GIAC
CERTIFICATIONS

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at http://www.giac.org/registration/gsec

Raptor Firewall 6.5 Running on NT 4.0
By Dennis Carter

Introduction:
   This paper is to help with customization of Raptor Firewall 6.5 Running on Windows NT 4.0. There were some major changes made in the 6.5 release from 6.02. I will get to them next. This paper I hope will help you to set up Raptor so you can get the best performance possible and setting in place some guide lines to help manage the many changes that take place in the Raptor product.

Raptor Firewall.

   Raptor Firewall is an application-level firewall based on application proxy server technology. Raptor also provides a packet filter engine on each network interface, local and remote administration, and strong user authentication with real-time monitoring. Raptor uses application–level processes (known as proxies) which only allows those services for which there is a proxy and all other services are blocked. One of the benefits of using proxies is that the protocol can be monitored. Raptor's easy to use windows GUI interface makes it easy to configure.

Raptor 6.02 VS Raptor 6.5

   The following were some major changes that took place in Raptor 6.5. One of the major changes was Secure Tunnels. In Raptor 6.02 you could create a secure tunnel for two local entities or between a remote and local entity. The services and protocols between the two did not matter, as all services and protocols were available. The changes in Raptor 6.5 change how secure tunnels are used. Tunnels between two local entities are now local tunnels. For a local tunnel you have to create a filter and specify the protocol and direction matched between two network entities. I will discuss setting filters later in this paper. In version 6.02, for remote tunnels you had to specify local and remote entity endpoints, local and remote gateways and the VPN policy being used.
   Another change in 6.5 is specifying the interface being used by local or remote entities when creating rules. In Raptor 6.02 you didn't specify interfaces in rules for the two entities you were creating the rule for. In 6.5 you have to specify the interfaces the two entities will be using. If you specify the wrong interface the entities will never get through the firewall.

Creating Entities

   When creating entities it is best to have some guidelines set up so that an entity does not show up two or three times with different rules assigned to that host. You can set entities up as a host, Domain, Subnet, Group, RaptorMobile, Security Gateway or a Workgroup. Where I work at we have about 5000 employees. Putting in that many entities can be over whelming and if an entity shows up two or three times with the same address under different names can cause overuse of needed resources by the computer

system, and cause confusion when trying to research a problem. If everyone is going to have Internet access then create three or four groups and create subnet entities. Depending on the size of your Network you will want more than one group, this will help balance the load as far as the firewall trying to read the entity database and the rule database to see if an entity is allowed Internet access.

Individual host entries are good to use when you are using VPN, RAS (Remote access system) or have outside consultants using the dialup system. The point here is that we have 3 information security people who add entities, rules, and setup remote security through our firewall; so it is important that each us add these entries to the firewall database the same way. You never know when one of us is going to be out of the office and have to trouble shoot why something is not working.

When creating individual host entities we use firstname_lastname if the host is using VPN or dialup connection we will add a VPN suffix. For instance John_Doe_VPN. We have different subnets for VPN and Dialup users, so when we are working on a problem we know what subnet they are on. In the description field of entity screen we have what department and their extension number. If the host is a consultant for one of our departments we use firstname_lastname_companyname. Ex: John_doe_fixit. In the description we will have the name and phone number of the consultant or someone in the department that hired the consultant, name and number as contact.

At present we do not use the Domain entity. The reason for that is we have one firewall and to create a Domain entity you also have to activate reverse lookup. Reverse lookup uses a lot of system resources and cause our firewall to run very slow. Unless you have several firewalls or a small enterprise network I would not use this feature.

We added a subnet entry for each of our segments. This works well when you want to create groups or you want to allow a whole segment access to a remote site. This way you don't have to add each person on the segment as an individual host and write rules for each one.

We created group entries so we group individual hosts or subnets together and apply one rule to the group. This works will if you allow everyone to access the Internet. You can create one rule add all the subnets as members. If your group is going to be very large let say a 1000 subnets you may want to break this down into more than one group. Depending on your IP address class you could easy have three or four times that many hosts.

Another use for groups is when we have more than one consultant from a company doing work for us. We create host entries then create a group using their company name and add the individual hosts as members. By doing this it allows us to control their access to our network. At present we do not use RaptorMobile, Security Gateway or Workgroups.

Creating Rules.

In the introduction I stated that in Raptor 6.02 defining the interfaces really did not matter but in Raptor 6.5 you must identify the "coming in" and "going out" interface or the rule will not work. In the rules window under services do not use the all* services option.

We have found this to also cause a slow down in our firewall when this was used for one of our rules. One of the reasons is that the all* will allow all the services in the protocol list to be used which may not be appropriate for all connections. By using all* in a rule drove our computer utilization up to the 70%-80% level. So in our guidelines this is a "no no". Some of the services like HTTP can be configured. We always configure HTTP to allow SSL over standard ports (443,563). We set this in all rules using HTTP services so that someone using SSL will not have a problem. Also in the HTTP dialog window you can add additional ports numbers when they are needed to make a connection to a remote site. We do not add additional ports here instead we add those ports to the HTTP demon. That way the additional port number is available to everyone instead of just the one rule.

Services and protocols.

Raptor 6.5 comes with a set of services and protocols to be used when creating rules. In our guidelines, before a service can be added, we have to verify that the port number is not being used by another service or one of Raptor's demons, like HTTPS. If the port number is in use by a service, we allow that service in the rule. If we create a new service the name of the service will match to the corresponding host, group or company that the service is going to be used by – i.e. Fixit Company. For example the new service is called fixit_company, the protocol is TCP, the destination port is 11000 and the source port range is 1024-65535. The reason to set this up this way is so whenever the Fixit Company is done we will remove this service. Otherwise you may have open ports for services that are not being used anymore. Make sure when you create a new service that you check the box "display in rule window". If you don't the new service will not be displayed when you go to create a rule to use the new service. You will also need to select the proper protocol for the service you created such as TCP, UDP or IP.

Using Time Ranges

Raptor 6.5 allows you to set up time ranges for existing rules. When we have an outside consultant working for us we explain to them that they will be allowed to access our system from 8:00 a.m. to 5:00 p.m. If they need to stay late or want to start early or need access over the weekend they have to call us and we will set that up. We also explain this to a director or manager who hired the consultant. You don't want to leave your system open after hours or on the weekend when nobody is there to monitor what going on in your network. There are time range templates ready to use or you can create a custom template to be used. We normally will create a custom template using the company name. This way, when someone calls in, we can locate the time template for that company. Remember we have three information security personnel that could get a call to make changes to the time template so the consultant's can work. You can also use time ranges to block people who work for your company and access your network from a VPN or Dial up connection to gain access to the Internet.

Redirect Services

We have a lot of web servers that are used by the public as well as company employees. To protect the actual IP address of these web servers or other systems we use redirect services. If you have a Virtual IP address on the same subnet as the Raptor system's address, Raptor will automatically take care of routing ARP (Address Resolution Protocol). If not, you can add a static route in your router configuration, specifying that services destined for the virtual address be sent to the Raptor system. We use the redirect whenever we have a web server or System that is going to be accessed from the outside and we want to protect the actual address of that system. You will find Redirect Services in the access control folder. Ex: the system you want to protect has an address of 10.10.1.4; you have a virtual address of 10.10.1.6. To set this up we use service HTTP, The requested address is 10.10.1.6 (virtual address) address mask is 255.255.255.255 and the redirected address is 10.10.1.4. So who ever accesses 10.10.1.4 will only see 10.10.1.6 as an actual address for that system.

Using Filters

Earlier in this article we mentioned that Local tunnels were replaced with filters. If you converted your Raptor 6.02 database to Raptor 6.5 all your local tunnels were upgraded to one forwarding filter sequence called 6.5forwardingfilter. One thing to note here is that a forwarding filter provides NO security over the Internet and is not as secure as proxies (rules). Using filters is another way to take the load off your firewall and provide some security. Forward Filters allows us to fine-tune our security between two entities, by controlling the services and protocols used and who can initiate the connection. When a packet matches a chosen filter, it is not sent up the protocol stack for authentication. It is allowed through the Raptor system, bypassing normal security checks. A good guideline for creating Forward filters is to use a subnet or groups whenever possible this keeps the number of filters you have to keep track of to a minimum. If that is not possible, then you will have to add the filter by individual host. Here is an example of how a forward filter works. Let's say you have a host (10.10.1.4) that needs to FTP to a host going through different interface of the firewall (10.10.24.2). The 10.10.1.4 represent your VPN segment and the 10.10.24.2 represent a system behind the firewall. Use the name of the from host" and the name of the to "host" as a title for your forward filter. Ex. John Doe VPN to Mainframe. Then under services and protocols in the filter window we select A>B FTP. This opens a one directional connection from 10.10.1.4 to 10.10.24.2 and only allows the FTP protocol to be passed between entity A and entity B. If any other protocol or service is tried, or if host B tries to connect to host A, it will be denied. So forward filters can be very powerful when used and providing excellent through put.

Another note to make here is do not select A>B all and B>A all from the protocol and services allow list when setting up a filter. Doing this will allow all the services and protocols between the two entities to pass.

Log files.

As a general guideline we check our log files daily. We keep at least five days worth of log files. In the rules window under miscellaneous there is a box you can check that says, "Log normal activity". We do not allow logging of individual rules because of the amount of disk space this would take up. The only time we turn logging on for a specific rule is when we are trouble shooting a problem. The Raptor firewall system will log any alerts, warnings and any attempted break-ins. So individual rule logging is not necessary. It is important to keep and eye on your log files and delete then periodically.

Address Transforms

By default the Raptor system overwrites packets with its own address for outgoing connections. Whenever a computer system sitting behind the firewall makes a connection through Raptor's interface, Raptor will change the source address on the packet. This is done so the client's IP address is not revealed. If the client is making the connection using a secure tunnel then the source address is untouched, revealing the client IP address. The address transform can help when you are having problems trying to make a connection between entities. There are situations when an address transform is going to be needed. 1) whenever the raptor system is the default route for servers behind the firewall that want to see original address of the connecting clients. 2) When the Raptor firewall is not the default route and the client making a connection to the Raptor system is unable to route the packet to it final destination because of seeing the Raptor system IP address. 3) If you are doing static one-to-one mapping of addresses on your network in order to conceal addresses or to handle problems with address overlapping. 4) When you are using Network Address Translation (NAT), a poo
 of addresses that are designated as a replacement addresses for a client address.
 Example; entity A needs to connect to entity B. But entity B will only accept entity A's original assigned IP address. You would go into the Address Transform property page and select "Use Original Client Address". Then select entity A for the "from client" and entity B as the "to server". Remember you have to also select the "coming in interface" and the "going out interface". If you put the wrong interface for any one of the entities the Address Transform will not work.

Conclusion

To make Raptor an effective firewall in your company you need to have a good security policy. Set some guidelines and naming conventions on now information is to be entered in the different configuration windows. Create groups and add subnets as entities so that you don't have to add very person. By creating groups you can limit the number of rules you have to create. Setting configuration standards not only makes troubleshooting easier, but can help to create a more secure network.

References:

Internet References:

  http://www.symantec.com/

  http://enterprise.cnet.com/enterprise/0-9567-707-3748534.html

  http://www.techreviews.com/sections/topReviews/article/TT20010410S0014

  http://www.complus-arg.com.ar/Press_Releases/New_Raptor_Firewall.htm

  http://www.firetower.com/

Book Reference:

  Raptor Firewall and PowerVPN 6.5 Configuration guide for NT.