IP (for Internet Protocol) addresses. These are the addresses of the originating machine, and the recipient system.

When such a packet is sent over the Internet, it is passed first to the nearest router; commonly this router is at the point where the local network connects to the Internet. This router is usually referred to as "border router". In larger organizations the configuration is more complex; assembling its own collection of LANs, interconnected into an in-house internet, cross-connected at one or more points (often with firewalls) with the Internet ("edge routers"). The border or edge routers pass the packets upstream to a core router, which interconnects with many other core routers all over the Internet; passing the packet on until it reaches its destination. Due to the inherent insecurity of the network communication protocol, the source address is normally ignored by routers; it

normally only provides the final destination address where the request is coming from.

The packets used in today's dDoS attacks often use forged source addresses, thus hiding (or "spoofing") the true originating source (i.e. the compromised systems). The very first router to receive the packet could technically catch a packet with a spoofed source address; it has to know which address spaces are configured on every network attached to it, so that it can correctly route packets to them. If a packet arrives, and the source address does not match the network it is coming from, the router should discard the packet. This style of packet checking is called Ingress or Egress filtering, but not widely implemented. If the packet is allowed past the border router, discarding packets with spoofed source addresses is nearly impossible.

From the victim's point of view, instant outages occur on routers, firewalls, or servers as thousands of compromised systems all over the world commence to attack. Indications are unavailability of routers, firewalls, or servers, and traffic simply stops flowing between systems and the Internet. Upon closer inspection it may be discovered that one or more targeted systems are being overloaded by the small fraction of the traffic that actually gets delivered, but the failures extend much further back.

Analyzing the traffic of a dDoS attack does not provide useful information. A capture of a sample of the dDoS attack packets continuously arriving at the target site show the correct destination address, with a random number as a source address. There is no trace of the compromised hosts continuing the attack. The only information available is the low-level, hardware (MAC) address of the last router forwarding the packet; these low-level addresses are used to handle distribution of packets within a LAN. So it can only be determined which router passed the packets last, but nothing else. Identifying the last router may identify the Internet carrier that passed the traffic to the target site, but requires as a next step to capture another packet on the other side of the forwarding router, and determine where that packet came from. Each step of the trace back to each compromised system requires repetition for every network interchange passed (15 or more on average).

Every time the back-trace crosses an administrative boundary, between the target and its Internet provider, between them and the next backbone provider on the path, all the way back to the compromised machine, the aid of another team of administrators is necessary to collect the trace data and carry the trace further back. Each trace requires typically several hours to complete.

This manual trace-back applies to each of the thousands of compromised machines that are participating in this attack, making it practically impossible to perform the trace.

Depending on the duration of a dDoS attack, it is unlikely to find more than a few of the thousands of machines used to launch the attack; the remainder will remain available for further attacks. And the compromised machines actually identified will contain no evidence that can be used to locate and identify the original attacker; the trace ends with the compromised system.

Many computing systems and software come with inherent vulnerabilities, which can be exploited to turn them into dDoS attack systems. This includes for example the network of 180,000 Internet Relay Chat (IRC) client systems, the Seti-at-home clients, or the millions of computers vulnerable to Visual Basic Script (VBS) exploits; where a possible compromise of the software agent converts all of these machines into dDoS attack handlers.

Previous dDoS Defense Mechanisms

To compare your newly installed CyberWarfare Defense Layer with the inefficient methods of prior attempts to defend against dDoS attacks, understand their background:

After the first widely reported dDoS attack in 2000 (initiated by a 17-year old Canadian, "Mafiaboy", who was caught only due to his bragging about the incident), solutions were sought to defend against the multiple vulnerabilities allowing dDoS attacks to be successful.

However, the technology hurdle to provide an effective solution has not been overcome. Traditional technologies have proven unsuccessful in creating a working defense (see below for traditional Intrusion Detection Systems, Quality-of-Service systems, Firewalls, and the section on Competition).

Applying Ingress/Egress filters on every router on the Internet (at the source) has proven to not be practical.

The ideal solution for a working dDoS defense addresses the root cause to determine valid ("good") traffic from all forms of dDoS (or "bad") attack traffic at the target site as much upstream as possible, react in real-time, and defend against both known and unknown forms of Denial-of-Service attacks.

Signature-based, Traditional Intrusion Detection Systems (IDS)

Traditional Intrusion Detection Systems ("IDS") are able to recognize some forms of TCP/IP, Denial-of-Service, and application attacks, if those attacks were observed before, and could be transformed into definitions (also known as "signatures") by which IDS compare the network traffic they observe against "known" forms of attacks. Not all attacks can be defined in a signature, and every attack not defined in a pattern match (or signature) is not recognized by a traditional IDS.

Updates to signature databases to address newly developed and recognized forms of attacks need to be done by qualified system administrators or by vendors and require the regular download and integration into each deployed Intrusion Detection System. This delay introduces an additional risk factor between the observation of a new attack, its definition as a signature file, and the successful download and implementation for comparison against live network traffic.

A traditional IDS is limited to only alert to a recognized, "known" attack, leaving it to other security systems to defend against the attack.

Often the time to compare traffic patterns against the signature database, determine a match, and send out the alert is longer than it takes for a dDoS attack to generate the damage. Most dDoS attack cause the damage within a few seconds, and sustain the damage until the attacker decides to stop the attack.

In June 2003, the Gartner group classified IDS as a technology with little value to contribute to computer and network security, creating a passionate discussion in the security community. Gartner predicts that by year-end 2003, 90 percent of intrusion detection system deployments will fail if false positives are not reduced by 90 percent.

IDS systems are, however, still seen as a factor in identifying known forms of irregular network traffic to justify IT budgets. While network troubleshooting tools ("sniffers") show all traffic without providing any analysis on their own, IDS systems match (known forms of) irregular network traffic – even though creating large logs which are difficult and time-consuming for administrators to interpret.

Quality-of-Service (QoS) Systems

The development of Quality-of-Service systems was geared to overcome drawbacks in the TCP/IP protocol for network-based applications, which require a consistent stream of data (for example, IP telephony or video-streaming applications). While other network protocols such as the Asynchronous Transmission Mode ("ATM") communication protocol have built-in provisions to allocate and reserve bandwidth for applications, this feature is not present in TCP/IP.

With the emerging Denial-of-Service attacks, an attempt was made to use the same QoS technology to throttle traffic observed in these attacks. This approach has not proven successful, as it effectively aids a DoS attack. The "bad" DoS traffic generates significantly more bandwidth usage than the regular, valid (or "good") traffic. While QoS contributes to some extent to keep the infrastructure connected by throttling traffic during DoS attacks on specific TCP ports, it practically "drowns out" the good traffic.

During a bandwidth-flooding DoS attack on web port 80 (HTTP), the "bad" traffic consumes 95% or more of the bandwidth to this port. If throttled, both "good" and "bad" traffic is treated equally by a QoS system, practically allowing almost none

of the "good" traffic through.

In a low-level, application-based DoS attack (such as "Synk4") or a TCP/IP protocol attack, a QoS system provides no value, as the overall bandwidth is not increased in measurable terms, necessary to trigger QoS settings.

Router Security: Access Control Lists (ACLs)

Routers direct traffic to the attached networks, and pass on traffic bound to other networks on the Internet. These routing functions are either static (passing packets to hard-configured IP number ranges within the networks attached to the router) or dynamic (using protocols such as BGP4, the dynamic routing protocol providing the redundant infrastructure of the Internet, or OSPF, by which a router learns and updates changing IP number ranges for connectivity to other routers).

Routers can add security to networks by applying filtering techniques such as the Ingress/Egress configuration to validate source addresses (though not implemented widely, as it does not provide many benefits for the owner of the router and slows down performance) and Access Control Lists ("ACLs"). While ACLs can provide some basic restrictions which IP ranges can connect to a router (thus decreasing the processing load on a firewall), they add a significant processing load on the router.

A router is typically not equipped with an abundance of processing power itself; as each incoming packet requires to be checked by the entire Access Control List, router performance is significantly reduced the longer the ACL becomes. Deploying routers with excessive CPU processing power for the network connections a router serves is a costly task defying business validation per connection, and can still be exceeded by Denial-of-Service attacks.

The number of TCP and UDP ports in use is limited. Among those, the most popular ports are 80 (HTTP, the "web" port, which is also the most attacked one), 443 (encrypted HTTP), 25 and 110 (sending and receiving e-mail), 22 (SSH for secure remote access), and 20 and 21 (FTP). On the UDP side, 53 for DNS and the ports used for Voice-over-IP (VoIP) are the most used, running critical services.

While only a fairly extensive ACL has the ability to provide more than basic security settings, it invalidates the function as a router. In addition, most TCP/IP protocol and Denial-of-Service attacks cannot be detected by routers and their Access Control Lists.

While dynamic routing protocols such as BGP4 and OSPF are necessary and critical for a redundant connectivity architecture within enterprise networks and the Internet itself, they are also often themselves the target for attacks. Due to the lack of authentication procedures, so far there is no known remedy for BGP4 poisoning attacks, which feed invalid data to the dynamic configuration of a BGP4 router, thus disabling the function.

"TCP Intercept" is another attempt to filter traffic at the router level (for example as an option in the Cisco IOS router operating system). Due to the nature of routers, this function requires the router to build a stateful table of intercepted connections, which creates a severe vulnerability. Attackers can saturate this table easily; the severity of the vulnerability outweighs the benefit of the feature.

Firewalls

Firewalls were the first steps towards improved security, and have undergone a long development cycle. The fundamental purpose of a firewall is to limit the availability of TCP and UDP connection ports from 65,535 each to those actually in use. These ports fall into three categories: "well known ports" are those from 0 through 1,023, "registered ports" are those from 1,024 through 49,151, and "dynamic" or "private ports" are those from 49,152 through 65,535.

The number of TCP and UDP ports in use is limited. Among those, the most popular ports are 80 (HTTP, the "web" port, which is also the most attacked one), 443 (encrypted HTTP), 25 and 110 (sending and receiving e-mail), 22 (SSH for secure remote access), and 20 and 21 (FTP).

A CSI/FBI study found 90% of all attacks bypass firewalls.

Firewalls limit the access to the ports allowed, and monitor both valid and invalid connection attempts, thus generating large logs. The content of those logs today is largely represented by low-level scans, which make is very difficult for administrators to find relevant entries to detect serious attempts to compromise a firewall or the infrastructure it protects, and thus remain largely unchecked.

The development of firewalls went from simple "port allow/deny" firewalls to "application proxy" firewalls and then to the "stateful inspection of established connections". None of these approaches provides a final solution; they represent an adaptation to the evolving threat levels.

An application proxy firewall is a software application that runs on a server between a network and the server (for example, an HTTP proxy for a Web

server). The proxy firewall runs an HTTP server that looks like the destination server when viewed externally; it looks like the requesting client browser when viewed internally. This "server in the middle" offers the security administrator an opportunity to establish rules about what type of traffic can be accepted or blocked.

A robust application proxy firewall would need to run an instance of every application that it defended, including Web servers, database servers, Internet relay chat, File Transfer Protocol, Telnet, e-mail, Morpheus, and an enterprise's custom applications. Application proxies failed to succeed in the firewall market because of the need to address the multitude of applications. In addition, invoking an application adds too much latency to enable network processing at wire speeds.

Stateful inspection firewalls face the challenge of throughput. Stateful inspection was developed to improve throughput by allowing a rule to be applied to the initiation of a session. A stateful inspection firewall looks at packet headers to make a decision based on a rule set. The decision to block or allow access is applied to all subsequent packets in the session (a session is defined by source, destination, protocol and a time factor). A stateful inspection firewall is a packet processor, which is a network device that can scale to the multi-gigabit speeds that are needed in Internet data centers and some enterprises.

However, application defenses require more awareness of payload content and the ability to inspect it at wire speeds. Web services will require high-speed parsing of XML and Simple Object Access Protocol (SOAP) objects. Applying security policies to these applications will require the ability perform 100 percent inspection of packet payloads.

The current trend in firewalls (starting in late 2002) is to develop In-Line Scanner devices ("deep packet inspection"), which should work in real-time (i.e. wire speed) and apply security policies based on application content, source, destination, and port; essentially the determination between "good" and "bad" packets by inspecting deeper into the packet stream.

Gartner predicts that by 2005, enterprises will no longer use software-based application proxy firewalls, and that by 2006, enterprises that rely on only proxy and stateful packet inspection will experience successful application-layer attacks at twice the rate of enterprises that use leading deep packet inspection approaches.

The leading commercial firewalls (Cisco PIX, formerly NTI; Checkpoint Firewall, Secure Computing, etc) do not incorporate this technology yet.

Web services are forcing perimeter defenses to be more aware of what types of traffic they allow to access the network via port 80 as code (such as SOAP elements) and control messages (such as XML statements).

In recent developments, Check Point Software Technologies' add-on product "SmartDefense" enables their product to look for common attacks and drop sessions that are associated with them. This is the same approach as used in IDS and Anti-Virus solutions, where "known" form of attacks are defined, downloaded to each firewall, and compared to observed traffic. "Unknown" (or new) forms of attacks are not recognized and defended against.

Previous Check Point add-on products (such as "SynDefender") made the problem worse. While adding some capabilities, a Check Point^[1] also equipped with "SynDefender" not only becomes unavailable for the duration of a "Synk4" attack, but needs to be completely re-installed, including license keys and a restore of the user-defined access rule base.

None of these solutions incorporate the ability to truly differentiate between "good" and "bad" packets, and do not provide a solution for an effective, real-time (i.e. wire-speed) defense against (distributed) Denial-of-Service attacks.

Firewalls remain vulnerable themselves against TCP/IP protocol and Denial-of-Service attacks – *see the section on "Is a firewall enough?"*

Melior CyberWarfare Defense Technology – The Product History

To understand the significance of the development of the product you just purchased, you may be curious about its origins:

Attacks in the past

Throughout the duration of Melior's Internet Security consulting business, it became apparent there was no available defense against Denial-of-Service attacks, or any protection against Penetration Testing probes. After the 1999 'Mafiaboy' attack, which shut down many prominent websites for several days, several companies set out to develop solutions against the dDoS threats, funded overall in excess of \$250 million, but it became increasingly obvious none of them came up with a new approach to solve the problem, but rather re-used already existing technology, such as bandwidth throttling/QoS, signature-based detection mechanisms, etc. - all of which were doomed to fail by design.

With the Internet Root Server attacks on October 21st and 22nd, 2002, which almost shut down the entire global Internet with a simple ICMP dDoS attack (had they lasted longer than the 5-hour DNS cache expiration), it became clear that something had to be done. The Internet had evolved to be a key communications network for Governments, businesses, and manufacturing; and networks started to face threats from internally compromised systems.

Melior's developments

Melior decided to start the development of a CyberWarfare Defense product line, and identified an interesting technology approach in Germany, which seemed to provide a head-start. Ultimately, this technology did not live up to expectations either, and the patents filed for it turned out to be fake.

Melior continued the development of a product as an inline-scanner, sitting undetectably "on the wire" in front of the protected infrastructure or at critical internetworking connection points, inspecting all traffic flowing through the system and deciding on the fly, which of this traffic is valid, and which constitutes either a Denial-of-Service condition, a Penetration Testing probe, or simply invalid 'junk' traffic.

The solution was to pass on all valid traffic and discard the invalid packets, thus providing uptime, keeping the cost of dDoS attack traffic down, and eliminating huge log sizes on the firewalls and intrusion detection systems behind the system, so IT and security administrators can find and act upon critical log

entries.

Throughout the development process towards the final product - now realized with Barbican RNP -, Melior kept the public informed about the goal and the R&D progress, and was approached several times to help out; dDoS attacks became more frequent, organized crime took over from the amateurish initial attacks for extortion attempts, and increasingly, dDoS attacks kept shutting down large e-commerce businesses and payment processing providers.

In summer 2003, Melior deployed a pilot system for a popular website at a US Internet Service Provider (ISP), for which those first, TCP-only code versions worked to successfully defend against the attacks, thus keeping the web site alive.



First "iSecure" pilot system

Later during the same summer 2003, Microsoft was the victim of a coordinated attack with the MS Blaster worm, which threatened to cause a Denial-of-Service attack against the company; in defense, Microsoft moved their online presence to Linux-hosted Akamai at great expense, and found that the traffic to download the MS Blaster patch almost exhausted Akamai's resources. The subsequent dDoS attack would likely have been successful in terminating the systems serving the 'Windows Update' domain name. Melior offered the prototype dDoS solution to Microsoft, but the company decided to abandon the domain name instead.

In fall 2003 the Spammers, by now mostly organized crime groups, contracted several virus authors, and launched a coordinated attack against the mostly non-profit anti-Spam blacklists. All six non-profit blacklists went instantly down, and the only commercial one moved to Akamai at great expense, yet was still taken down. Melior was approached by the non-profits, and provided Spamhaus with hosting, a prototype, and hands-on implementation of other code in development; units were also shipped to SORBS & SPEWS in Brisbane/Australia and the US,

while Melior engaged also in other activities to defend them against the dDOS attacks.




Second revision of "iSecure" as used in the dDoS Defenses against the anti-Spam blacklists Spamhaus, Sorbs, and Spews

The Spamhaus Project

As a result, Spamhaus, SORBS, and SPEWS survived, while other blacklists, among them Osirusoft in New Zealand and monkeys.net disappeared from the Internet forever. The combined user base of the three surviving blacklists (Spamhaus alone with 260 million, serving many multi-national governments and corporations) continued to be able to eliminate about 40% of incoming spam through Melior's efforts in the middle of the development cycle.

The direct involvement with Spamhaus gave Melior hands-on experience from a user perspective, and allowed to follow the progress on various dDoS attack attempts in a real-world scenario:



(External attack traffic against the Spamhaus Web site)

Continuously, new dDoS attacks were launched against the Spamhaus web site, spiking at 100 Mbit/s; the attackers realized those were without effect, seized the attacks after a day or so, and tried a different method, only to find it did not work again. After several months, they gave up.

Other prototypes were installed at various locations, providing feedback to R&D during the development process.

In February 2004, Melior introduced the productized version of the first, TCP-only code under the label 'iSecure'.



Standard 1U „iSecure“



HA-2U version for Fail-Safe Configurations

The general public saw the product first at the RSA Show in San Francisco.

In June 2004, our reseller in Japan, DTC, introduced iSecure at Network + InterOp in Tokyo:

Now, Melior finally introduces the CyberWarfare Defense Layer with the full product, which bears the permanent name for the product line: Barbican RNP. With daily news of increasingly sophisticated dDoS attacks, and greater damages by each attack, the threat of Denial-of-Service has peaked as the Number 1 threat, and the product could not come at a better time.



Barbican RNP v1.0

Melior will continue to expand the product line with additional modules, for other bandwidth requirements, telecom-grade fail-over solutions, and products, which address other threats creating Denial-of-Service conditions, such as eDoS (aka: Spam).

Trouble Shooting

Barbican RNP is a „Plug & Protect“ device, which cannot be compromised.

If you followed the instructions in the Quick Start Guide, there is very little, which can go wrong – unless a physical malfunction occurs.

Hardware Failure – Replacement Process

Prior to shipment to you, all Barbican RNP systems undergo a rigorous hardware and „burn-in“ process.

Each unit has only very few moving parts (no hard disk drive, only cooling fans for the CPU and the power supply) to decrease the risk of a hardware failure.

In the event of a malfunction (failure), you are covered for three (3) years or as long as you maintain your service & maintenance contract with Melior, Inc., for replacement units.

Please contact Melior for an RMA (Return of Merchandise Authorization).

Melior will quickly provide you with the RMA, and provide you with the address to which to return the unit (your cost); within 24 hours upon receipt of the malfunctioning unit (faster if you chose the maintenance upgrade) you will receive a new, replacement Barbican RNP unit directly from the manufacturer.

Melior, Inc. CyberWarfare Defense Solutions is proud to provide you with Barbican RNP and help you protect your online presence, network and systems security with this new security technology!

If you wish to demonstrate your commitment and support for the integration of a new CyberWarfare Defense Layer, Melior makes its „no dDoS“ logo available:



(see www.dDoS.com for URL links)

January 2005

Melior, Inc.
CyberWarfare Defense Solutions

1501 Beaumont Street
Columbus A. Langley Building
Dallas, Texas 75215
USA

www.dDOS.com

Telephone: +1-214-421-5975 (or toll-free in the USA: 888-4MELIOR)
Facsimile: +1-214-421-5951 (or toll-free in the USA: 888-TOFAXUS)