



Refreshing ftp.heanet.ie

Rob Gallagher

Managed Network Services, HEAnet

robert.gallagher@heanet.ie




Overview

- Refreshing ftp.heanet.ie - why?
- Basic requirements
- Architecture
- Software load balancing methods
- LVS
- New hardware
- Things to come..

Why ?

- Existing hardware was EoL
- Single point of failure
 - Difficult to carry out maintenance without downtime
 - Publically visible service
 - “is heanet down?”
- Bottleneck
 - Release of Ubuntu Hardy caused an extremely high load on the system.



Basic Requirements

- Remove the bottleneck
- Remove the single point of failure
- Scalability



Architecture

- Some questions...
 - Stick with our existing arch. and just upgrade our single, powerful machine?
 - Some advantages, but obviously doesn't fit with our new requirements.
 - Single point of failure.
 - Difficult/impossible to carry out maintenance without (very visible) downtime.
 - Popular software gets released, machine gets DoS'ed into oblivion.



Architecture

- Some questions...contd..
 - **Load balancing + failover ?**
 - Multiple off-the-shelf machines.
 - Easy swap out of faulty components.
 - Easy to carry out maintenance.
 - Popular software releases won't take everything down.
We have the bandwidth :)
 - Move to high-performance shared storage system in May 2008 made this easier.



Architecture

- Some questions...contd..
 - **Hardware load balancers ?**
 - Advantages:
 - *Support*
 - *Ease of administration (debatable, ever tried to use a CSS?)*
 - *Speed (again debatable, see below)*
 - Disadvantages:
 - *€€€€€€€€'s!, custom hardware with ASIC components.*
 - *Speed, throughput etc.. is often geared towards accelerating e-commerce sites and such. Doesn't really fit with our situation.*
 - *Locked into a single manufacturer*



Architecture

- Some questions...contd..
 - **Software load balancing ?**
 - Quick setup
 - Easy to test
 - Easy to administer
 - Easy to monitor

Software Load Balancing

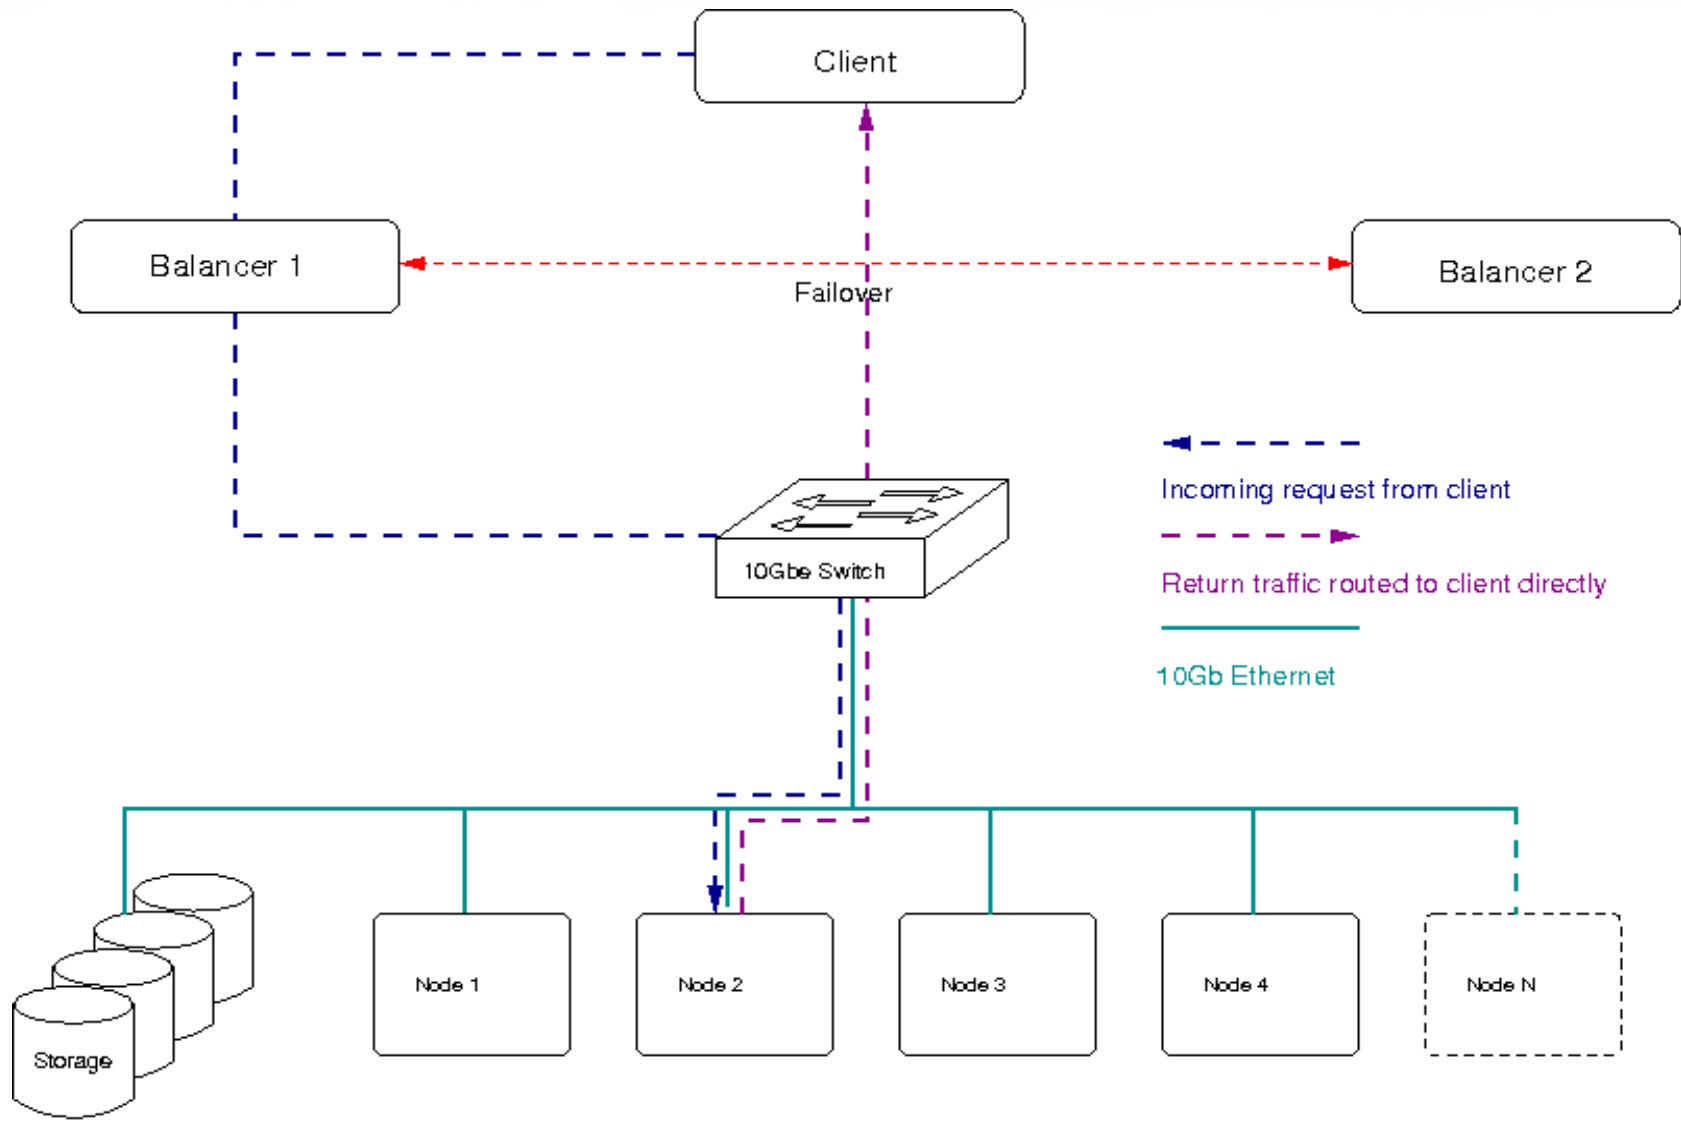
- Various options
 - Round robin DNS
 - Advantages:
 - *Very straightforward to deploy.*
 - *Supports IPv6 ;-)*
 - *No special hardware or software required.*
 - Disadvantages:
 - *Not true load balancing, far from fine grained.*
 - *There would always be a bias to the machine with ftp.heanet.ie's original IP addresses.*
 - *Broken resolvers that don't honour TTLs.*

Software Load Balancing

- Various options...contd.
 - **OpenBSD relayd**
 - Advantages:
 - *Extremely versatile and powerful.*
 - *Full IPv6 support.*
 - *Subject to OpenBSD's stringent security requirements.*
 - Disadvantages:
 - *We've standardised on Linux.*
 - *The operating mode we are interested in, Direct Server Return, is very new (May 2008) and largely undocumented.*

Software Load Balancing

- Various options...contd.
 - LVS (Linux Virtual Server)
 - Advantages:
 - *Does direct routing*
 - » *Server talks back to the client directly (more later).*
 - *Very straightforward to setup.*
 - *In combination with keepalived, can get a full load balanced system with LB failover and node monitoring up and running within minutes.*
 - Disadvantages:
 - *No IPv6 support.*



LVS

- LVS: Load balancing support within the Linux kernel.

- Quick to get running

- On the load balancer:

```
- ipvsadm -A -t 193.1.219.54:80 -s wlc
```

```
- ipvsadm -a -t 193.1.219.54:80 -r 193.1.219.91 -g
```

- On the server:

```
- sysctl net.ipv4.ip_forward=1
```

```
- ifconfig lo:0 193.1.219.54 netmask 255.255.255.255 broadcast  
193.1.219.54 up
```

```
- route add -host 193.1.219.54 dev lo:0
```

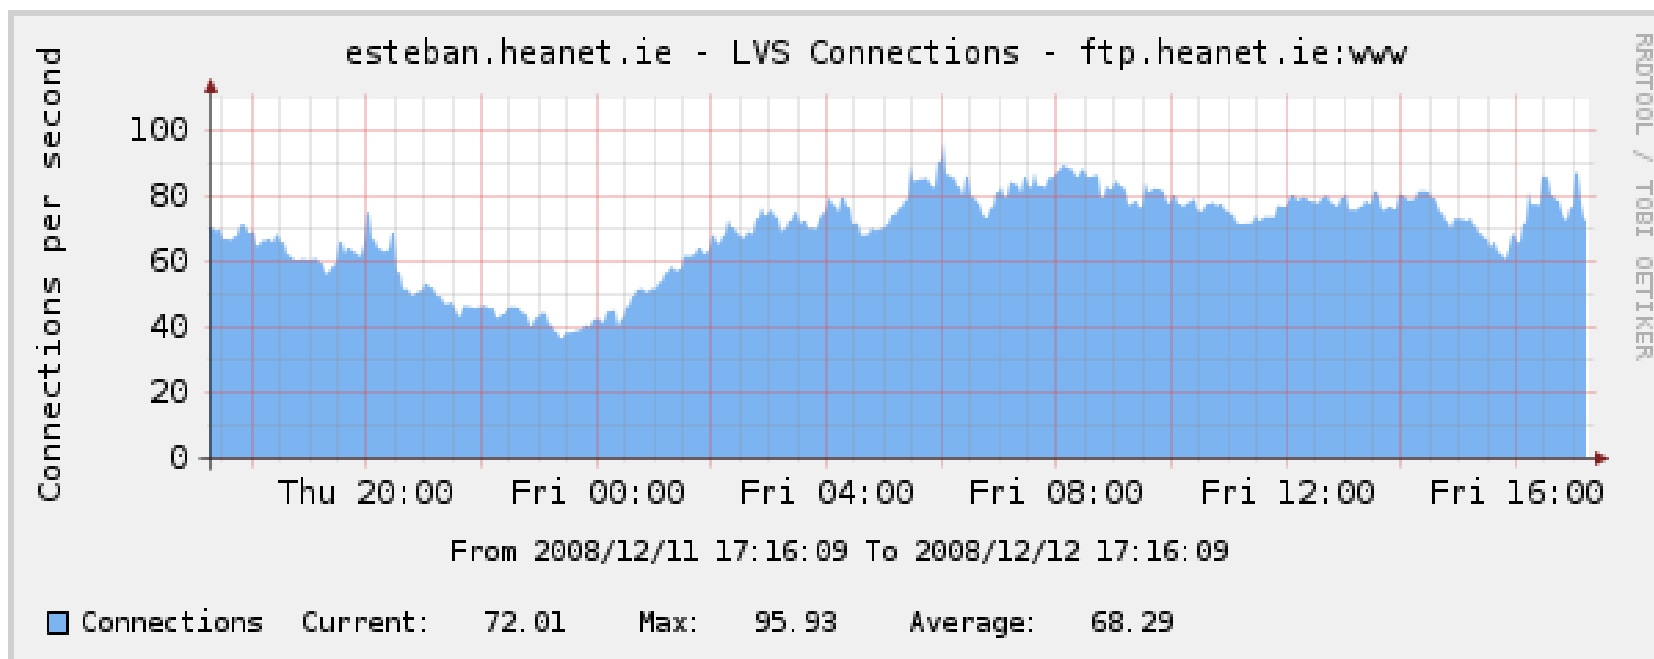
LVS

- Direct Routing

- Essential for a busy site like ftp.heanet.ie
- Load balancer sends incoming request to one of the real servers as normal.
- Real server responds to the client directly, without involving the load balancer.
- Traffic profile of ftp.heanet.ie: Comparitively small requests generate massive amounts of return traffic.
- ~~Bottlenecks.~~ :)

LVS

- LB failover and real server monitoring.
 - **Keepalived**
 - VRRP gives us floating IP addresses on the load balancers.
 - Keepalived monitors real servers, evicts malfunctioning hosts.
 - However, no IPv6 support currently.



Hardware



- Acquired new high-performance storage in April/May 2008.
- Isilon IQ NAS.
- This facilitated the new, clustered system.

Hardware

- Real servers:
 - 4 x PowerEdge R905s:
 - 4 x Quad-Core AMD Opteron
 - 64GB of RAM
 - Dual 10Gig Ethernet
 - /proc/cpuinfo scrolls off your terminal
- Load balancers:
 - 2 x PowerEdge R300s
 - Not as interesting, spec wise, LVS needs very little grunt, even for a site as large as ftp.heanet.ie

Things to come...

- 10 Gigabit Ethernet
 - Each machine has 2 x 10 Gig NICs
 - HEAnet core is capable of 10 Gig, INEX etc..
 - Arista 24-port 10 Gig switch.
 - One node already operating with a 20 Gig port channel to the Arista switch.

Things to come...

- IPv6 load balancing
 - As of kernel 2.6.28 IPVS has (mostly complete) IPv6 support.
 - Node and load balancer failover:
 - Migrate to ldirectord and heartbeart, both of which support IPv6.



?