

Speed Daemons

Lars Noldan, Systems Administrator

BSDCAN 2011

six feet up



where sophisticated web projects **thrive**

Who am I?



- Lars Noldan
- Six Feet Up
- Systems Administrator

- e-mail: lars@sixfeetup.com
- IRC: LarsN on Freenode
- Phone: 317-861-5948 x609
- Twitter: @Absenth

What Requirements?

Every Project Starts With Requirements

thrive

Technical Requirements

- Site Must Maintain 99.9% Uptime.
- Pages Load in 5-10 Seconds on Broadband Connections.
- Process User Registration in 15-20 Seconds on Broadband Connections.
- Must Support Between 5,000 and 50,000 Registered Users.
- Must Integrate With Existing Systems



Software Requirements

- Zope / Plone (Enterprise CMS)
 - With RelStorage (MySQL Backed)
- Solr (Powered By Apache Lucene)
 - Apache Tomcat Engine

Open Source Solutions!

Using FOSS To Solve The Puzzle

thrive

Web Proxy Chain

- Nginx - Web Server
 - Light Weight
 - Extremely Fast
- Varnish - Reverse Proxy Cache
 - High Performance
 - Extremely Flexible
- HAProxy - Load Balancer
 - Session Aware
 - Configurable



Application Servers

- Python 2.4
- Zope 2.10 - Python Application Server
- Plone 3.3.5 - Enterprise Content Management
- Apache Tomcat - Java Application Server
- Solr 1.4 - Enterprise Search Software



Database Servers

- MySQL 5.1 - Database Server

Network Technologies

Can't Build a Cluster Without a Network

drive

Network Technologies

- FreeBSD 8.1-Release AMD64
 - PF - Enterprise Grade, FOSS Firewall
 - CARP - Common Address Redundancy Protocol
 - LAGG - Link Aggregation via LACP
 - ifstated - Network Interface State Engine





Speed Daemons!

High Performance Clusters



to drive

Simplified Cluster

Proxy01

Simplified Cluster

Proxy01

Application01

Simplified Cluster

six feet up
BSDCan, 2011



Proxy01

Application01

Database01

to
rive

Simplified Cluster



Proxy01

Nginx - 80/443
Varnish - 3180
HAProxy - 3380

Application01

Database01

to drive

Simplified Cluster



Proxy01

Nginx - 80/443
Varnish - 3180
HAProxy - 3380

Application01

Zope - 8080-8088
Solr - 6080

Database01

to
rive

Simplified Cluster



Proxy01

Nginx - 80/443
Varnish - 3180
HAProxy - 3380

Application01

Zope - 8080-8088
Solr - 6080

Database01

MySQL - 3306

to
rive

Somebody Order a Double?

Let's Kick This Cluster Into High Gear

drive

Expanded Cluster

six feet up
BSDCan, 2011



Proxy01

Proxy02

Application01

Database01

to
drive

Expanded Cluster

six feet up
BSDCan, 2011



Proxy01

Proxy02

Application01

Application02

Database01

to
rive

Expanded Cluster

six feet up
BSDCan, 2011



Proxy01

Proxy02

Application01

Application02

Database01

Database02

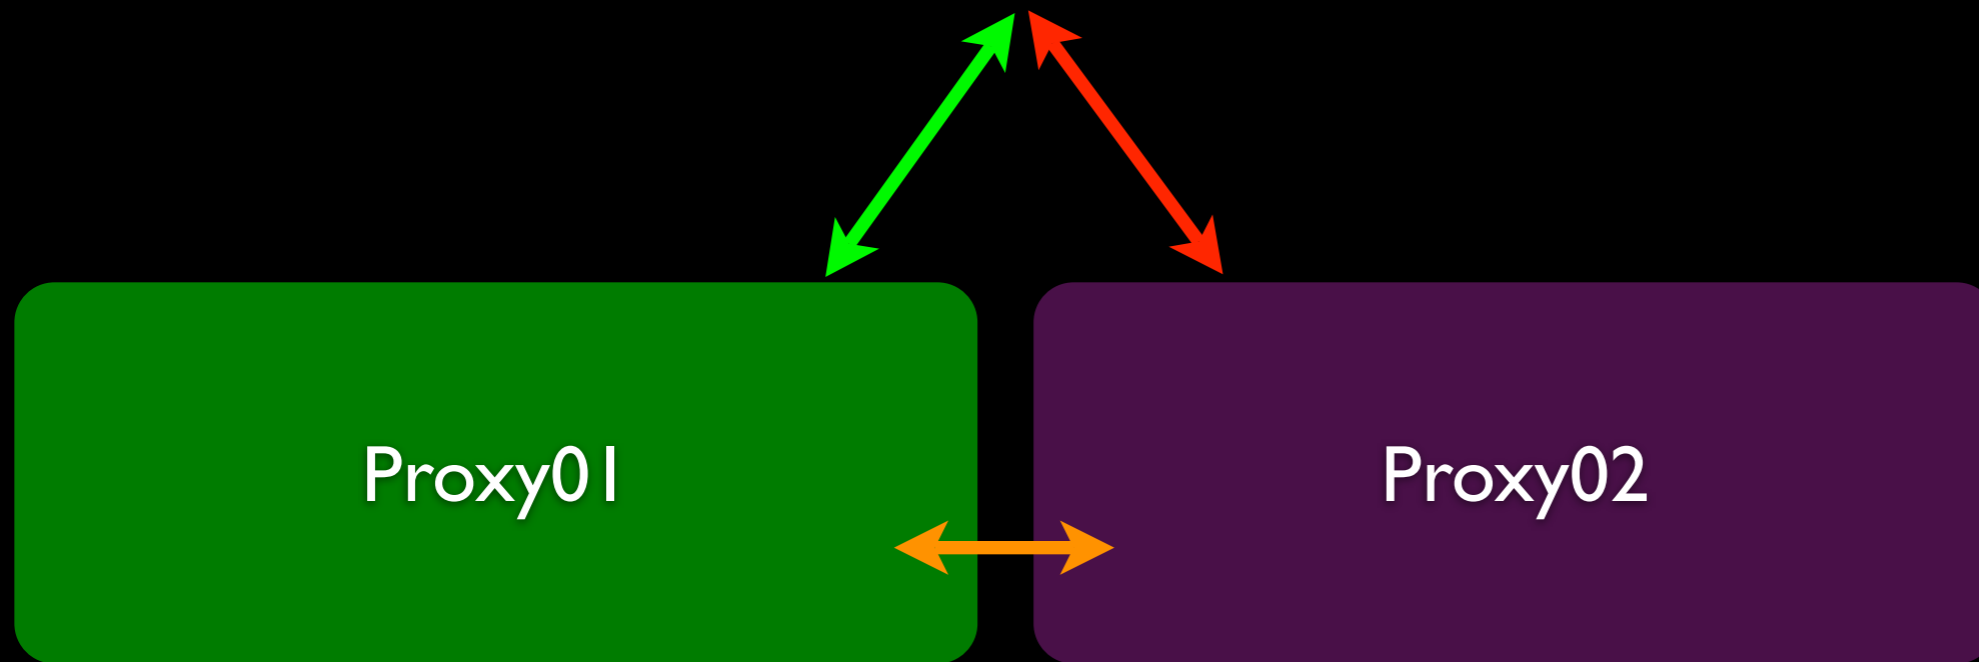
to
drive
e

Wire It Up


A Little Cable Never Hurt Anyone. Right?

thrive

Proxy Servers




Proxy Servers




- Each Proxy Server Has a Unique Public IP Address.
- Both Proxy Servers Share a Public IP Address via CARP

Proxy Servers




- ifstated - Runs On Both Servers
- Both Servers Constantly Monitor Each Other

Proxy Servers




- In The Event Of Any Service Interruption if stated Forces CARP To Fail Over Making Proxy02 Primary
- This Works Bi-Directionally

Proxy Servers




- pfsync Runs On Both Proxy Servers.
- This Keeps The Statefull Firewall In Sync Between The Two Servers

Application Servers

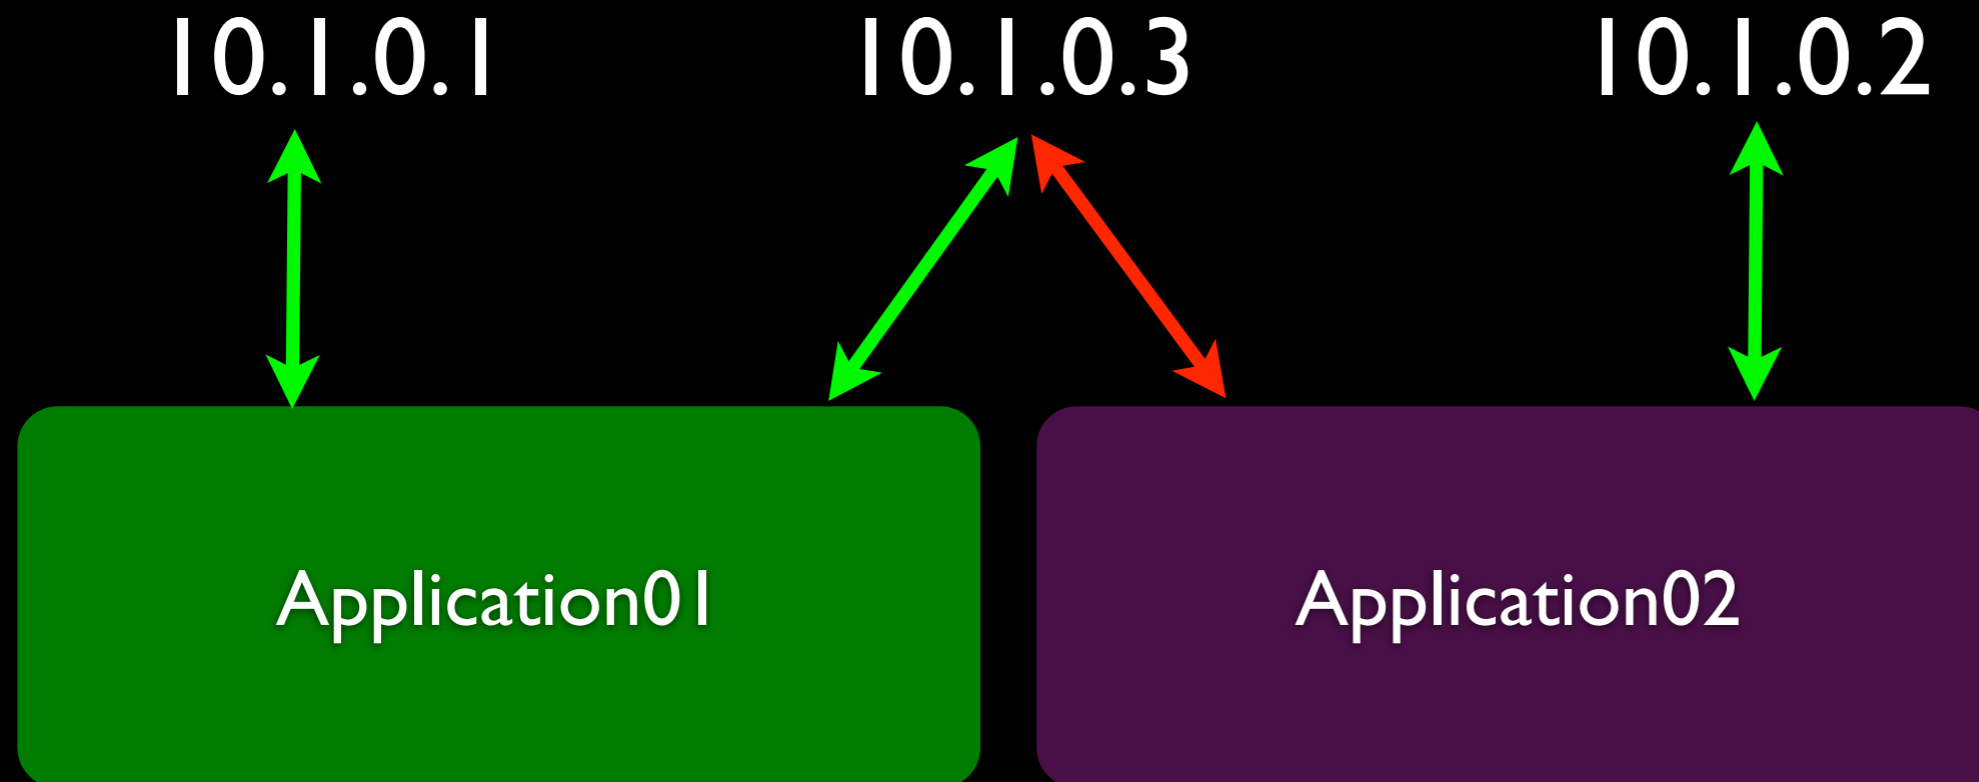


Application Servers



- These Servers Only Have Private Addresses
- Zope Listens On All Addresses
- Solr is Configured To Listen On The Shared CARP Address


Application Servers



- These Servers Share Solr Search Data
- ifstated - Does Not Monitor These Servers


Database Servers

six feet up
BSDCan, 2011



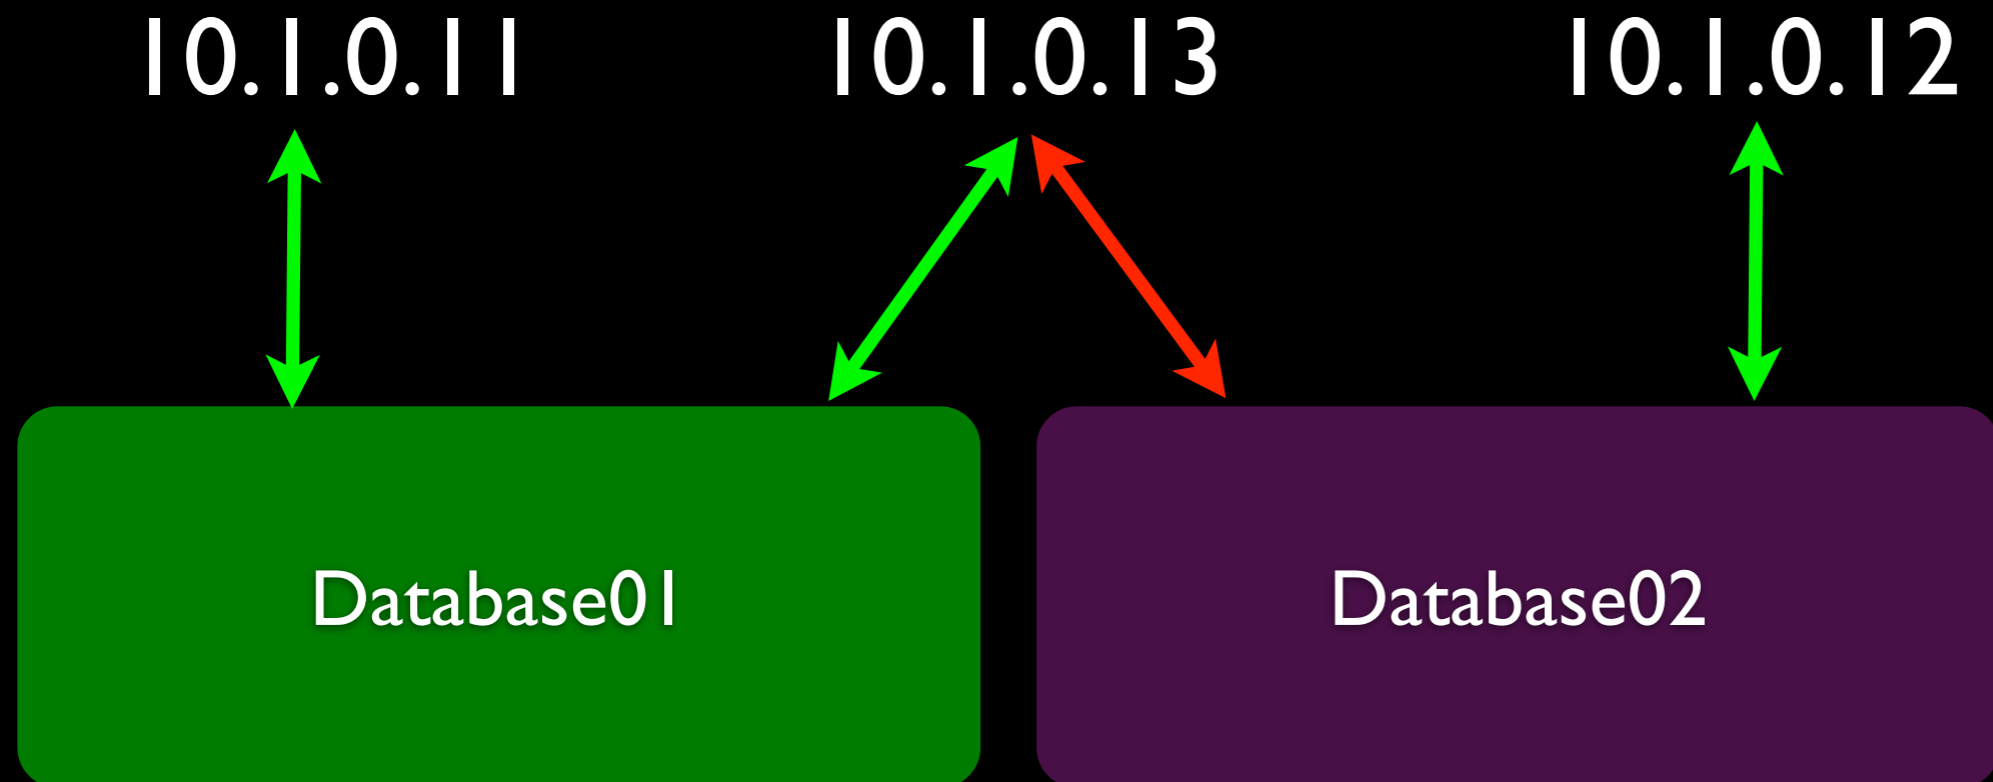
to drive

Database Servers



- These Servers Only Have Private Addresses
- MySQL Listens On All Addresses
- PF Is Used To Limit Access To MySQL

Database Servers



- ifstated - Runs Custom Scripts To Facilitate MySQL Master/Slave Replication State Changes
- Data Is Kept In Sync By MySQL Replication

Data Storage

Our Solutions For Storing and Shipping Data

thrive

Storage Solutions



Proxy01

Proxy02

Application01

Application02


Database01

Database02

Stor01

Stor02

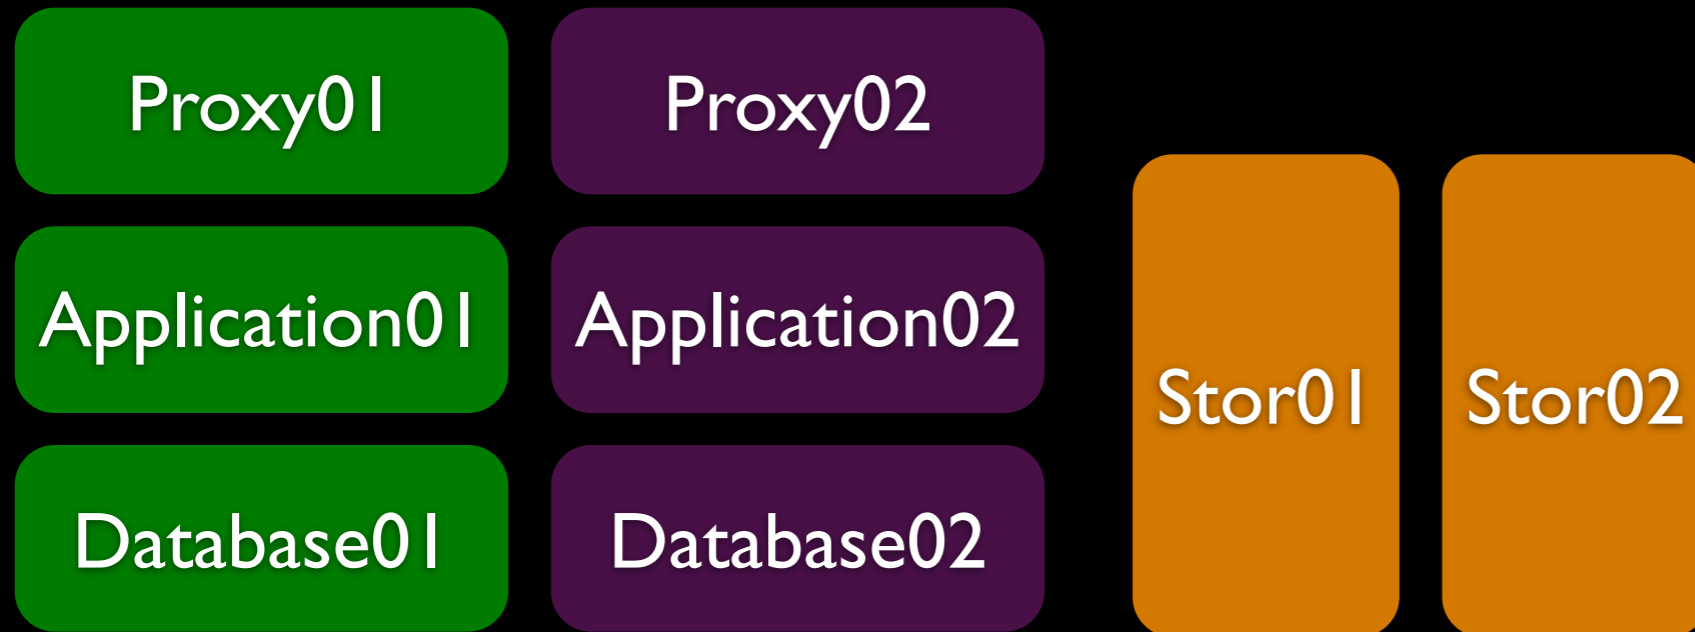
Storage Solutions



- Hardware:


- LSI MegaRaid - Raid Card Using JBOD
- 2 OS Drives - Mirror
- 2 Intel SSD Drives - Mirrored ZIL
- 7 1TB Spindles - RaidZ2 (5TB)

Storage Solutions




- Configuration v1.0:
 - FreeBSD 8.0-Release
 - Several File Backed iSCSI Mounts
 - Numerous NFS Shares
 - Snapshots At 5m, 15m, 1h
 - Stor01 Primary
 - Stor02 Replication Target

Storage Solutions




- Problems Encountered:
 - iSCSI Performance Abysmal
 - Kernel Race Condition During Snapshot Replication

Storage Solutions




- Configuration v2.0:
 - OpenSolaris - 2008.11
 - Several File Backed iSCSI Mounts
 - Numerous NFS Shares
 - Snapshots At 5m, 15m, 1h
 - Stor01 Primary
 - Stor02 Replication Target

Storage Solutions



- Future Configuration v3.0:
 - FreeBSD 9.0-Release
 - Several File Backed iSCSI Mounts
 - Numerous NFS Shares
 - Snapshots At 5m, 15m, 1h
 - Stor01 Primary
 - Stor02 Replication Target

Storage Solutions



- Future Configuration v3.0:
 - iet - iSCSI Target Replacement
 - zpool - v28
 - Kernel Race Conditions Fixed in 8.2

Does It Work?

A Quick Review

thrive

Quick Review

- Customer Requirements Have Been Met
- Over 1.5 Years Of Hosting 99.995% Uptime
- High Performance / High Availability Cluster
Can Be Built Using All OpenSource Solutions
- FreeBSD Rocks!

Who is Six Feet Up?

six feet up
BSDCan, 2011



- Web Development & Hosting Provider
- Focus on CMS, KMS, and Web Apps
- Located in Fortville, IN
- Major Player in Open Source Technologies such as:
 - Plone (CMS) - <http://www.plone.org>
 - KARL (KMS) - <http://www.karlproject.org>



Where Can You Find Us?

- [HTTP://www.sixfeetup.com](http://www.sixfeetup.com)
- Twitter: @sixfeetup
- Conferences and Events
 - IndyPy - Indianapolis Meetup Group
 - PyCon - Atlanta, GA
 - BSDCan - Ottawa, Canada
 - KMWorld - Washington, DC
 - Plone Symposium East - State College, PA
 - Plone Conference 2011, San Francisco, CA
 - OSCON - Portland, OR
 - Ohio Linux Fest - Columbus, OH

Want To Work With Me?

<http://www.sixfeetup.com/about-us/jobs>

or

E-Mail a Resume To: jobs@sixfeetup.com

thrive

where sophisticated web projects **thrive**

Check out
sixfeetup.com/demos