Proceedings of the 1st Conference on Network Administration

Santa Clara, California, USA, April 7-10, 1999

Tricks You Can Do if Your Firewall Is a Bridge

Thomas A. Limoncelli
Lucent Technologies, Bell Labs

Tricks You Can Do If Your Firewall Is a Bridge

Thomas A. Limoncelli
Lucent Technologies, Bell Labs
Murray Hill, NJ, 07974
http://www.bell-labs.com/user/tal
tal@bell-labs.com

Abstract

Firewalls that forward packets like a bridge, rather than as a router, have many operational benefits. By decoupling routing from filtering, the firewall becomes a pure filter, unburdened by routing table or interface configuration. The result is increased flexibility. This paper explores some of the benefits we have found. Most of the benefits stem from the fact that a bridged firewall requires fewer transit subnets. Sometimes transit subnets are completely eliminated. It can be placed between any two network devices and act like a line filter without needing to change the logical routing of the network. It is easy to put one in series with another firewall for testing. Our examples include replacing an old firewall with a new one, moving a firewall from one router to another with zero downtime, firewalling off an individual office or lab, and others. In many cases topology changes are made without service interruptions. The operational procedures become much more simple. The paper also suggests future directions for research in this area.

1 Introduction

Firewalls filter packets by sitting between two network points and deciding whether or not to pass each packet based on a set of rules. Sitting between two points requires the device to pass packets like a router (a Layer 3 device) or a bridge (a Layer 2 device). The fact that a firewall is bridge-like or router-like is not usually emphasized by vendors because the same firewall features can be provided either way. However, we have found interesting operational advantages to bridge-like firewalls.

2 Background and history

A firewall is a device that filters TCP/IP packets based on a set of rules. As each packet passes through the system, the rules are processed to determine a "pass" (forward the packet) or "no pass" (drop the packet) decision. Depending on the security policy required, different kinds of rules may be constructed. For a general discussion about firewalls refer to [Cheswick94] or [Chapman].

Many of the early firewalls were routers modified to perform certain filtering functions. As firewalls became more complicated, vendors began using UNIX workstations with two network interfaces. The workstation would be configured to route packets between its interfaces and in some cases would run a modified kernel and software that was capable of performing some kind of filtering or proxying functionality. Later "standalone firewalls" became popular because of their simplicity and the advantages that network "appliances" have.

In these early generations, the device connects two IP networks and therefore must also contain routing functionality to know how to forward packets. They must be configured with routing tables that describe which IP subnets are where. These routing tables are usually configured statically rather that relying on potentially insecure dynamic protocols.

A few new firewalls such as the Lucent Managed Firewall [LMF] and others forward packets like a bridge. That is, they act as a learning bridge between two devices. These two devices are usually routers which are much more capable of performing the complicated routing tasks required in modern networks. By decoupling routing and filtering into separate boxes, each can focus on one task and

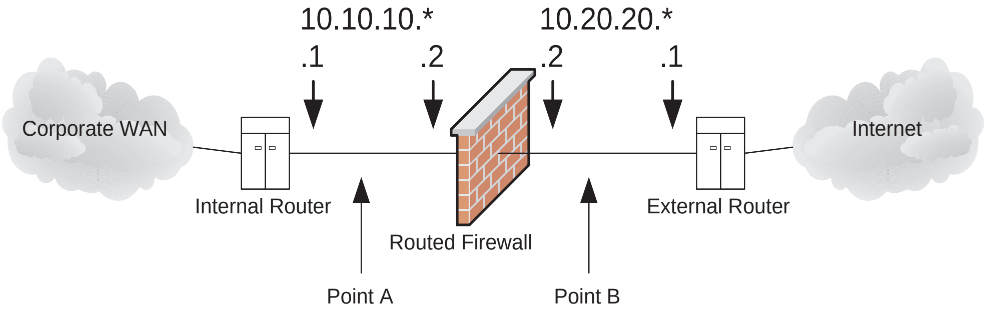

Figure 1: Router-like firewalls require two transit IP subnets.

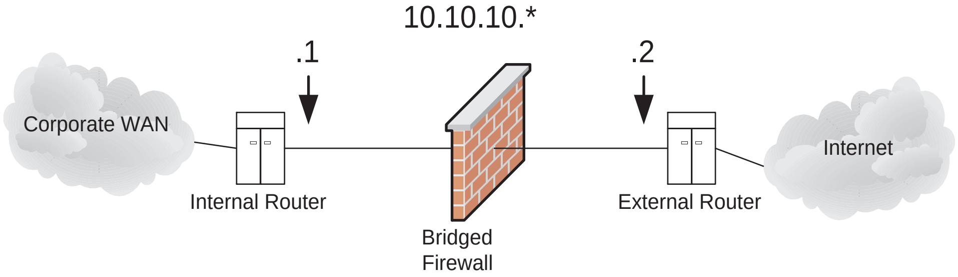

Figure 2: Bridge-like firewalls require no additional transit IP subnets.

do a better job of it. The firewall still bases its "pass/no pass" decisions for each packet based on the Layer 3 information within the packet. Even though the term "bridge" is used, it does not make its decisions based on the L2 ("MAC") addresses in the packet. If we were to compare the ruleset of a bridged firewall vs. a routed firewall we should find the exact same rules if they implement the same policy.

I will use the phrase "routed firewall" to mean a firewall that passes packets the same way that a router does, as a L3 device. That is, it routes packets between two IP subnets. It is not necessarily meant to imply that an off-the-shelf router being used to filter packets. Similarly, I will use the phrase "bridged firewall" when I mean a firewall that passes packets like a bridge or other L2 device, not necessarily an off-the-shelf bridge that has been configured to filter packets.

3 Transit subnets

Most of the techniques in this paper rely on the fact that a bridged firewall can be inserted between two points without changing the IP routing topology of the network. Without a firewall it only takes one "transit" subnet to connect two devices. To insert a router between two network elements one must use an additional transit subnet.

Figure 1 shows that two small transit subnets are required when a routed firewall is between two routers. Figure 2 shows how only one transit subnet is required when a bridge-like firewall is used.

The first benefit of reducing the number of transit subnets is a reduction in the use of IP address space. Allocating one fewer subnet allows you to use that address space elsewhere. Conserving IP address space is important. The second benefit is less obvious. When adding an additional transit subnet one must reconfigure one device to use the new subnet. This interface configuration and the routing table adjustments is an additional burden on the installer. With a bridged firewall the two network devices need not be reconfigured. This latter benefit is the basis for much of what is described in the remainder of this paper.

4 Trick: Ease of deployment

Replacing our old (routed) firewall with our new LMF firewall was made considerably easier because the LMF is a bridged firewall. Because the new firewall was a bridge, it was programmed with its filters and then was tested in various positions in our network with limited disruptions. Backing out would have been quick and easy if problems had been discovered.

Our previous firewall was named "`stile`."[1] It was a prototype that used stateful inspection filtering[2] like the firewall that was replacing it. However, it was a routed firewall.

We were very cautious when we deployed the LMF. It was a prototype and not even an announced product at the time. We were testing it in an environment with hundreds of users that would be very unhappy to lose Internet service.

The initial test of the LMF prototype was done by putting it in series with the old firewall. Our firewall network configuration was similar to Figure 1. First the LMF was programmed with the same filter rules as `stile`. The LMF was then installed behind the old firewall, at Point A in Figure 1. Now all traffic was filtered twice, but if any problems were discovered with the LMF, the problem would be inside the zone protected by the firewall.

Since the same developers had created both prototypes, the logs that each generated were extremely similar, or similar enough that simple pre-processing via `awk` should produce output that could be compared with utilities such as `diff`. Intrusion attempts would be seen in the `stile` logs but not the LMF prototype logs, otherwise they should be the same. The LMF prototype's logs should not include any entries that were not in `stile`'s logs. If this did happen, it would be a sign that the LMF was not following the rules the same way. The LMF prototype logs should also not be missing any entries from `stile` with the exception previously mentioned.

Installing a bridge is a matter of two quick cable changes that takes seconds. No interfaces needed to be reconfigured and no routing tables had to be adjusted. If the LMF were a routing firewall, inserting it into the path would have taken more effort and would have involved a longer outage. This would have slowed down insertion and (more importantly) the possible removal that would have followed if we discovered the LMF prototype wasn't working properly.

Turning a multi-minute outage into a few short seconds is important because users did not feel an outage. A 3-second pause when accessing the Internet is not noticed by most users. Modifying router configurations to insert the firewall would have been a serious disruption.

Similar time would be required if the device had to be removed. If a problem was discovered, being able to rapidly remove the device is appreciated.

Once validation was completed, it was time for a more serious test. We moved the LMF prototype outside the old firewall, at Point B in Figure 1.

We did similar log comparisons. This time, only the LMF logs should show any intrusion attempts and the remaining log entries should be similar (after pre-processing).

Moving the LMF to position B was also a matter of a few quick cable changes. The same LMF, with no changes to its configuration, was used. No routing tables had to be changed on the LMF, `stile`, or the other routers. Again, if these tests found problems with the prototype, it could be removed quickly. In fact, it could even be removed by operational staff with minimal training rather than requiring network or security administrators.

Once the LMF was validated, removing `stile` was relatively time consuming because it involved removing a transit subnet and the making the appropriate routing table and interface updates. Luckily that would be the last time we would have to deal with a routed firewall!

Later when the LMF prototype was replaced with the actual LMF product, we repeated the same process of placing the new firewall behind, then in front of, its predecessor. However, now removal of the old model didn't require any routing changes because we were replacing a bridge with a bridge. Removing the prototype was a snap.

[1] a stile is a ladder or set of steps that goes over a fence.

[2] Stateful inspection is a filtering technique where packets are examined in the context of the entire session. For example, rather than permitting all telnet packets through, the firewall maintains a list of which telnet session have been initiated, and only permits telnet packets associated with those sessions. For more details, see [Chapman].

Now we have enough confidence in the product to reduce the testing. As we receive new versions of the product we only test it by installing it outside of the current firewall. This saves a lot of time. We currently have two complete sets of hardware. One tests the "even releases" and the other tests the "odd releases." As a result, we can switch back and forth rapidly with very little disruption. With a routed firewall the interruptions would be much larger and we would lose our flexibility.

It should be noted that sometimes the LMF must be rebooted when it is moved. As a learning bridge, it learns which ethernet (MAC) addresses are on each side. If a MAC address moves from one interface to the other, it is blocked. The theory here is that machines usually don't physically move around your network. The LMF does not time out a MAC address like a real bridge, the only time the table is cleared is when the device is rebooted. While none of the steps described in this paper require a reboot, other experiments we performed did require a reboot. If you plan on doing your own experiments please remember this. There are other bridged firewalls on the market that may or may not work the same way. Contact your vendor for more information.

5 Trick: Moving to a new router

We were able to move the firewall from one router to another without any downtime, without any changes to the firewall, and (until the very end), without bothering the owner of the first router.

Figure 3 is a simplified diagram of how our routers and firewall are connected for our LAN. Our LAN involves a cluster of routers connected at a central backbone. Each router serves a separate customer group and in many cases is owned and controlled by that customer group. The firewall is attached to one of the routers. This means that one group was slightly "closer" to the Internet than others.

Each of our internal routers contains a dynamic routing table that contains all the subnets ("prefixes") within our corporation. The default route on each router points towards the firewall (either to the external router or to the router attached to the firewall). The assumption is that if a destination host is unreachable via our internal route tables, it

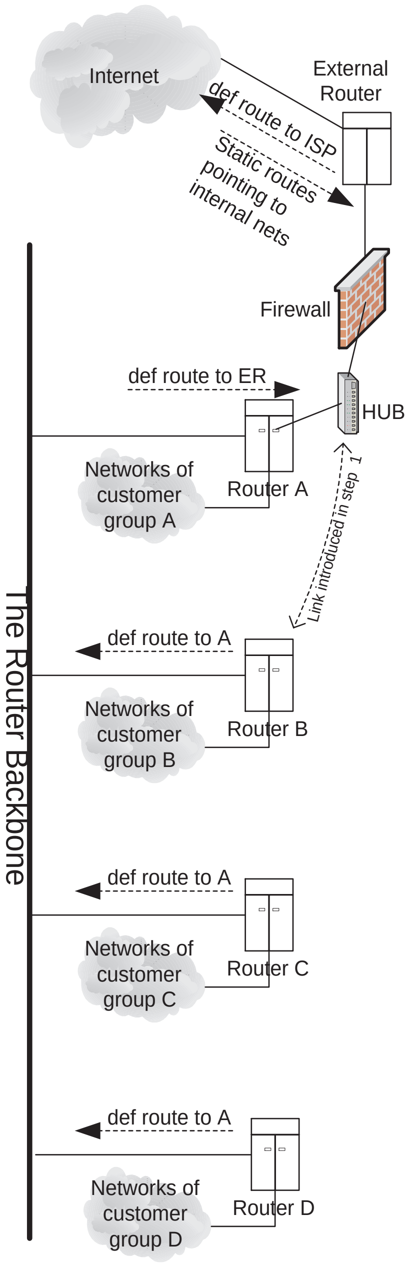
Internet
External Router
def route to ISP
Static routes pointing to internal nets
Firewall
def route to ER
HUB
Link introduced in step 1
The Router Backbone
Networks of customer group A
Router A
def route to A
Networks of customer group B
Router B
def route to A
Networks of customer group C
Router C
def route to A
Networks of customer group D
Router D

Figure 3: Our router inter-connections

must be on the open Internet. This is more efficient than having our internal routers maintain the entire corporate routing table as well as the Internet's route list. In Figure 3, the firewall is connected to Router A. Router A's default route points to the external router. The remaining routers' default routes point to Router A.

Due to operational issues, we needed to move the firewall off Router A onto Router B.

The move consisted of a sequence of carefully planned steps. After each step, we ran a battery of tests to verify our work. The tests consisted of sending packets from a host in each customer group to a number of places on the internal network, our extranet, and on the open Internet. The test was automated into a simple shell script that completed in a matter of seconds.

One more point must be understood by the reader before continuing. IP routes can be asymmetrical. That is, the path a packet takes to get from host X to Y may be different than the return path. The difference in the two paths might be slight or huge. The hosts do not care. During the following steps, there will be times when outgoing packets take a different route than incoming packets. Keep this in mind.

The first step was to achieve physical connectivity between Router B and the hub. Since these devices were in very different parts of the building, this would verify a series of cables, jumpers, and wire-fiber adaptors. Once physical connectivity was achieved, we configured the interface on Router B and verified that we could ping from Router B to Router A. The battery of tests was run as a baseline. At this point nothing should have been different. Packets from Router B to the Internet would still go through the Router Backbone, to Router A, through the firewall, to the external router.

The second step was to have outgoing packets from Router B go to the external router. Incoming packets would still come in via Router A. This step was achieved by changing the default route on Router B to point to the External Router (though the firewall). We re-ran the battery of tests and everything was fine. There was no service interruption due to this step.

At this point we took a short break. We waited to see if any users in Customer Group B reported any problems. Would they notice that their outgoing packets were taking one fewer hop than incoming packets? Nobody noticed anything and no service problems were reported. We could continue.

The third step was to configure more routers to send their outgoing packets via Router B. We changed the default route on Router C and D. We continued running our battery of tests. Again, no problems, no service interruptions, and nobody noticed that their packets were going out via Router B but returning via Router A.

Step four focused on the incoming packets. The External Router has a static route for every internal network, each pointing to Router A. All other networks are routed towards our ISPs. In this step, we change the static routes to direct packets to Router B instead. By doing this, we now had symmetrical routes for all the customer groups except for group A which is still sending packets out directly to the External Route via the firewall. Again, the tests were run. Again no service interruptions were caused by this step.

The fifth and final step was to contact the owner of Router A and request that the default route be changed to point to Router B so that it is the same as all the others routers (except Router B itself). The owner was incredulous at what we had done but agreed to make the one change required of him. Once the default route of Router A was pointing at Router B, all traffic was flowing symmetrically and was going to the firewall via Router B.

The final step was to disconnect the cable between Router A and the hub. Before we did this we ran various `traceroute`s to make sure we were disconnecting a cable that had no active traffic.

These changes were made without scheduled downtime. At some point packets were traveling in asymmetric paths, but users wouldn't notice this. Also, the stateful inspection process of the firewall was not disturbed by our live changes.

Table 1 depicts the route table settings after each step of the process. The columns for Router A–D show where the default route pointed. The last column shows how the External Router's static routes for internal subnets were directed. An asterisk ("*") denotes that packets for that customer group were routed asymmetrically at this stage.

	A	B	C	D	ER
At start/Step 1	ER	A	A	A	A
After Step 2	ER	ER*	A	A	A
After Step 3	ER	ER*	B*	B*	A
After Step 4	ER*	ER	B	B	B
After Step 5	B	ER	B	B	B

Table 1: Route table settings after each step.

This technique could have been done with a routed firewall but it would have required service interruptions and would have been much more labor intensive, as each step would have required routing table and interface configuration changes.

6 Trick: Firewalling off my office

With a firewall that is a bridge, I can firewall off my office without the network administrators knowing or caring. I simply place it between the network jack in my office and my workstation. Because it is a passive filter, no routing changes are required.

If my firewall was a router, I would have to request a dedicated subnet, have a dedicated router port with dedicated connection to my office, deal with routing issues, etc. It wouldn't be nearly as fun.

We often isolate an office with a learning bridge or etherswitch to provide more bandwidth. Simply connect one port of the bridge to your office's network jack and another port to what used to be plugged into the network jack. Doing this with a bridged firewall is very similar, except this bridge can have a complicated security policy. If I use a ethernet hub, I can even have more than one machine in my office. This creates a many-to-many relationship within the same subnet. This can not be done with a routed firewall.

The security policy that I used was quite simple. TCP connections from my machines could connect to anyplace else. Incoming TCP connections could only come from certain machines on certain ports. For example, I decided that only one particular machine could `telnet` to me, and a larger list of machines could reach me via `ssh`. UDP protocols such as DNS and the like had to be decided on a per-protocol basis. I started with a very strict security policy and added exceptions as I discovered what services I needed (i.e. what broke due to my strict security policy). Reviewing the policy violation log on the firewall was helpful to find out what protocols I needed to add to the policy to regain functionality.

This was much easier than constructing the policies for the firewall that connects us to the Internet. For example, there was no need to set up internal and external SMTP gateways to funnel mail though known "safe" software. DNS did not need to be passed though something like `dnsproxy` [Cheswick96] as I was fairly confident in the quality of the DNS data I was receiving. (And if I had been wrong, chances are other machines would fall to the attack before mine). Another difference is that one would never pass NFS through a firewall to the open Internet. However, in this scenario, I could permit my machines to mount from file servers outside my office and live with the risks.

Since I can have more than one machine on either side of a bridge, this technique could also be used to easily firewall off an entire cluster of machines, say, in a lab.

I could also develop rules that were a little more sneaky. For example, I could hide my machines from the subset of machines used by my boss. I could also configure the firewall to log, but not reject, sessions from particular groups of machines if I suspect they are probing me. However, since my firewall is transparent, they would not be able to see that I am monitoring for their packets.

The real benefit gained from the firewall being a bridge is that I was able to do this without any awareness or adding to the workload of our network administrators. If this were a routed firewall, connecting the dedicated transit network to my office would have involved a long wait as changes were made in wiring closets, etc.

7 Trick: Hiding machines

Inspired by the ability to do the previous trick without the involvement or knowledge of the network administrator[3] I thought there might be some other interesting things I could do without the network administrator's knowledge. For example, I could steal IP addresses. Normally using the IP address of your own choosing is dangerous and the author

[3] I'm exaggerating. I am the network administrator!

recommends you only use IP addresses allocated by a centralized authority. If you use a random IP address that is later allocated to someone else on the same subnet, the address collision will make both machines unusable.

However, with a bridged firewall, filters can be constructed so that nobody outside your firewall will ever see packets from machines using the stolen IP addresses. Of course, this means that these machines can not talk with any servers outside of your firewall, or even the default route for the network.

The DNS data for your hidden machines would be that of the hosts whose stolen IP addresses you are using. At least one machine would have to use a non-stolen IP address so that it could gateway or relay for the others. This machine would have to provide services such as DNS, SMTP, DHCP, web proxy, and other protocols. Clients could conceivably have all the services they need either locally (compute and file service) or via proxies on the gateway machines. In fact, with `dnsproxy` you could even fake the DNS data so that the machines inside your little kingdom see the DNS data that you want them to see.

Those proxy services aren't needed if you do not interact much with the outside world. Then again, if you don't need connectivity to the rest of the world, you don't need a firewall, you can just pull the plug. That doesn't require any intervention from your network administrator either. In this case one may want to use the special IP space reserved in [RFC1597].

This is a lot of work to go through to hide a couple machines. Certainly this is more work than getting a transit network connection set up to your router. Unless you have a real vendetta against your network administrator or are James Bond, I don't think this is very useful besides an interesting theoretical discussion.

8 Trick: Firewalling off a lab

Firewalling off an office is interesting but in the future we plan on taking this idea even further. If we can firewall off an office, can't we firewall off an entire division?

The challenge here is that a large group of machines usually involves many subnets. Things become easier if one can pass a routing protocol through a firewall. A routed firewall could not pass a routing protocol like OSPF without re-implementing the entire routing protocol. That would be very difficult.

A bridged firewall could do this more easily by just blindly passing all packets from your chosen routing protocol and letting the surrounding routers process the packets, something they should be very good at. There is a loss of security because we are now depending on the routers to properly process the packets. We wouldn't be protected against routing attacks. However, we're already inside the main firewall, so we can take that risk. A routing attack could be thwarted by our firewall rules.

There are two unsolved problems we face when firewalling off a group of machines as large as a division. One is technical, the other is policy. The technical issue is that the network of a very large division usually has multiple connections to the rest of the company. In this situation we would need to put a firewall at each entry. If there are more than one connections to the division asymmetrical routing may be involved: a packet may enter via one connection and the reply might leave via another. Dealing with asymmetric routes through two different gateways requires the rapid sharing of state information. There are performance and security issues to getting this right. There aren't any products that do this today. This may change as demand for it increases.

The other problem is simply policy. Before one can define a set of firewall rules, one must determine the security policy. That is, you must define "what are you trying to protect?" and "how much risk are you willing to take?" So far, we have not been able to determine if, in our particular situation, we are protecting ourselves from the rest of the company or are we protecting the rest of the company from us. There are many arguments for both. For example, my division is very research focused. We try new protocols and take different risks than the rest of the company. Are we trying to protect the rest of the company from the risks we take?

On the other hand, there is an excellent argument that states we should be protecting ourselves from the rest of the company. An internal network census found 260,000 live hosts corporate-wide. When the Internet was 260,000 hosts large, we had a firewall to protect us from it. If our internal corporate network is that large, should we be considering a firewall to protect us from it? What's the difference between 260,000 Internet hosts run by people that we've never met, and 260,000 corporate hosts run by people that we've never met, but have the same source of paychecks? What's to say that one of our 150,000 employees hasn't accidentally (or on purpose) given access to outsiders? Who's to say one of our local users hasn't accidentally done the same?

In the future we hope to investigate these issues.

9 Trick: Connecting directly to the backbone

We put our firewall directly on our router backbone, making it zero additional hops from each customer group. Previously, all but one customer group was an extra hop from the firewall. This change was not done for performance reasons as the additional hop would go unnoticed in today's huge Internet. Instead, this was done to prevent the one router from being a single point of failure. We do not want an outage in one customer group to cause outages in other customer groups.

Astute readers will point out that the firewall itself is a single point of failure. We deal with that in other ways (such as a hot spare, etc.).

To achieve this goal we will perform similar steps to when we moved the firewall from Router A to Router B (see Figure 3).

First we connect the Router Backbone to the hub. We assign a secondary IP address to the external router so that the same interface is both on the hub's IP subnet and the Router Backbone's IP subnet. The same battery of tests is used as before.

The next step is to configure outgoing packets to go directly to the External Router. The default route of Router A, B, C and D is changed to point to the External Router's backbone address. Now all outgoing packets are going directly to the External Router after, of course, being filtered by the firewall.

Now let's take care of the incoming packets. The external router must be configured with a static route for every subnet to the appropriate customer groups' router. We can save some time by using a default route for the router with the most routes and providing static routes for all the other routers.

Finally we can remove the secondary IP addresses that enable the hub to be both on the Router Backbone and the old transit net simultaneously. Before we do this, we can do a final test to make sure everything was done right. We can disconnect the wire between Router B and the hub. If our battery of tests succeeds, then we know it is safe to remove the old secondary IP addresses.

It is possible to do this with a routed firewall, but downtime would be required as we made the routing changes. Many routed firewalls do not support secondary IP addresses, which means the changes could not be done live without a service interruption. This is an excellent example of one of the benefits of decoupling routing from the firewall. A full-featured, dedicated router is more likely to have the routing features we need than a routed firewall.

10 Conclusions

Bridged firewalls decouple routing from filtering. This provides interesting operational benefits.

As a bridge, it can placed between any two network devices. If those two devices are hubs, a subset of machines on the same IP subnet can be firewalled from each other, something a routed firewall could never do.

Its extremely easy to put a bridged firewall in series with other firewalls for testing or other purposes.

With a bridged firewall, many network topology changes can be done without service interruptions. These physical and logical changes are not on the firewall, but on the routers that exist in the network. Thus change is localized and therefore simplified. Fewer modifications on the firewall is also a plus because often security measures to "lock down" a firewall make changes to its configuration bothersome.

In our rapidly changing network, these benefits have been a remarkable advantage. Many of our examples involved a total elimination of service disruption as major changes were made. Some required disruptions on the order of seconds. We reduced or eliminated the number of transit networks from two to one, or from two to zero.

In the future we hope to explore firewalling off large labs of machines and experiment with what policies are appropriate when firewalling off larger and larger numbers of machines within a corporate intranet.

We also hope to explore the opposite direction: firewalling off extremely small groups of hosts. A bridged firewall might be just the right technology for a firewall in the home, especially given that IP addresses are not allocated in abundance in that market without higher charges.

Being able to place a firewall between any two network devices changes the way we think about firewalls. They are no longer routers with special filtering abilities. They are independent filters that can be put between any two devices, anywhere, at any time, even between large groups of devices (point to point, one to many, and many to many). They can even be put in series with each other. Decoupling routing from filtering lets us take advantage of full-featured, dedicated routers which may have features that a routed firewall may not provide. We have only begun to scratch the surface of the new possibilities introduced by this new paradigm.

11 Acknowledgments

This paper couldn't have been written without the help and support of Bill Cheswick, Lookman Fazal, Paul Glick, Darren Shaw, and everyone involved with the LMF prototype team: David Majette, Mike Coss, and Ron Sharp. Thanks to Josh Simon, Tommy Reingold, Adam Porter, and William LeFebvre for their editing and feedback.

References

[Chapman] "Building Internet Firewalls". D. Brent Chapman and Elizabeth D. Zwicky. O'Reilly and Associates. Cambridge, MA. 1995.

[Cheswick94] "Firewalls and Internet Security; Repelling the Wily Hacker". W. R. Cheswick and S. M. Bellovin. Addison Wesley. Reading, MA. 1994.

[Cheswick96] "A DNS Filter and Switch for Packet-filtering Gateways". W. R. Cheswick and S. M. Bellovin. Proceedings of the 6th UNIX Security Symposium. July 1996.

[LMF] The Lucent Managed Firewall. `http://www.lucent.com/security/`

[RFC1597] "RFC1597: Address Allocation for Private Internets". Y. Rekhter, B. Moskowitz, D. Karrenberg, G. de Groot. March 1994