

Clustering PHP and MySQL

php|tropics

May 12th, 2005. Cancun, Mexico

Derick Rethans <dr@ez.no>

<http://derickrethans.nl/talks.php>

Scalability

- ! Processing Speed
- A scalable application allows you to keep adding servers without adding (much) overhead
- In other words: your application's performance does not suffer from more load (or data size)

Reasons To Cluster

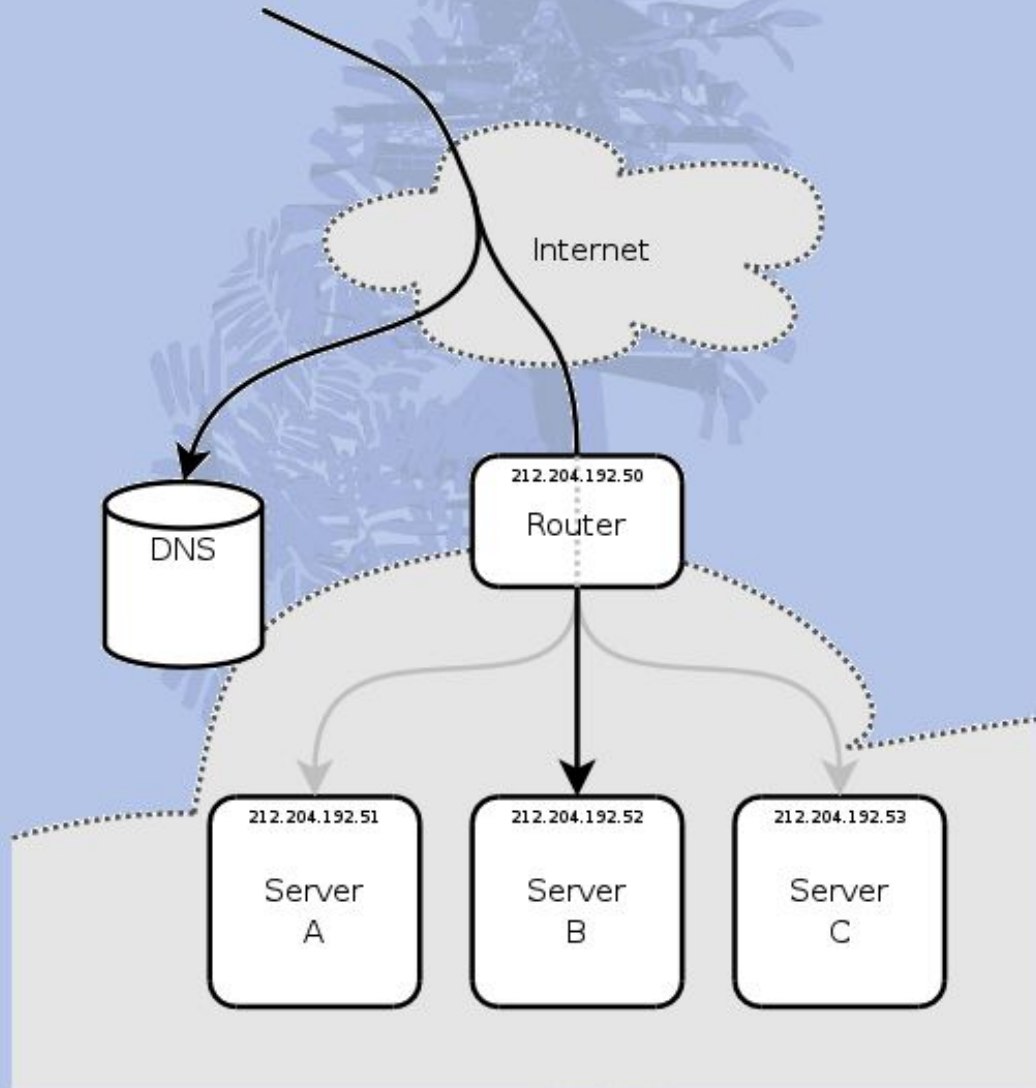
Load Balancing:

- A single server can not handle the load
- Faster machine does not always work
- Different bottlenecks: Disk IO, CPU or Network
- SMP has inherent problems

High Availability:

- Failover Capabilities
- High Availability (the 9s)

Round Robin DNS



Round Robin DNS

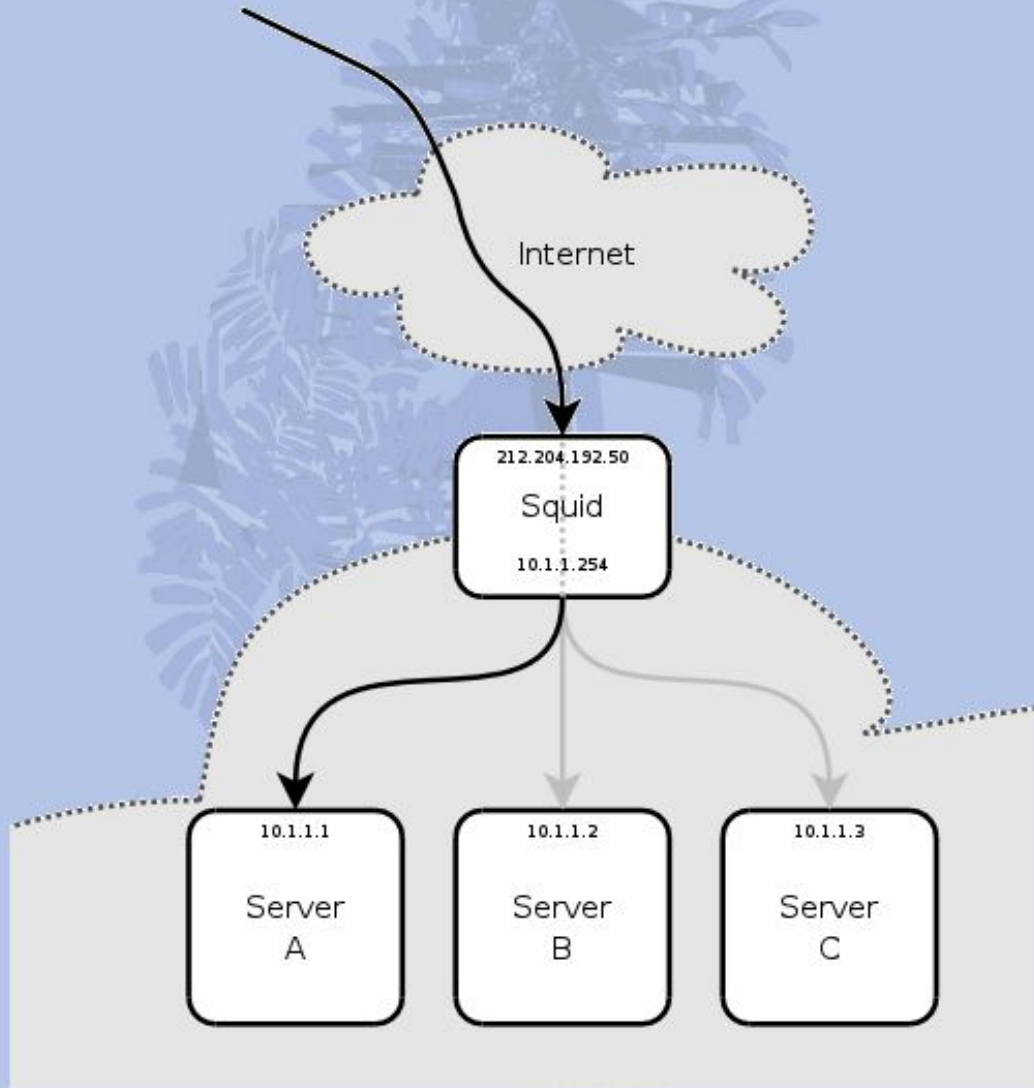
- Uses DNS to give out different IPs in a circular fashion
- Easy to set-up
- Limited control over which server gets requests

DNS config:

```
www.example.com. 60 IN A 10.45.67.1  
www.example.com. 60 IN A 10.45.67.2  
www.example.com. 60 IN A 10.45.67.3  
www.example.com. 60 IN A 10.45.67.4
```

- A short TLL is important to avoid caching

Squid as Front Proxy



Squid as Front Proxy

- Squid is a HTTP Caching Proxy
- It can also be used as front end proxy
- With a redirector script, you can implement a load balancer

Squid setup:

```
http_port 80
httpd_accel_host 127.0.0.1
httpd_accel_port 8080

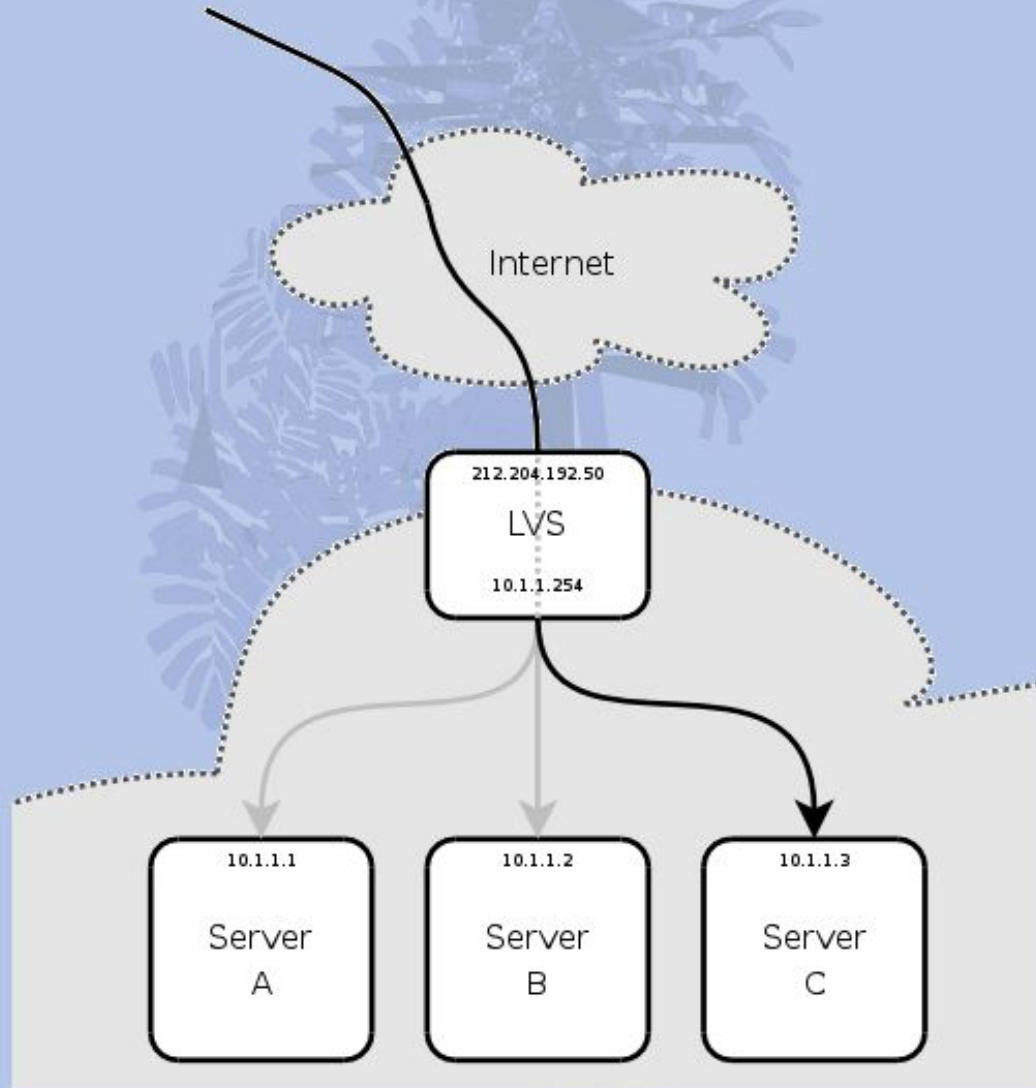
httpd_accel_single_host on
httpd_accel_uses_host_header on

redirect_program /usr/local/bin/redirect.sh
redirect_children 20
```

Redirector script responsible for:

- Rewriting URLs
- Distributing the load over the backends

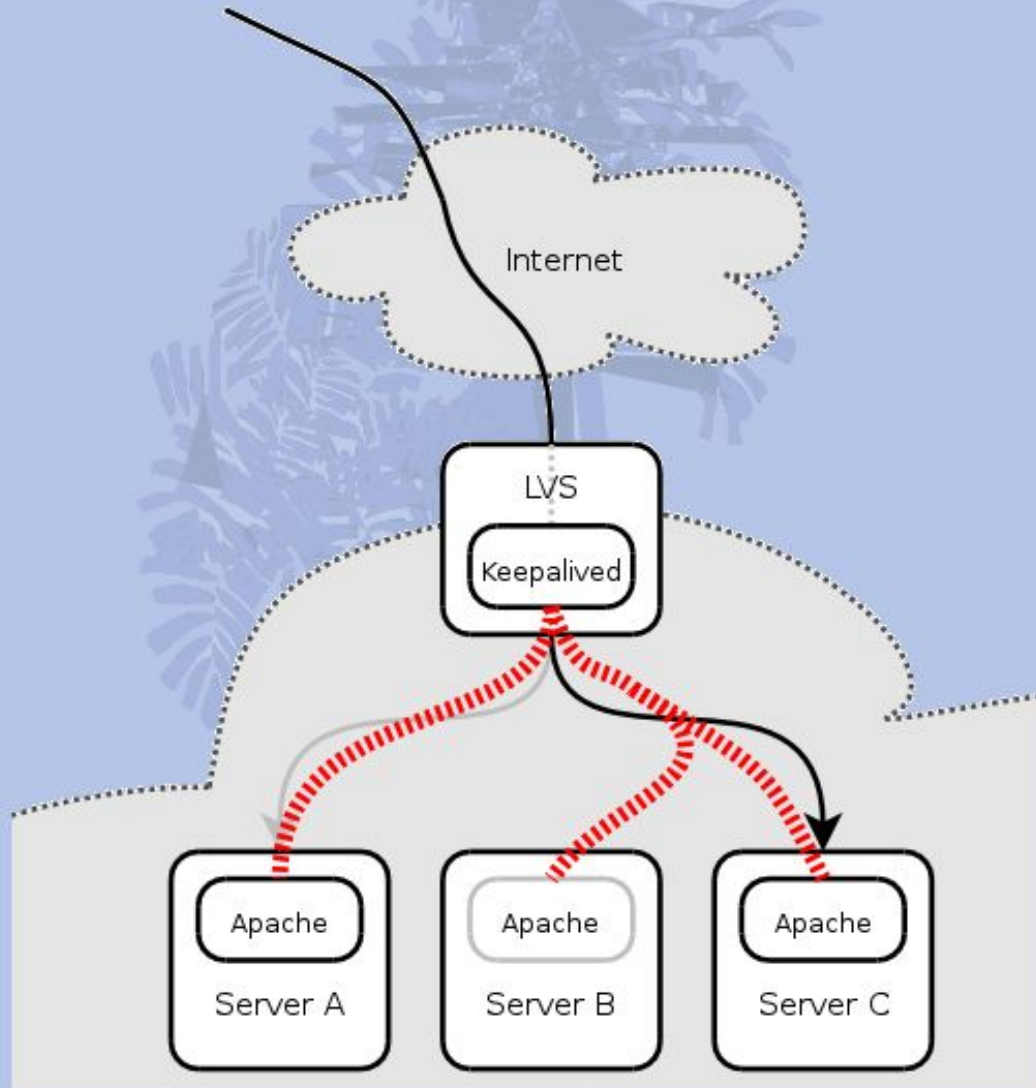
Linux Virtual Server



Linux Virtual Server

- Front End Server: Redirector (Load Balancer)
- Uses the load of specific worker servers to determine how to forward
- Uses a kernel module for implements transport-layer load balancing inside the Linux kernel, so called Layer-4 switching
- Can use different scheduling algorithms
- Does **not** provide fail-safe functionality

Keepalived

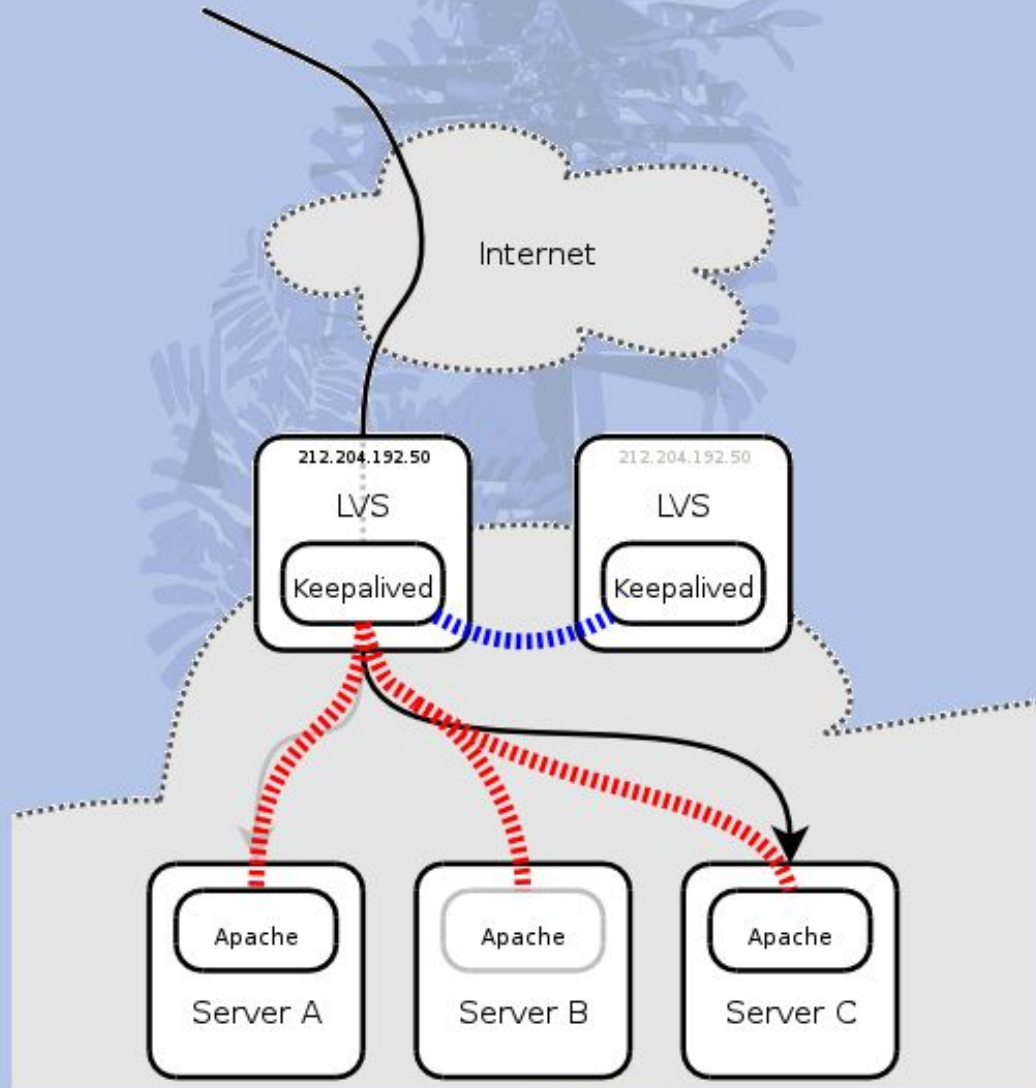


Keepalived

- Hooks in into LVS
- Provides fail-over functionality
- Checks at different level: TCP, HTTP Get and Script



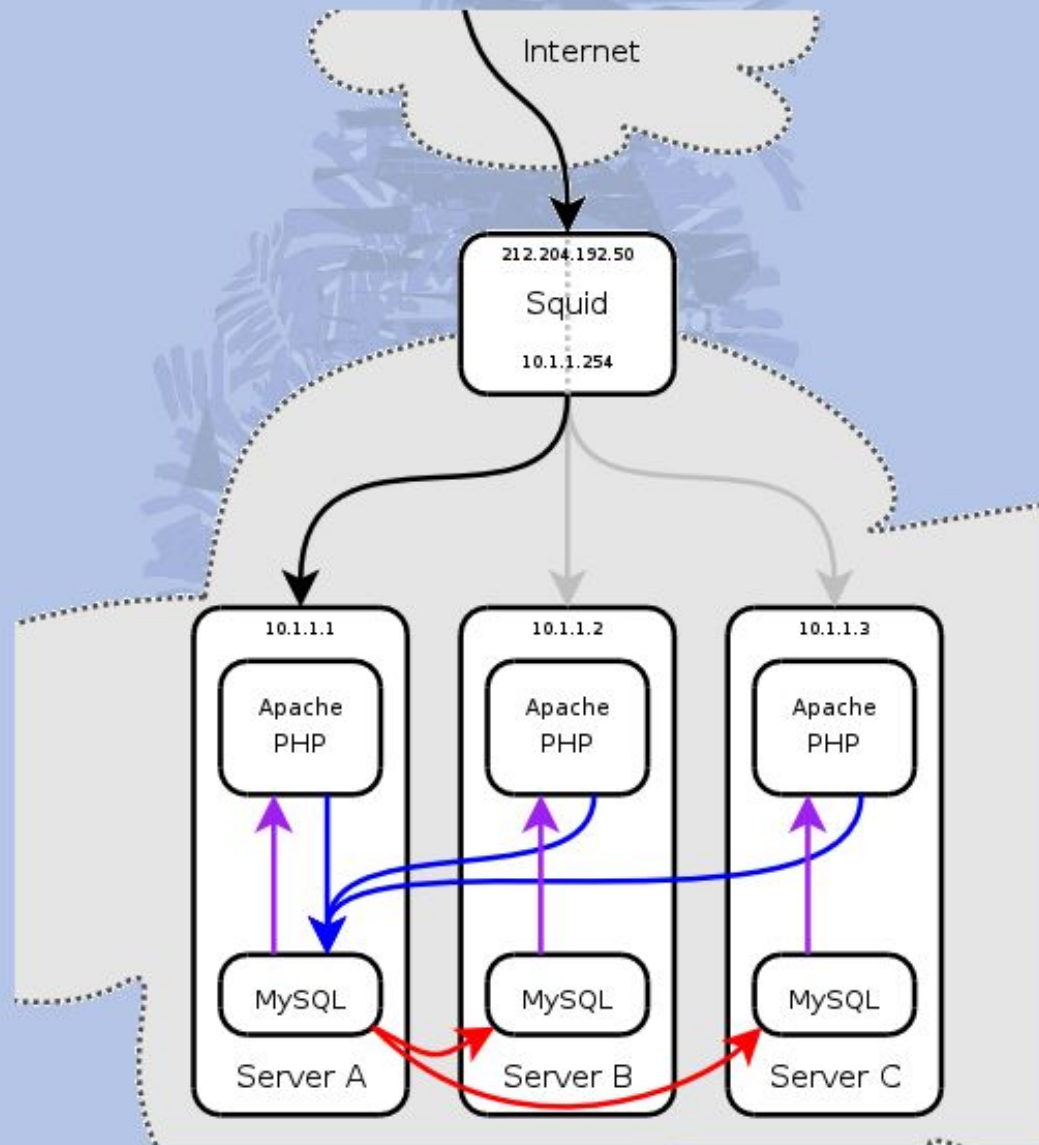
Back-Up redirector



Back-Up redirector

- Keepalived also implements failover with VRPP
- This allows you to have a Back-Up redirector to take over when the primary goes down

Database Replication



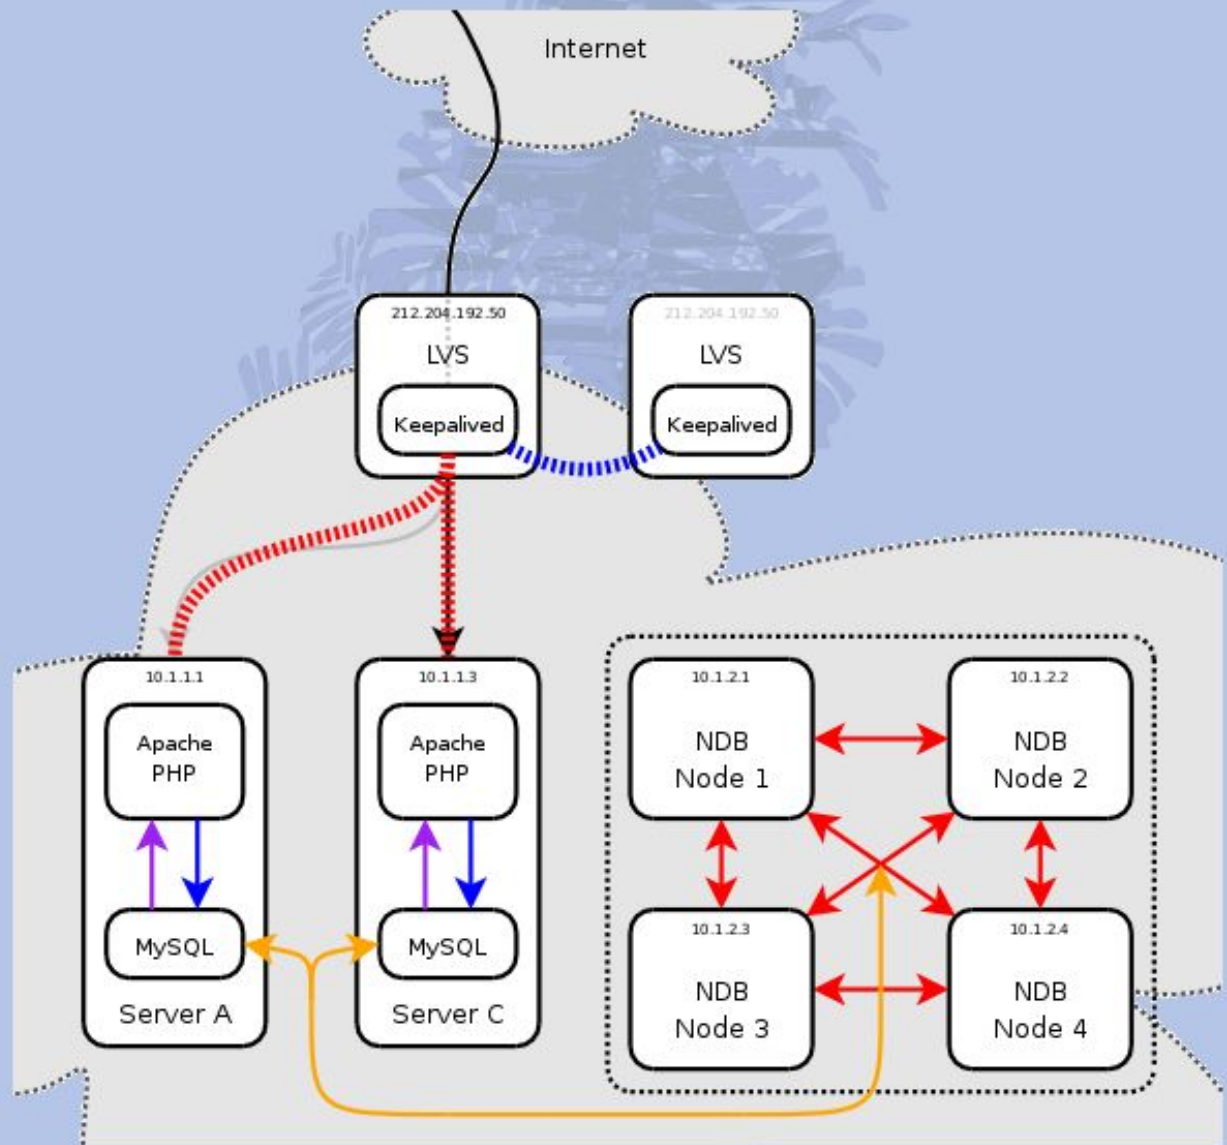
Database Replication

- Master with multiple slaves
- All writes go to the master
- All reads go to one of the slaves

Problems:

- No real fail-safe
- Transactions are not atomic over all nodes

MySQL Cluster



MySQL Cluster

- NDB Support
- Multiple replica's for fail-safe
- Transactions are atomic cluster wise


Different nodes:

- Management Node
- SQL Nodes
- Storage Nodes

Problems:

- Everything is in-memory
- It likes using a lot of memory

PHP 5 Power Programming



- Written by Andi Gutmans, Stig Sæther Bakken and Derick Rethans
- Meant for people familiar with PHP (or another language) wanting to know how to use PHP in the most efficient way
- ISBN: 0-13-147149-X

Resources

These Slides: <http://derickrethans.nl/talks.php>

PHP: <http://www.php.net>

Squid redirecting: <http://www.squid-cache.org/Doc/FAQ/FAQ-15.html>

Linux Virtual Server: <http://linuxvirtualserver.org>

Keepalived: <http://keepalived.org/>

MySQL Cluster: <http://dev.mysql.com/doc/mysql/en/ndbcluster.html>

Questions?: <mailto:dr@ez.no>