

^{ltd}
Esgob



Anycast on a shoe string

Nat Morris

10th May 2014 @ DNS-OARC

About me

- Consultant, **Esgob Ltd**
- Engineer, **Cumulus**
- Board member, **UKNOF**
- Based in West Wales
- Aspiring lighthouse keeper



Summary

- Fancied deploying a DNS Anycast service
- Motivated by:
 - Bill Woodcock
 - SANOG8: “Best Practices in DNS Anycast Service-Provision”
 - Dave Knight
 - RIPE64: “Dense Anycast Deployment of DNS Authority Servers”
- Gain more experience automating distributed environments.

Thoughts

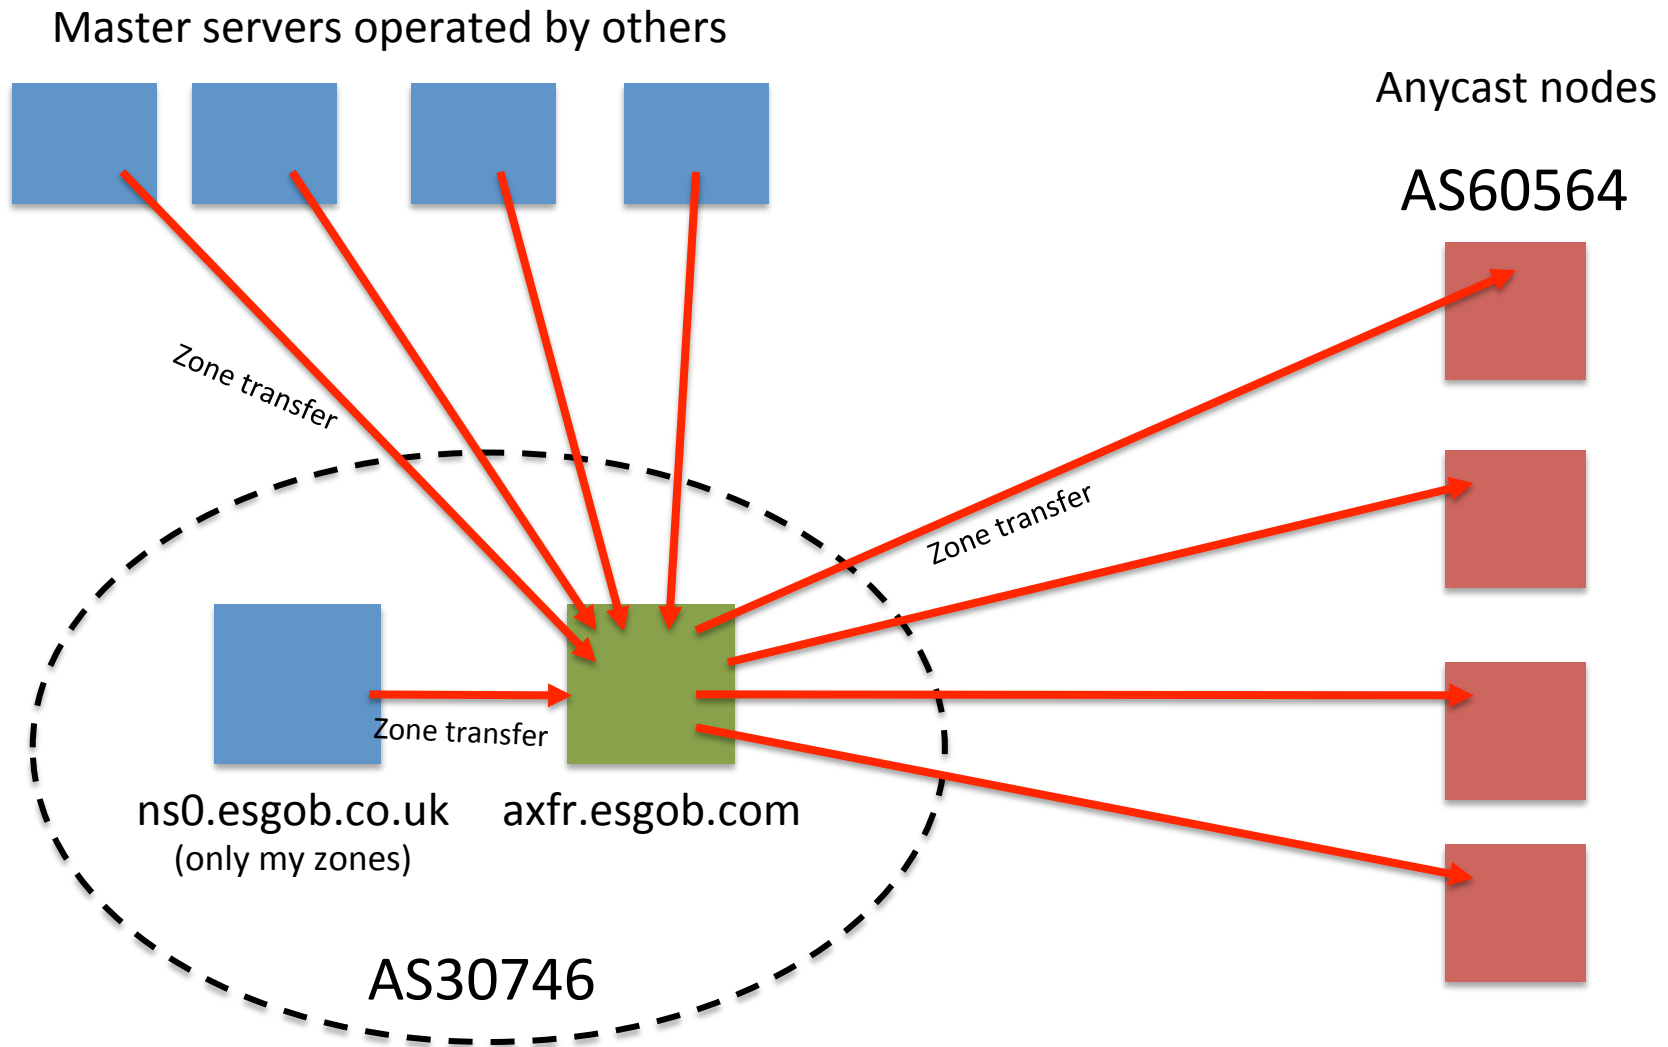
- Can I do it without spending too much?
- Does Claire need to find out?
 - Must be sub \$1000/yr running cost
- What to offer?
 - Initially only secondary DNS
 - IPv4 + IPv6 (yes Martin!)
 - Free service - no SLA

Requirements

- Separation from existing management network - AS30746
- Had spare PI /24 + /48
- Applied for new ASN - got AS60564
- Need highly automated framework
- Wanted to play with new tools
- Open source everything
 - <https://github.com/esgob>



DNS zone transfers



Supports notifies from masters

Store everything in JSON

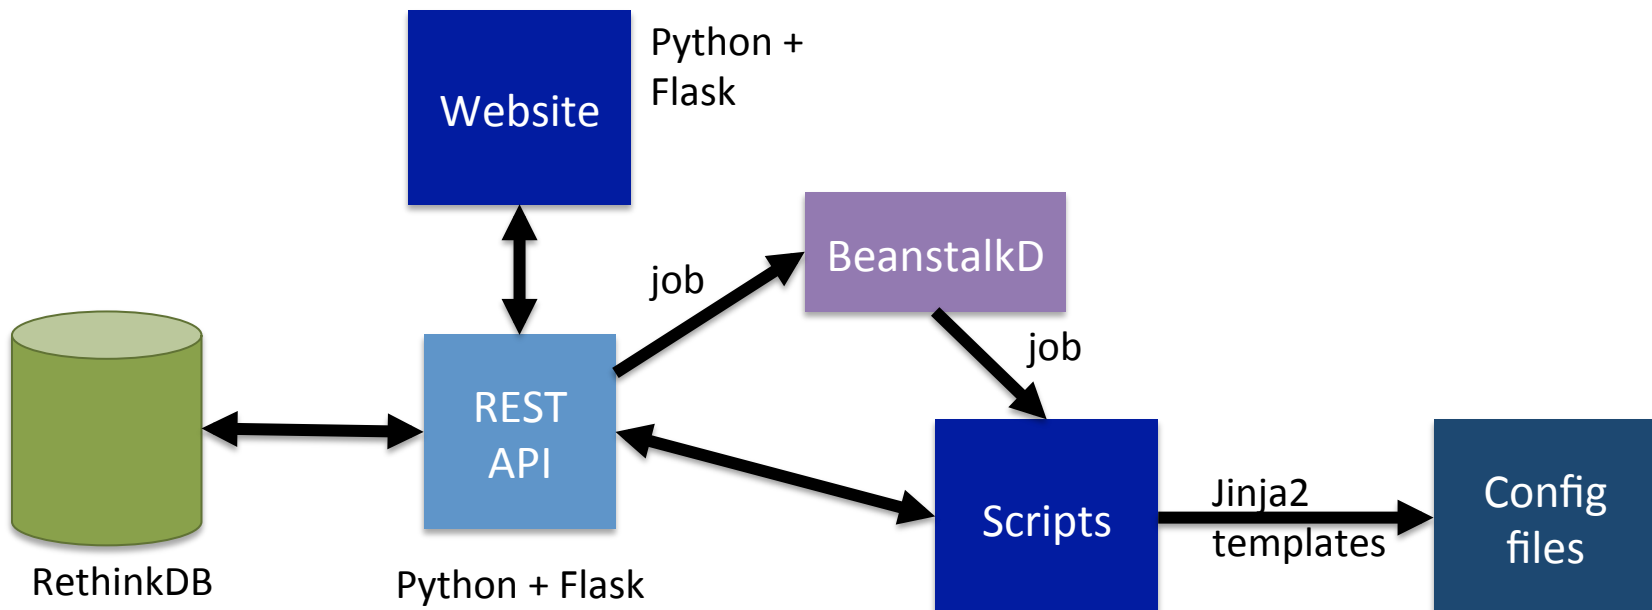
- Data stored in **RethinkDB**
 - Distributed JSON document database
 - Free form, easy to add attributes

```
{
  "account": "esgob",
  "added": 1391807801,
  "id": "esgob.co.uk",
  "masterip": "195.177.253.166",
  "type": "slave"
},
{
  "account": "mike",
  "added": 1392340310,
  "id": "grepular.com",
  "masterip": "2001:470:1f15:d7::2",
  "type": "slave"
},
```

```
{
  "city": "Edinburgh",
  "country": "Scotland",
  "countryiso": "gb",
  "flag": "scotland",
  "fqdn": "alban.esgob.com",
  "geolat": 55.929959,
  "geolng": -3.294011,
  "hostdesc": "VPS at Fluency",
  "id": "alban",
  "livedisplay": "25th Oct 2013",
  "locationdisplay": "Scotland, Edinburgh",
  "mgmtip4": "46.226.2.182",
  "mgmtip6": "2a00:a600:0:6::2",
  "peers4": [
    {
      "asn": 56595,
      "communities": [
        "56595:9500",
        "56595:9510",
        "56595:9551",
        "60564:10003",
        "60564:20044",
        "60564:30001"
      ],
      "name": "fluency",
      "neighbor": [
        {
          "ip": "46.226.2.181"
        }
      ]
    }
  ],
  "peers6": [
```

Application design

- REST API - Python + Flask framework
- Beanstalk - Message queue
 - Add/change/removal of zones via API triggers job to be added to queues for each node
 - Daemons on AXFR + Anycast nodes monitor for jobs
 - New zones live on all nodes with 5 seconds



Network

- Each anycast node:
 - Static default route to the provider
 - Announce /24 + /48 via BGP
- All management traffic inside OpenVPN
 - Additional RFC1918 loopback per node
 - No need for TSIG from AXFR to anycast nodes
 - Protects beanstalk traffic, collectd, zone transfers

Where to host the anycast nodes?

- Can't cost too much each month
 - Ruled out renting co-lo / dedicated servers
- Virtual machines
 - Swap with friends?
- Low cost hardware

Mythic Beasts

- UK clueful hosting company
- £7/month = £84/yr = \$141/yr
 - 512mb RAM, 1 CPU, 10gb HDD, 75gb BW
- Bought a VM, opened a support ticket
 - BGP not listed on the website.
 - Me: “I have a /24 + /48 of PI, can you set me up a BGP session?”
 - Them: “Peer with these IPs, sessions are ready!”
- Too easy!



Quest to find more friendly VM hosts

- LowEndBox.com
 - Blog / adverts for VM providers, all < \$10 month
 - Usually OpenVZ based
 - Need KVM/XEN/VMWare to support Quagga/BIRD
- LowEndTalk.com
 - Message board, various small scale VM hosts
- Google
 - “vps bgp session”

Anynode in Detroit



anyNode <support@anynode.net>

06/10/2013 ☆



to Nat ▾



Dear Nat Morris (ESGOB LTD),

We have received your order and will be processing it shortly. The details of the order are below:

Order Number: **6758645372**

Product/Service: Detroit KVM - DK-512

Domain: esgob

First Payment Amount: \$80.00 USD

Recurring Amount: \$0.00 USD

Billing Cycle: Annually

Total Due Today: \$80.00 USD

to Nat ▾

- LeapSwitch
- New
- \$120.00



Dear Nat Morris (Esgob Ltd),

This is a notice that an invoice has been generated on 29/09/2013.

Your payment method is: PayPal

Invoice #44717

Amount Due: \$48.00

Due Date: 29/09/2013

Invoice ItemsKVM-UVPS-1 - cyri.esgob.com (29/09/2013 - 28/09/2014)

Control Panel: NONE(Default)

FTP Backup: No Backup

IP Address: 0 x IP Addresses (In addition to default 1 IP) \$20.00

cPanel Addon: None

Monitoring: No Monitoring

~~Operating System: Debian 7.0.0 64bit \$120.00~~

Promotional Code: INKVMVPS512MB48 - \$72.00 Recurring Discount \$-72.00

Sub Total: \$48.00

Credit: \$0.00

Total: \$48.00

You can login to your client area to view and pay the invoice at <https://service.leapswitch.com>








Growing

- Up to 4 VMs
 - London A (on existing KVM server)
 - London B
 - Detroit
 - India
- Spinning up VM taking about 10 minutes
 - Add JSON entry into RethinkDB
 - Install Debian
 - Install Puppet (Collectd, BIND, Quagga, custom daemons etc)
 - Done

Friends

- Offers to host VMs for free from:
 - Edinburgh @ Fluency
 - Boston @ TorwardEx
- One swap:
 - Bremen, Germany @ Fremaks
- Quickly up to 7 nodes after 1 month



#	Location	State	Protocols	In service
1	 England, London, A	Live	IPv4 IPv6	24th Sept 2013
2	 England, London, B	Live	IPv4 IPv6	2nd Oct 2013
3	 United States, Detroit	Live	IPv4 IPv6	9th Oct 2013
4	 United States, Boston	Live	IPv4 IPv6	13th Oct 2013
5	 India, Pune	Live	IPv4 IPv6	14th Oct 2013
6	 Germany, Bremen	Live	IPv4	23rd Oct 2013
7	 Scotland, Edinburgh	Live	IPv4 IPv6	25th Oct 2013

Friends

- “We can host something, but not a VM”
- Various issues:
 - Network engineers didn’t have access to VM hosts
 - No VM infrastructure at all
 - All VMs routed by hypervisor, can’t bridge to BGP routers
- Solution...
 - Raspberry PI, \$35
 - 512mb / 16gb SD Card
 - 800mhz (overclocked)



Raspberry Pi

- Sent RPI to Belfast – David Farrell @ Tibus
- Ran NSD3 great, ~200qps
- Adding / removing zones with NSD3 required service restart ☹️
- Swapped to PowerDNS, testing went ok
- pdns_control segfault'ing on ARM
- Gave in and moved to BIND



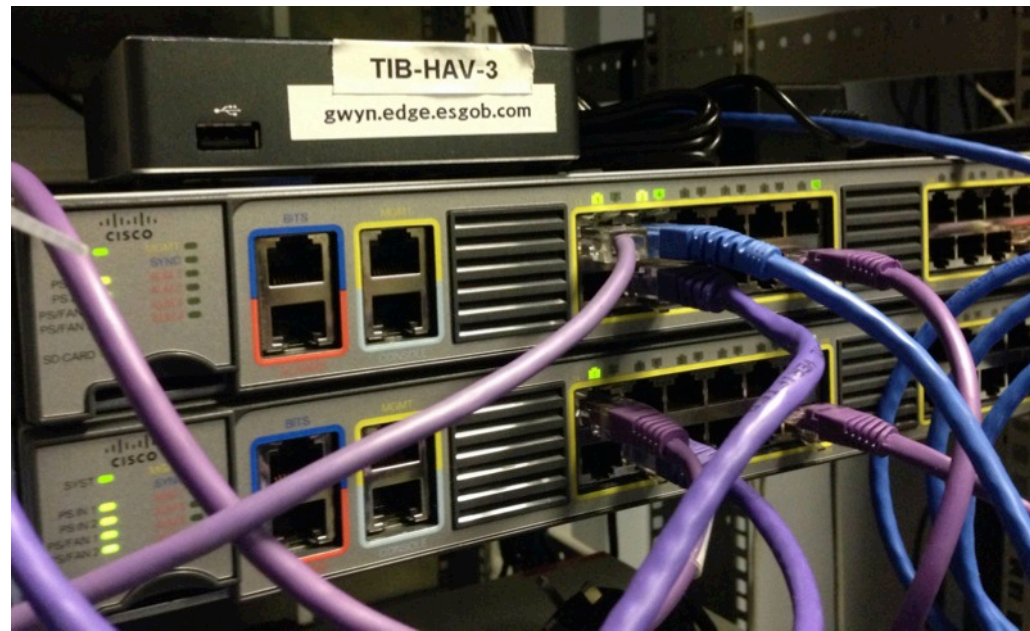
Time to peer

- Offer of a node at SFMIX from Matt Peterson
- San Francisco based IXP
- Couldn't turn this down.
- Zero U install
- Needed 2 NICs, IX + OOB
 - FitPC2i - perfect 😊
 - Atom 1.6ghz, 1gb RAM, 16gb SSD
 - \$150 on eBay
- Peering with:
 - HE.net, ISC, Unwired
 - Layer42, Lookout



Pi problems

- Puppet slow + loading zones taking too long
- More offers to host h/w nodes
 - Manchester, Andy Davidson @ Allegro
- Gigabyte BRIX
 - 1.8ghz, 4gb RAM, 30gb MSATA
- Swap out Belfast Pi



Still growing

- 11 live nodes
 - London C, VPS from Thomas Greer @ TSONE
- Bahrain on its way



Fun along the way

- First BGP customer for some of the LowEndTalk advertisers
 - Assisted educating them...
 - Prefer transit routes over customers, eek
- Some hosts don't have communities
 - Helping people get those implemented
- RIPE Atlas
 - Scheduled measurements to look at latency + CHAOS id.server

Market for VMs with BGP sessions

PORTLANEHOMEIP-SERVICESHOSTING

You decide what you wish to pay for, no terrifying invoices each month. You configure your server with the specifications you desire, adding the extra addons and then easily pay in advance wherever you are.

► CLOUD VPS

Mini 12€ /month	Medi 72€ /month
CPU 1 CORE	CPU 4 CORES
RAM 512 MB	RAM 4096 MB
Storage (SAN) 50GB	Storage (SAN) 50GB
Bandwidth 1000GB	Bandwidth 1000GB
Buy now	Buy now

Addons	
Storage (SAN)	Up to 200GB per drive (multiple hard drives per VPS)
Offsite Backup	Up to 100GB
Bandwidth	Up to 100TB
Private Network	10€ /month
BGP session	10€ /month
Additional IPv4	1€ /month for each additional IP

zettagridHOMECATALOGTRIALSCARTLOGIN

- CATALOG
- Server**
- Data Centre
- Backup
- Licences
- Network

Virtual Server

A dedicated single virtual machine for any purpose, on demand, ready in minutes. [Read More...](#)

Processor (CPU) ② 0.4 Ghz

Memory (RAM) ② 0.25 GB

Storage ② 100 GB

Operating System ②

Bandwidth ② 20 GB

Location

Summary

Computer
0.4 Ghz / Processor \$3.00
0.25 GB / Memory \$4.00
Storage
100 GB / iOSTOR-250 \$10.00
Operating System
CentOS Server 6 (\$0.00)
Network
20 GB / Traffic \$2.00
Location
Melbourne
Machine Name
Auto-generated

\$27.73/month
GST Inclusive

BGP Service

If you want to assign your own IP ranges to a ZettaGrid Cloud Server then you will need this option. Please note that you will require your own AS (Autonomous System) Number and IP Range.

BGP Service

AS Number

IP Range

SUMMARY

Monthly Cost	\$0.00
Setup Cost	\$0.00
Total Cost	\$0.00

Discoveries

- Easy to find budget anycasters
 - Webhost forums / LowEndTalk / bgp.he.net
- “Interesting” deployment methods
 - **Shared /24**
 - Dedicated IP, custom fwd/rev DNS, slaved zones
 - **Hosted /24**
 - Same as above but customer provides /24 to announce
 - **Shared /24, /32 tunneled**
 - Dedicated IP, tunneled by anycaster to customer via GRE
 - **Single /24**
 - Anycaster only has one /24, website + mail + mgmt in same space. DNS answered at edge, other IPs tunneled to another VM/dedicated box.

Discoveries

- Not all budget DNS hosting companies host every zone at the edge
 - Some host zones centrally and cache at the edge
- Many tunnel all DNS traffic back to a single location
 - Looks like anycast
 - Poor DNS performance

Not just DNS?

- Simple HTTP/HTTPS service to assist debugging
 - local.esgob.com
 - local4.esgob.com / local6.esgob.com
- Returns JSON via lighttpd on each node

```
nat@enw:~ $ curl http://local.esgob.com
{
  "city": "London",
  "country": "England",
  "countryiso": "gb",
  "flag": "england",
  "locationdisplay": "England, London, B",
  "ref": "q17f823b"
}
```

What next?

- Keep honest – host every zone at the edge
- Finish API + web interface
- Open source more code
- Support:
 - Multi master
 - TSIG inbound
- Public SOA checker
- Mix of routing and DNS daemons
 - configurable per node in RethinkDB
 - BIND, NSD4, KNOT, Quagga + BIRD

What next?

- Move to BGP sessions to full table
 - Archive MRT dumps
 - Maybe BIRD, don't need full table in the kernel
- Crypto loopback file system on remote nodes
- Fix v6 routing – peer with HE in Europe
- Find more friendly hosts:
 - Eastern Europe
 - Transit in west coast US
 - Especially in Africa, Asia + South America
 - Wales!

Currently costing just under \$480/yr

Questions?

<https://noc.esgob.com>

NAT-RIPE

@natmorris

@esgobltd

