
Transparency

The various network architectural issues involved when deploying load balancers is one of the toughest aspects of load balancing itself, especially for those with a server background. This document is meant to help bring some of the more common issues to light, and to help LoadMaster users navigate the options for deploying load balancing.

One of the most common issues that comes up, and one that is a bit difficult to grasp, is the issues surrounding L7 transparency. This document serves as an explanation of network transparency, its implications, and other related concepts.

Implications of Network Transparency

All of the issues with network transparency can be simplified with this single question:

Do you need the IP address of the client requests in your logs?

If the answer is yes, then you need network transparency and you'll need to configure your LoadMaster and architect your network in a certain way, which this document will describe.

If the answer is no, then you have a little more flexibility in how your network can be configured.

The table below shows a matrix of the advantages and disadvantages of transparency.

	Transparent	Non-Transparent
Pros	Preserves source IP address	Can browse from same subnet as real server
	Works with L4 and L7	No need to change the default gateway
Cons	Can't browse from the same subnet as Real Servers	Source IP address is not preserved (Though X-Forwarded-For header can be used)
	Default gateway must be LoadMaster	Only available for L7

The difference between transparency settings are based on making sure traffic moves from the real server back to the client through the LoadMaster. This type of symmetric routing, that is, going in and out of the LoadMaster, is an inherent

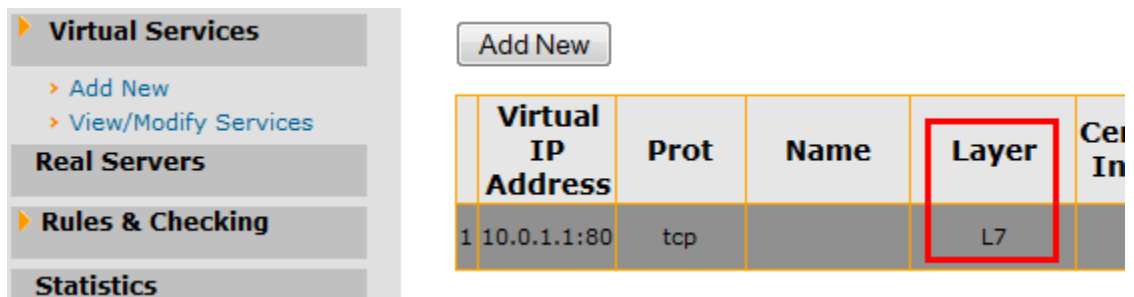
requirement of all load balancers (with the exception of employing direct server return, a feature which the LoadMaster supports, which has its own set of limitations).

Layer 4 and Layer 7

The LoadMaster makes a differentiation between L4 and L7 handling. They refer to Layer 4 and Layer 7 of the OSI model. Layer 4 involves TCP/UDP ports, and Layer 7 refers to the higher-level awareness of the LoadMaster, such as with HTTP cookies, SSL acceleration, and content switching. For all Layer 4 virtual services, the only behavior available is transparent networking.

What does Layer 4 mean? It's any load balanced traffic that does not involve cookie persistence, content switching or content switching rules, or SSL acceleration. Layer 4 does include SRC (source IP) address persistence.

The way to tell if a virtual service is using L4 or L7 handling is to look at the Virtual Service in the list when you click View/Modify. It will indicate what layer it is operating on in the Layer column.



The screenshot shows a navigation menu on the left with options: Virtual Services (expanded), Add New, View/Modify Services, Real Servers, Rules & Checking, and Statistics. An 'Add New' button is located above the table. The table has the following data:

	Virtual IP Address	Prot	Name	Layer	Cel In
1	10.0.1.1:80	tcp		L7	

Figure 1 Layer 7 Handling

Any time any cookie persistence, SSL acceleration, or content switching options are used, the traffic automatically becomes L7.

Transparency Requirements

In order to do transparency, the default gateway of your real servers must be the LoadMaster. This is true whether the network configuration is one-armed or two-armed. Without being the default gateway, there is no way to ensure that traffic passes through the LoadMaster on the way from the server to the client, and the LoadMaster can't do its job.

Another requirement of transparency is that you must be browsing from a subnet other than that of the real servers. Again, it has to do with making sure traffic

passes in and out of the LoadMaster. If you're on the same subnet as the real server, the traffic will simply go directly to the client, instead of through the LoadMaster. As a result, the client will be expecting to see traffic come from the IP of the virtual service, but instead will see traffic coming from the IP of the real server. When that happens, the client system ignores the traffic. For a more detailed explanation, refer to the appendix of this document.

Enable Layer 7 Transparency

Each L7 virtual service has the capability of being transparent or non-transparent. If the service is a L7 service, whether it's using some of the L7 handling features, or if it's forced, you will see the following checkbox.



Figure 2 L7 Transparency Option

This checkbox governs the transparency setting for this specific virtual service.

The LoadMaster can be configured to allow a global transparency setting. To enable this, navigate to System Configuration --> Miscellaneous Options --> L7 Configuration and set the L7 Transparency box to Transparent.

If this box is set to Non-transparent, you will be able to set transparency on a per-virtual service basis. This is the default setting.

Layer 7 Issues

When you're doing load balancing without any Layer 5-7 functionality, such as no cookie persistence and no SSL acceleration, then the only option is transparent. Even if transparency is disabled in the LoadMaster configuration, Layer 4 traffic is always transparent.

Transparency, SNAT, and Single-Arm Networks

If you've got a single-armed configuration, that is when the Virtual Services and the real servers are on the same subnet, and you're employing transparency, you'll want to disable SNAT (Source NAT). SNAT is the mechanism that allows servers behind the LoadMaster to make outbound connections in a two-armed configuration. It acts much like an office firewall, by "masquerading" the outbound connections as coming from a public IP address. In a single-armed configuration,

SNAT isn't necessary, although it normally doesn't interfere with regular operations.

The exception is when you're using transparency, the LoadMaster is the default gateway for your real servers, and you want to access the real servers directly. SNAT will “break” connections directly to the servers by attempting to masquerade those connections, so you'll want to disable SNAT.

To disable SNAT, go to the System Properties tab in the WUI. From there, select the “Miscellaneous” sub-menu, and you'll see a check box for SNAT. Simply uncheck the box, and SNAT will be disabled. You should be able to access your servers directly now.

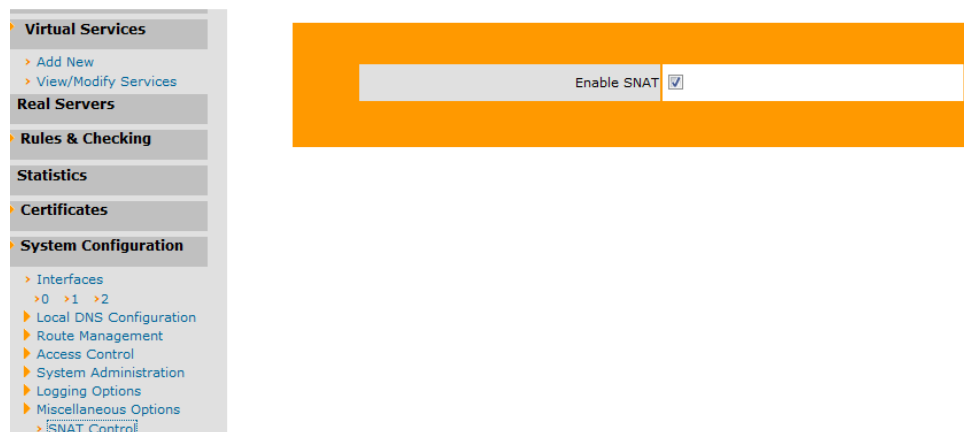


Figure 3 SNAT Control

Non-Transparency

There are two main benefits to using non-transparency. The first benefit is that it allows you to browse your virtual service when the client is on the same subnet. The other advantage is that you do not need to make the LoadMaster the default route in a one-armed configuration. Traffic is forced through the LoadMaster on the way out by making the request appear it came from the LoadMaster itself (which is why the IP address is hidden).

Non-transparency is the default setting for the LoadMaster, but it only affects Layer 7 traffic. To force Layer 7 handling even if you're not using cookies, content switching, or SSL acceleration, go into the properties for the virtual service in question, and select the option for “Force L7”.




Figure 4 Force L7 Handling

If you employ cookie persistence, content switching, or SSL acceleration for a given virtual service, this option goes away. And as discussed, the chief disadvantage is that the source IP address of the client is hidden, although it is forwarded in a separate HTTP header.

Additional L7 HTTP Header

While the source IP address is not preserved in the regular sense with non-transparency, the LoadMaster does provide a method to retrieve the actual source IP address through an HTTP header. For HTTP GET requests, the LoadMaster inserts an additional HTTP header, called X-Forwarded-For when L7 is used with non-transparency.

In order for these headers to be sent by the LoadMaster, the following conditions must be met:

- Virtual service is operating L7, non-transparent
- L7 is not achieved through 'Force L7' checkbox

What this means is the virtual service must be operating at L7 because it is using either some L7 persistence mode (i.e. NOT Source IP), content switching, or SSL acceleration for it to send these headers.

Depending on your web server or application infrastructure, you can configure the X-Forwarded-For value to be logged.

For Apache, the combined format in the httpd configuration file is as follows:

```
LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"  
\"%{forensic-id}n\" combined
```

Add the client side field by adding `%{X-Forwarded-For}`.

```
LogFormat "%{X-Forwarded-For}i %l %u %t \"%r\" %>s %b \"%{Referer}i\"  
\"%{User-Agent}i\" combined
```

There is another option available under the Additional L7 Header drop down box. This is the X-ClientSide Header. This is just an alternative to the X-Forwarded-For Header.

To log these in Apache, use the following code:

```
LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"  
\"%{forensic-id}n\" \"%{X-ClientSide}\" combined-ClientSide
```

Why can't I browse from the same subnet with Transparency?

With a network configuration with transparency enabled, the reason why you can't browse from the local network is because of the path that the traffic flows. As stated, in order for a load balancer to do its job, the load balancer must be in the path of both inbound and outbound traffic. Load balancing happens typically in four steps:

1. Client to the virtual service on LoadMaster
2. LoadMaster to real server
3. Real server to LoadMaster
4. LoadMaster to client

Take the example of a simple one-armed configuration, where the client IP address is 64.254.1.12, the virtual service address is 192.168.0.200, and the real server is 192.168.0.100. In the table below, you'll see what happens in a regular connection.

Step	Path	Source IP	Destination IP
1	Client to Virtual Service	64.254.1.12	192.168.1.200
2	Virtual Service to Real Server	64.264.1.12	192.168.1.100
3	Real Server to Client (before LoadMaster)	192.168.1.100	64.254.1.12
4	Virtual Service to Client (after LoadMaster)	192.168.1.200	64.254.1.12

Now take the same example, except this time the client will have the IP address of 192.168.0.10, which is on the same subnet as the real server.

Step	Path	Source IP	Destination IP
1	Client to Virtual Service	192.168.0.10	192.168.1.200
2	Virtual Service to Real Server	192.168.0.10	192.168.1.100
3	Real Server to Client (before LoadMaster)	192.168.1.100	192.168.0.10
4	Virtual Service to Client (after LoadMaster)	192.168.1.200	192.168.0.10

The last step doesn't happen, because the real server doesn't need to send traffic out its default gateway since the client is on the same subnet, so it sends its response directly to the client. The response comes back from a different IP address than the client was expecting, so the client drops the traffic entirely, and the page never loads.

So why can I browse from the same subnet with Non-Transparency?

Non-transparency replaces the IP address of the client with the IP address of the LoadMaster itself, thereby forcing traffic back through the LoadMaster on the way out. When the real server responds to the request, it responds to the LoadMaster. The LoadMaster then forwards the traffic along to the client.

Step	Path	Source IP	Destination IP
1	Client to Virtual Service	192.168.0.10	192.168.1.200
2	Virtual Service to Real Server	192.168.1.200	192.168.1.100
3	Real Server to Client (before LoadMaster)	192.168.1.100	192.168.1.200
4	Virtual Service to Client (after LoadMaster)	192.168.1.200	192.168.0.10

Notice that in the first table for transparency, either the source IP or the destination IP was rewritten, but not both. In the second table for non-transparency, both the source IP and destination IP were re-written. This is why the logs of your web server will only see IP address of the LoadMaster for all incoming connections when transparency is disabled (as it is by default).