

Technical White Paper

The IDNS module for incoming load balancing

For Elfiq Operating System (EOS) version 3.x

Document Revision 1.6

March 2008

Build **Smarter** Networks



Table of Contents

1. The iDNS module.....	3
1.1 DNS sub-domain requests (internal DNS scenario).....	4
1.2 DNS sub-domain requests (external DNS scenario).....	4

1. The IDNS module

The key element for the Elfiq Link Load Balancer (Link LB) that will allow it to balance outside traffic for incoming load balancing is to properly manage domain name resolution, or DNS, requests.

One of the most fundamental services of internet and IP networks is the DNS service. When accessing an IP service, users and applications will almost always try to access it with a fully qualified domain name, and not through its IP address. Standard DNS servers usually statically map resource records to a single IP address. Although some improvements have been made to allow DNS servers to link a resource record to multiple IP addresses, this will still limit incoming requests to a round robin balancing. The major problem is that DNS servers won't verify the availability of the link before sending DNS resolutions in a round robin balancing. In case of a link failure, this could result in a false unavailable service access attempt.

The Link LB, on the other hand, can dynamically answer DNS queries submitted by the client applications according to the balancing algorithms configured for each of the resources, which provides a very high level of flexibility and a much higher level of balancing.

The Elfiq Link Load Balancer has an Intelligent DNS module (called IDNS) for incoming load balancing. It monitors all traffic going through the Elfiq Link Load Balancer and watches for incoming DNS requests. When it sees one, the Link LB verifies if it is the right type and if it has a matching entry in its configuration and if it does, it intercepts the request and answers on behalf of the DNS server. The IDNS will intercept the DNS requests only for IP services you need to balance and that are configured with incoming rules. If there are requests for services that are not configured in the Link LB, they will simply pass through the Link LB, transparently, and reach your DNS server.

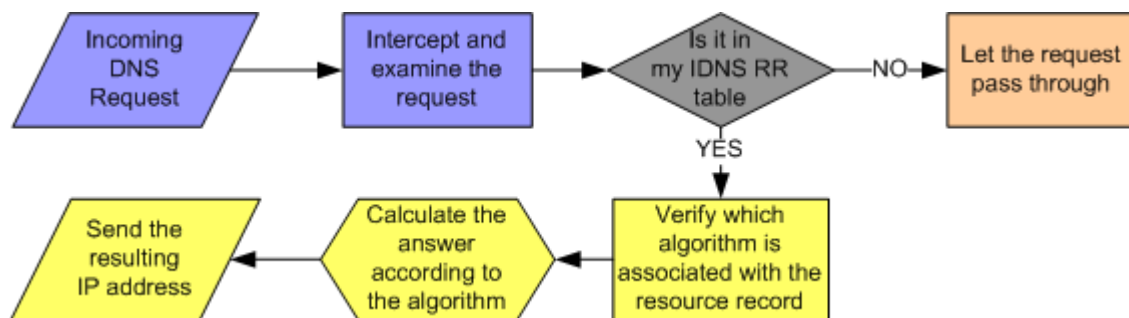


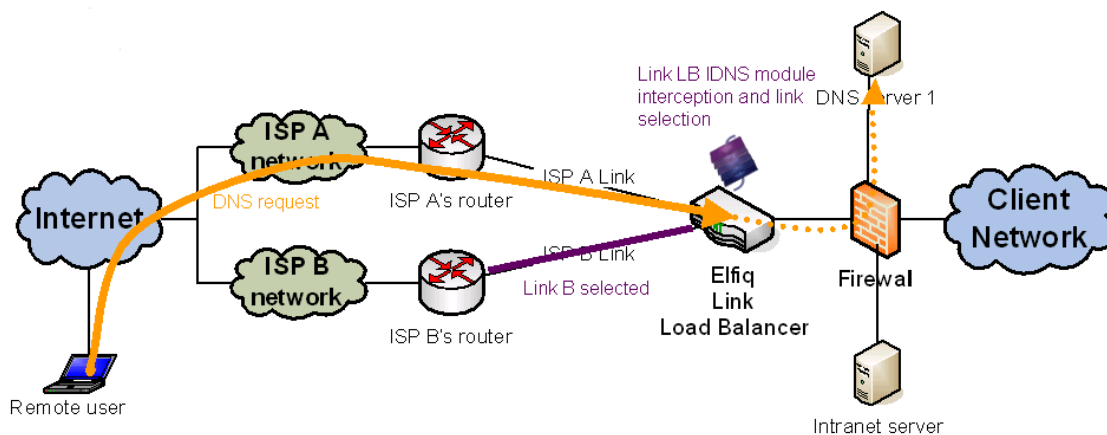
Figure 1.1: Flowchart of the IDNS facility DNS interception process

Public DNS servers can be hosted inside your organization or via your Internet service provider (ISP). Both scenarios are covered in this document.

1.1 DNS sub-domain requests (internal DNS scenario)

If your public DNS servers for each domain that will have balanced resources are hosted internally, behind the Elfiq Link Load Balancer, no modifications to your DNS servers are needed. They will keep “A records” pointing to your resource IP addresses on the primary link network.

Any query about a balanced resource within your domain will necessarily have to pass through the Link LB to reach your DNS servers. While all DNS queries pass through the Link LB, it will verify if it has an iDNS resource record associated with that DNS query. If it does, it will intercept the query and answer with the appropriate IP address according to the balancing algorithm. This means every



DNS query for which you create iDNS resource records will never reach your real DNS server.

Figure 1.2: The DNS request interception process and link selection process

1.2 DNS sub-domain requests (external DNS scenario)

If your public DNS servers are hosted by an external site, then some changes will need to be performed to the current DNS entries of each of the resources that you want to be balanced by the Elfiq Link Load Balancer.

In a typical configuration, each service that you want to provide to the Internet requires an “A record” DNS entry to resolve its server name to an IP address. This works well when your network is not load balanced and only accessible through one link. In this scenario we want to perform DNS queries and resolve a server name to its IP address on the best available link even in the event of a link failure. Such a resolution is impossible if all the DNS queries are to be answered by the external DNS, since that server has no information on the alternate IPs to give, let alone which link is the best to use. Moreover, because the DNS server is not behind the Link LB, there is no way that the Link LB can intercept those DNS queries before they reach the DNS server, like it is the case when all the DNS servers are hosted internally.

In order to resolve this issue, we need to force a redirection of the DNS queries for all resources or services that we want to balance to the Link LB, which will then be able to resolve the hostnames to the desired balanced IP addresses. This is performed by modifying the “A records” on the external DNS servers for the services that you wish to balance into “NS records”. Those “NS records” will be pointing to the Link LB. This enables your external DNS server to delegate the resolution of those specific records to the Link LB, which will now be able to properly resolve the DNS queries according to the balancing methods configured on your Link LB.

However, since the Link LB operates at the data link layer (layer 2 of the OSI model) it doesn’t have any IP address defined on its interfaces; in this manner it resembles an Ethernet switch or a hub. This means that you cannot send IP packets directly to the Link LB, instead, queries have to be redirected to an IP address in the network segment behind your Link LB. This way, the queries will have to pass through the Link LB unit, and at that point they can be intercepted. No real DNS server has to exist at this address. This IP address will be used in the creation of the iDNS interceptors in your Link LB, so that it knows that it should only intercept DNS queries destined for this virtual DNS server.

Let’s take a look at the following example, from a BIND DNS configuration, which illustrates how the zone file for your DNS could be configured, should it be hosted on a public DNS server:

```

$TTL 1h
example.com.      IN SOA      ns1.isp.com. hostmaster.isp.com. (
                    2005041301 ; serial yymmddxx
                    1800      ; Refresh 30 min.
                    3600      ; Retry 1hours
                    360000    ; Expire 100hours
                    3600 )    ; Minimum TTL 1 hour
;
                    IN NS      ns1.isp.com.
                    IN NS      ns2.isp.com.
www               IN A       33.33.33.9
    
```

No mention of the Elfiq Link Load Balancer is present in this configuration. There is only an “A” record of hosts that currently can be resolved: www.example.com

Now, let’s modify the DNS server to start to balance www.example.com. To do so, as mentioned before, we would have to change the “A” type record of the www entry, to an “NS” type record, pointing to a virtual DNS server behind our Link LB. Here is how the new DNS entry for www would look like:

```

www               IN NS      virtualdns1.example.com.
                  IN NS      virtualdns2.example.com.
    
```

The hostname that you give to this virtual DNS server can be anything you wish, however, the DNS server must know at which IP it can reach it. Therefore, you would need to add another entry, pointing to the IP addresses of the virtual DNS (one IP per link), like the following:

```

virtualdns1      IN A       33.33.33.5 ;Primary Link IP
virtualdns2      IN A       44.44.44.35 ;Secondary link IP
    
```

The same entry needs to be added once per link that you are balancing, pointing to the virtual DNS server of each of those links. This will enable your DNS server to direct the DNS queries of all resources that point to virtualdnsX.example.com to be reached on any of the available virtual DNS server IP addresses, ensuring that queries will always be properly resolved by the Link LB on each of your balanced links. **This method ensures, in case of a link failure, that the sub-domain request will always reach the Elfiq Link LB, as long as one link is available, because the external DNS server will try every virtualdnsX IP address to get a DNS resolution for the www.example.com domain resource.**

This configuration for incoming load balancing ensures to always return an accessible IP address on the best available link for DNS resolution. It solves the issues discussed in the introduction with a basic DNS resource record to multiple IP addresses in round robin.

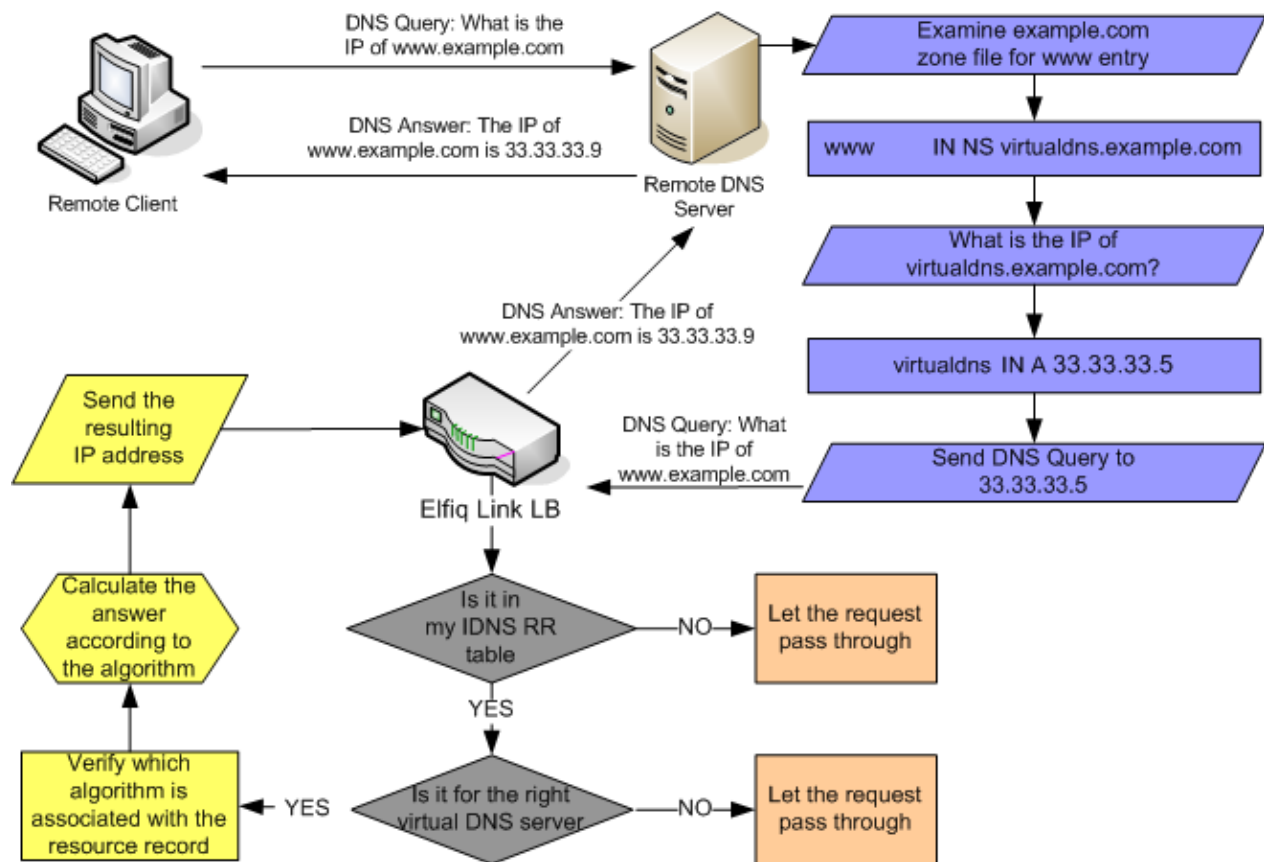


Figure 1.3: Flowchart of the DNS sub-domain request

Produced by Elfiq Networks

Elfiq Networks is a technology leader and innovator in the field of WAN link management and balancing. With hundreds of successful installations in over 49 countries, Elfiq's Link Balancer products help organizations of any type and size perform more competitively every day with the ability to use multiple Internet and private links easily and securely.

For more information on Elfiq Networks' products and technologies, please contact:

Elfiq Networks
1155 University, #712
Montreal, Quebec, H3B 3A7
Canada
Telephone: 888-GO-ELFIQ / 514-667-0611
Internet: www.elfiq.com
Email: info@elfiq.com

May 2009

© Copyright 2009, Elfiq Networks (Elfiq Inc.). The contents of this document are protected by copyright. Any modification of this document, in any shape or form, is prohibited. Any redistribution, publication or derivation of the contents of this document without written authorization from Elfiq is also prohibited. All rights reserved.