Web hosting services
ActionAid UK

Request for Proposal/Quote

ActionAid UK
33-39 Bowling Green Lane,
London EC1R OBJ

Tel: 0203 1220561
www.actionaid.org.uk

Table of Contents

1 Copyright information

This document is the proprietary and exclusive property of ActionAid UK (referred as ActionAid UK) except as otherwise indicated.

No part of this document, in whole or in part, may be reproduced, stored, transmitted, or used for design purposes without the prior written permission of ActionAid UK.

The information contained in this document is subject to change without notice.

The information in this document is for information purposes only.

2 Document history

[Complete the items below as to authorship and necessary Glossary]

Version	Date	History	Name
1	10/04/2019	Initial document created	Tijana Miletic

3 Glossary

AA	ActionAid
AAUK	ActionAid UK

4 Organization information

ActionAid is an international charity that works with women and girls living in poverty. Our dedicated local staff are helping end violence against women and girls and changing lives, for good. We won't stop until women and girls are out of danger, out of poverty and on track to create the future they want.

More information on the organisation will be made available upon request or can be accessed through our website www.actionaid.org and www.actionaid.org.uk

5 Current situation overview

ActionAid UK's production and staging environments for its main website, payment platform and three microsites are currently hosted on dedicated servers with Rackspace, but this contract is reaching its end and a review of services is required.

The websites hosted are using LAMP.

SSL certificates and DNS for AAUK are managed by the internal IT department.

5.1 Websites

1. ActionAid UK main site
   https://www.actionaid.org.uk/
   Our main and largest site running on Drupal 7.64

2. Gifts in Action (often referred to as GIA)
   https://giftsinaction.org.uk
   Custom-built CMS on Zend PHP framework

3. Magento checkout site
   https://support.actionaid.org.uk/
   Magento 1.9.2.4 site which supports a large chunk of the donations and purchases put through https://www.actionaid.org.uk/

4. Weddings
   https://weddings.actionaid.org.uk/
   A very small site that has a custom-built CMS on the Laravel PHP framework.

5. Change a life
   https://changealife.actionaid.org.uk/
   A very small static HTML site

See Appendices 1, 2 and 3 for the full description of current infrastructure setup.

5.2 Current web traffic and transactions

The number of monthly users on our main site (actionaid.org.uk) ranges from 40,000 up to 120,000 and page views up to 225,000 but we would like to have enough capacity for accommodate surges that happen around busy times such as Christmas, or when we run emergency appeals.

Average monthly users: 59,000
Average monthly page views: 124,000

Current number of transactions: c. 16,000 a year split between 11,000 through Magento (including Christmas cards etc) and around 5,000 through either Stripe, PayPal or through Gifts in action website.

Peak pageviews: 12,000 a day
Peak visitors: 7,500 a day
Peak transactions: 800 a day

Average session duration (actionaid.org.uk and support.actionaid.org.uk): 01:50 minutes
Average number of pages per session (actionaid.org.uk and support.actionaid.org.uk): 2.09

We are aiming to increase our web traffic and transactions by 20% this year.

6 Goals and objectives

This request for proposal (RFP) is established on behalf of ActionAid UK to provide managed hosting solutions for production and staging environments for its main website, payment platform and smaller microsites.

We require a solution that delivers:

- 99.99% uptime
- SLA ensuring rapid response times for critical functions and 24-hour support
- Full compliance with PCI standards and data protection legislation
- Managed solution to ensure appropriate efficient support
- Scalable solution to manage increased demand
- Good value

7 Requirements detail

ActionAid UK is looking to review its current production and staging hosting provision and procure a suitable replacement. We welcome proposals which suggest how we can optimise the hosting specification to meet our requirements of high availability, security, ease of management and good value. The new hosting provision should be fit for purpose and not be over- or under-engineered.

Respondents are asked to specifically answer each requirement listed in this section in their response.

7.1 Service requirements

Please provide details of how you will meet each of the requirements below:

1. ActionAid UK would like to ensure high availability of the websites on the hosting platform. Our target is to have 99.99% uptime. Respondent needs to provide details about how the uptime guarantee is calculated, and whether and how the refunds are provided if this standard is not met.

2. Infrastructure should be scalable to support potential growth and higher traffic periods. Respondent needs to provide bandwidth information.

3. Respondent needs to specify response and resolution times and SLA terms proposed within a full SLA matrix covering P1 to P5 incidents.

   P1 incidents should be responded to within 30 minutes and resolved within 2 hours.

   A typical example of a P1 would be assets not loading on the website due to a storage drive failure or customers seeing an SSL warning and not being able to proceed to the website due to SSL certificate expiry.

   P2 incidents should be responded to within 2 hours and resolved within 4 hours.

   An incident on the staging environment could also constitute a P1 issue, for instance if a deployment which fixes a business-critical problem cannot be tested on staging due to a hardware or software failure.

4. ActionAid UK needs a disaster recovery procedure in place should the hosting setup fail. Respondent needs to specify backup and recovery procedures that would be in place to make sure normal service is restored as soon as possible.

5. ActionAid UK needs to ensure that its hosting provider has appropriate system monitoring, logging and alert systems in place.

6. The hosting provider needs to keep the hosting platform up-to-date with server and application security updates. This needs to be accomplished in a way that minimises downtime, particularly during busy periods.

7.2 Technical requirements

7.2.1 System setup

ActionAid UK needs a production and staging environment to host their websites. Production needs to meet the highest standards of performance and availability as discussed previously. Staging needs to mirror production as closely as possible, but doesn't need to meet the same levels of performance, availability or support. In your response to requirements listed below, specify how your solution will be implemented for production and staging environments respectively:

1. Current websites need to function on the new infrastructure as they do currently, achieving 99.99% uptime and fast performance for users, and fully supporting existing websites, their applications (Drupal, Magento, Zend and Laravel php framework) and services used (Solr, Stripe, sending email etc.). Responder needs to propose an appropriate infrastructure solution based on existing websites, capacity and redundancy requirements.

2. Current websites, website data and payment transactions need to be secure.
   Set up HTTPS for all sites and detail any other hardening and security implementation in the proposed solution.

3. Performance of the new platform is expected to improve in relation to what we have currently in a predictable and measurable way.

4. Provide full technical specification for hardware and software in your proposal (including OS, storage, processor, RAM, upgradability) and explain the level of virtualisation.

5. Testing on staging needs to provide an accurate representation of behaviour on production, which is to include the full test transaction journey.

6. Access to the staging platform needs to be limited to ActionAid UK staff (accessing the site from any location) and selected third parties. Access control to the staging platform needs to be setup through firewall/access list control to limit availability to specific IP addresses/networks (company office, third parties etc) or suggest an alternative solution.

7. Access to both production and staging platforms for internal and external developers who need to deploy changes, manage and troubleshoot the application layer (i.e. Drupal, Magento etc.) needs to be secure and easy to use. Optionally, ActionAid UK would like to ask respondents to enable developers to take snapshots and roll-back changes if necessary. If this is possible, please explain how this can be done.

8. Access to both production and staging platforms needs to be secure and managed. Recommend a secure and sensible user management policy where the respondent is responsible for adding and removing users and setting their privileges.

9. Developers need SSH and SFTP/SCP access on production and staging environments as we have at present.

10. Respondent needs to specify what type of files can be uploaded to the hosting platform and if there are any exclusions, due to malware protection for example.

11. ActionAid UK needs the production and staging platforms setup fully documented before the new platform becomes active.

7.2.2 Migration

1. The respondents are required to fully manage the production and staging platform migration from Rackspace for all sites.

2. Migration and the new hosting setup for production and staging environments needs to be completed, fully tested and ready to be put into use by 1st December 2019.

3. Disruption needs to be minimised during migration and provision needs to be made for making sure all the latest edits to content and assets are preserved.

7.2.3 Managed solution

1. ActionAid UK needs rapid and efficient response from its hosting service provider, with feedback of appropriate frequency, particularly for business-critical issues.

   For the production platform, ActionAid UK requires 24/7 support, particularly for business-critical issues.

   For the staging platform, ActionAid UK needs support during standard office hours, 9.30am to 5.30pm, Monday to Friday excluding bank holidays.

   ActionAid UK needs to be able to initiate support tickets by email or telephone.

2. ActionAid UK requires the respondent to manage all aspects of backup and restore management.

3. ActionAid UK needs flexible support for the managed hosting solution, fully supporting system administration and willing to assist with any overlap with the application layer.

4. Developers need to be able to deploy from the repository from AAUK offices and remotely. The current deployment method uses Capistrano scripts to deploy from the GitLab repository which is hosted internally. This method doesn’t currently allow remote deployment. Propose a new deployment method which will be secure and easy to use and is manually managed (we are not looking for a CI/CD solution at this time). AAUK is open to migrating the repository elsewhere to fulfil this requirement.

5. ActionAid UK needs regular automated backup processes appropriate for production and staging websites, which would include code and database data.

6. ActionAid UK needs to be able to add new sites as needed to the new hosting platform which will be supported by the hosting partner in the same way as the existing ones.

7. Optional: It would be useful if the hosting service provider can set up scripts for on-demand synchronisation of the database and database assets between the production and staging environments.

8. Optional: It would be a bonus to have a separate stand-alone development environment set up at the same time. This would need to be a low-grade environment that ideally mirrors the setup of the production/staging environments.

7.3 Legal requirements

1. We must be able to ensure continued compliance with PCI DSS standards.

2. Respondents must be prepared to sign data protection contractual clauses that cover the mandatory Data Processor requirements of Regulation (EU) 2016/679 (commonly and hereafter known as the GDPR) and the UK Data Protection Act 2018, as outlined by the Information Commissioner's Office (ICO). These include (but are not limited to):
    - Maintaining appropriate organisational and technical measures so as to ensure the security of the sites to ensure the protection of the rights of the data subject.
    - Providing assistance to ensure that ActionAid UK can meet its requirements with regards to data security breach reporting.
    - Accommodating any reasonable request by ActionAid UK to conduct an audit

3. Hosting should ideally be located in the UK, although European Union hosting can be considered.

7.4 Cost requirements

ActionAid UK would like to procure a hosting service which is good value and represents a cost saving in relation to the current hosting arrangement.

Respondents are asked to:

1. Provide 3 cost options, low-cost, medium-cost and high-cost, explaining benefits and compromises of each option

2. Specify fully itemised monthly costs, listing prices both with and without VAT

3. Specify fully itemised one-off costs, listing prices both with and without VAT

4. Proposals with costs less than £40,000 (incl. VAT) a year will be considered.

7.5 Term requirements

Respondents are asked to specify:

1. The initial term of contract. ActionAid UK will consider 3-5 year contracts.

2. The notice period.

8 Vendor instructions

Please provide a detailed proposal demonstrating:

- Understanding of the context and situation
- Specific answers to all requirement details in the section 7 of this document
- Sampling considerations and methods
- Three cost options, low, mid and high (listing prices both with and without VAT)
- Suggested timelines to meet the deadline of 1st December 2019
- Relevant project and personnel experience

Further information is available upon request (subject to NDA compliance).

8.1 Schedule for evaluation process

The expected timeline for the evaluation and decision-making process is as follows:

Process step	Date
RFP & Tender document information posted	11/04/2019
Opportunity for vendor questions	07/05/2019 – 31/05/2019
Deadline for RFP responses	By 05/06/2019
RFP response review	June, July 2019
Contract	End July 2019

8.2 Proposal submission & contact information

Please submit electronic copy of your proposal, including all supporting documentation to:

Name	Katherine Griffis
Company	ActionAid
Address	33-39 Bowling Green Lane, London, EC1R OBJ
Phone	02013 122 0561
E-mail	commercialuk@actionaid.org

Contact Person(s) for any questions about the RFP by email only.

Sally O'Connell	Sally.oconnell@actionaid.org

Proposal Format
A vendor's internal template for responses to RFP will be accepted. The New Supplier Questionnaire provided herewith has to be answered and attached along with the proposal.

Notifications
Vendors will be notified regarding requests for additional information, formal presentations and the outcome of the selection process through email.

9 Basis of award

All proposals become the property of ActionAid UK and will be evaluated by the RFP Review Team. Evaluation and selection of vendors to provide products and services as defined in this RFP to ActionAid UK will be based on the following criteria, which are given in no specific order.

- Quality of proposed solution
- Expertise of proposed team
- Price
- Cultural fit

Weighting to scoring will be determined along these areas:

Quality 40%	Experience 20%	Price 30%	Cultural Fit 10%

The RFP Review Team reserves the right to accept or reject any or all RFPs received.

The RFP Review Team reserves the right to negotiate with respondents to this RFP, within the requirements of the RFP, to best serve the interests of ActionAid UK. However, vendors **must not assume** an opportunity to negotiate and are cautioned to submit their proposals on a best and final basis since an award or decision is likely to be made without further negotiation based on pricing and terms of the original submittals. Accordingly, all requirements must be included with your initial offer.

All proposals submitted will be considered to be proprietary by ActionAid UK and will not be released to any outside party, in part or in total unless required by law. Neither the transmission of this RFP to a prospective bidder nor the acceptance of a reply shall imply any obligation or commitment on the part of ActionAid UK.

If vendor needs to take exception to anything under the RFP, these exceptions must be clearly identified on the RFP response.

All prices and conditions must be shown.

10 Appendix 1: Rackspace production cluster: detailed architecture overview

10.1 Rackspace cluster: outline

This section outlines the basic technical details of the ActionAid UK Rackspace server infrastructure. It is intended to give a broad overview of how things have been set up and why they've been set up that way.

ActionAid UK have purchased managed hosting on physical servers from Rackspace. The package consists of four servers, a firewall, a load balancer, and access to a SAN.

Rackspace are contracted to manage all the details of the network, server hardware and the software from the OS up to Apache. ActionAid UK, or a subcontractor, are expected to manage anything else.

The basic design of this cluster prioritises resiliency. To this end everything except the load balancer and firewall are duplicated. In terms of performance, the amount of hardware available should be considerably more than is needed to serve the volume of requests AAUK sites receive.

10.1.1 Basic hardware and software details

Figure 1 (below) gives a diagrammatic overview of the server hardware and software components. There are two frontend web servers (web1, web2), a primary database server (db1), a backup database server (db2), a firewall, a shared SAN and a load balancer with SSL termination.

Access through the firewall is permitted on ports 80 and 443 to the load balancer for all IPs. Direct access to any servers is only permitted from AAUK approved IP addresses for administrative purposes.

The Load Balancer is also the termination point for SSL requests. Requests are distributed evenly across the two web servers (web1, web2) by the load balancer.

The two web servers contain most of the web application software. That is the front-end HTTP cache for Drupal (Varnish), Apache and PHP with the various necessary extensions, and Couchbase, which serves as a clustered high speed in memory cache, similar to Memcached. The reason Couchbase is clustered is to absorb writes to the Drupal cache as well as reads. The cluster is perfectly capable of functioning on only one web server. Requests should automatically go to the other server in the cluster if one of the servers fails.

The two backend servers serve as an active/passive MySQL cluster, and an active/passive NFS cluster. DB1 is the primary MySQL node. DB2 the primary NFS. Both are configured with virtual IP addresses, so that in the event of a server failure, the other server assumes the relevant IP address and carries on serving requests. Failure of either of the backend servers should not affect NFS or MySQL requests for more than 30 seconds or so, after which, operations should transparently move to the other server in the cluster.

Additionally, an Apache Solr instance, running inside Tomcat, is configured on db2 to provide the Solr backend for the Drupal site search. This is the only part of the system that's not currently clustered.

Rackspace manage all hardware, including firewalls, servers, load balancers and networking equipment. They manage all the software on the load balancer, and the SAN. They are contracted to manage only some of the software on the servers. This includes The Operating System (Red Hat Enterprise Linux 6), including applying necessary software/security patches to:

- Apache
- MySQL
- NFS
- PHP

They are not contracted to provide support for (though often willing to look at if asked)

- Varnish
- Couchbase
- Solr
- Drupal
- Any other applications we might run

10.1.2 Basic hardware diagram

Figure 1: Basic Hardware Diagram

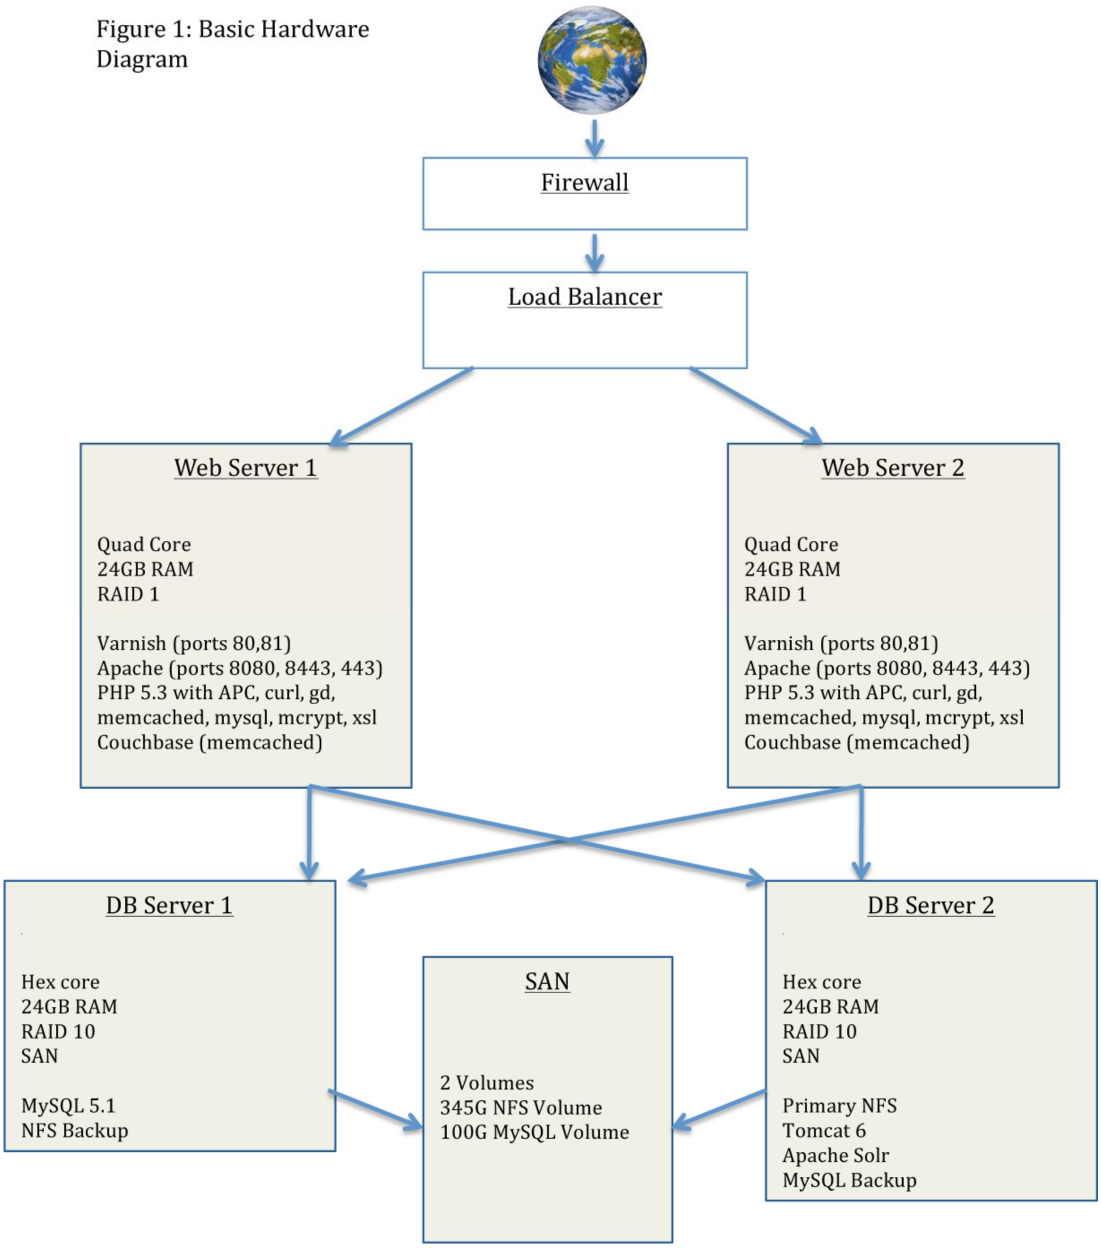

10.1.3 Anatomy of an HTTP(S) Request

It's probably a good idea to go into a little more detail and look at how HTTP(S) requests will actually flow through this cluster at this point, in order to establish what this design is intended to accomplish and why.

Figure 2 (at the end this section) illustrates how a request processes through the various parts of the cluster. Further detail is included below. The configuration of these individual components is discussed in more detail later.

Not illustrated in the diagram, but mentioned here for completeness is the firewall. This is a Cisco 350Mbit hardware firewall, which can be administered through the my Rackspace portal. Currently, only requests on ports 80 and 443 to load balancer IP addresses are globally allowed. Direct access to any other servers in the cluster is locked down to ActionAid UK IP addresses.

Once an HTTP request is through the firewall, it hits the load balancer. This device is administered solely by Rackspace support staff. The load balancer distributes requests across the frontend web servers evenly. It should not be assumed that all requests in a session, or from a given IP will end up on the same web server. Any shared state must therefore be stored in Couchbase, on the NFS layer, or in MySQL.

SSL connections terminate and are decrypted on the load balancer. Once decrypted, SSL requests are sent to port 81 on a web server (with a X-Forwarded-Proto header added indicating that the protocol was HTTPS), standard HTTP requests are sent to port 80.

The web servers do most of the heavy lifting. Requests on ports 80 and 81 are not served directly by Apache, instead an HTTP accelerator/cache named Varnish listens on those ports. Varnish is used because Drupal is extremely slow. As the majority of requests to actionaid.org.uk are for static content an HTTP cache can effectively mitigate this. When a request hits Varnish, it can basically do one of 3 things. It can serve the request entirely from its internal cache. It can proxy the request to Apache and cache the result. Or it can simply proxy the request and not cache. The "vcl" Varnish configuration file determines which of these actions is performed on a per server basis. Very broadly, requests that contain cookies (except those set specifically to be ignored at the vcl level) cannot be cached. Certain paths (e.g. admin pages) are excluded from caching at the vcl level. Other paths (e.g. urls ending in .jpg/gif/png), are set to always be cached.

The majority of requests will (hopefully) be served out of the Varnish cache. If an object is not found in the cache, the cache will proxy the request on to Apache. Apache is set up in a pretty standard way, except that it's configured to run on ports 8080 and 8443. Standard HTTP requests should go to port 8080, decrypted SSL requests should go to port 8443. Requests received on port 8443 have certain environment variables are set in order to notify PHP applications (or any other web applications) that they should be using https URLs and so forth.

Apache runs with the standard mod-php module to enable PHP files to be served. PHP has the APC opcode cache enabled to speed up compilation.

Drupal is configured to use the PHP Memcached extension to access Couchbase and store volatile cache data there. Couchbase is clustered across the web servers to ensure that data retrieved from it will be consistent across servers.

Shared file storage between the web servers for uploaded files and so forth is implemented using an NFS mount at on all frontend servers. This mount points to the virtual IP of the active/passive NFS cluster. The actual files are stored on the SAN, which is accessible from both backend servers (Db1, Db2), and usually mounted on Db2.

MySQL requests are served from a virtual IP as well, which will point to Db1, the primary MySQL server, in normal operation. Again, the actual file system level storage for the database is on the SAN to enable failover to Db2. This should be transparent to any web app, which simply needs to point itself at the correct IP.

The final piece of the puzzle is Solr, which is a Lucene based search system written in Java. This is served through Tomcat on port 8080 and is accessed directly by IP from Drupal.

Figure 2, a basic visualisation of this is shown below.

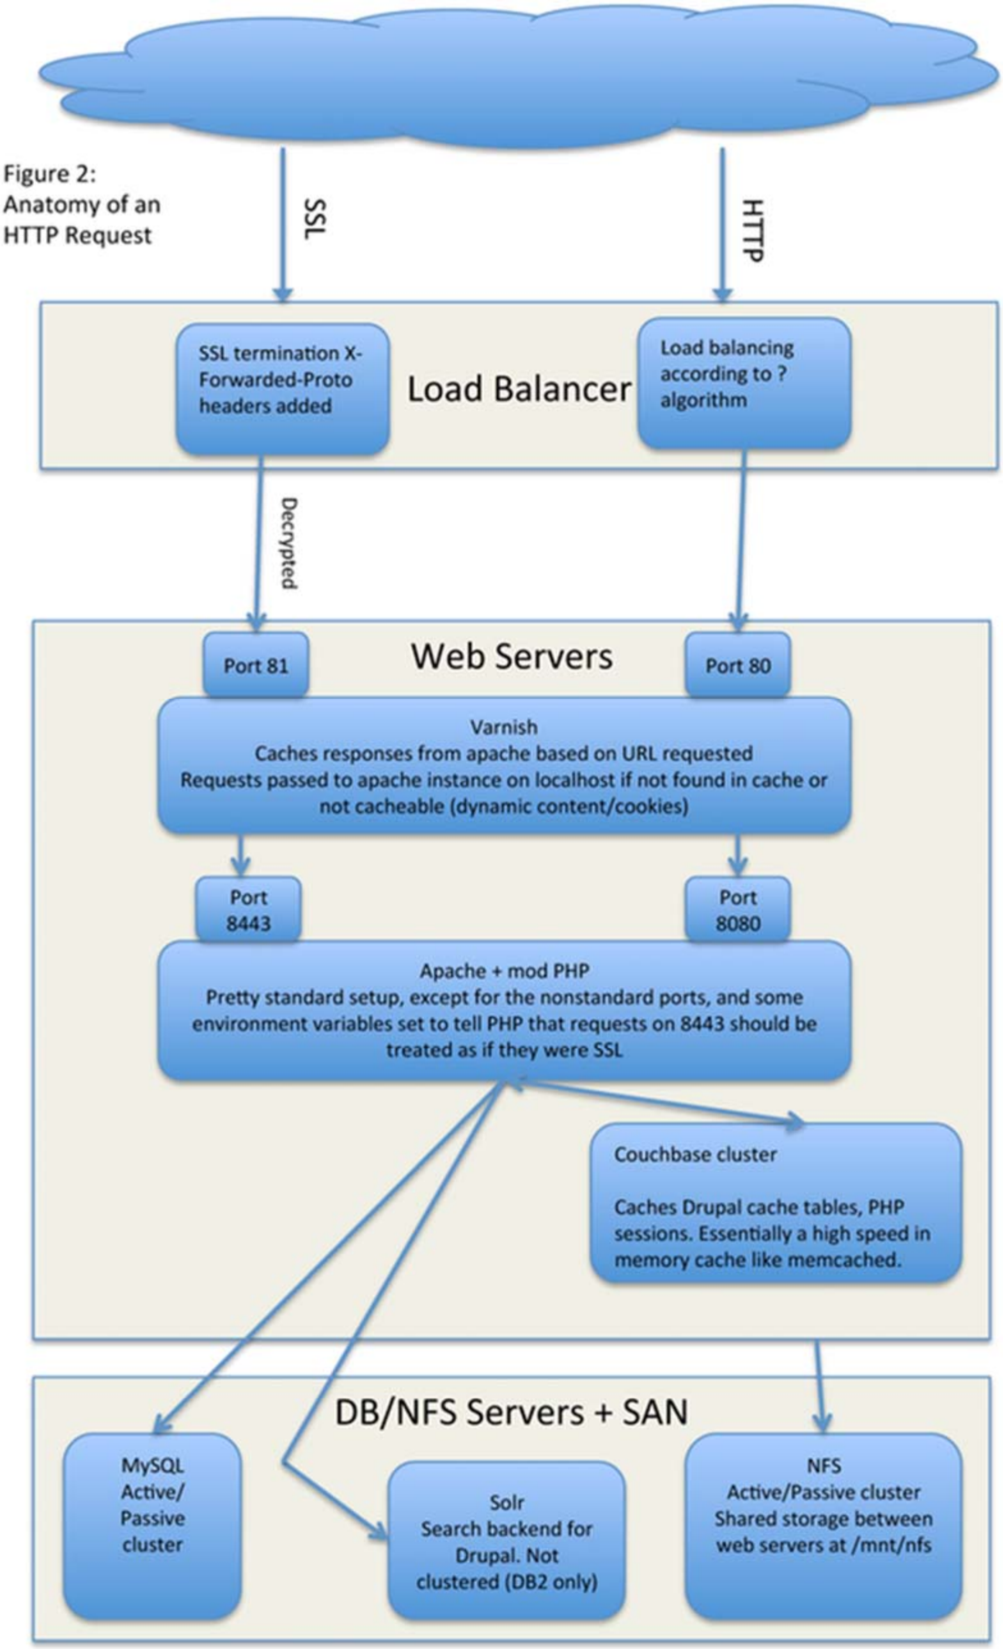
SSL
HTTP
SSL termination X-Forwarded-Proto headers added
Load Balancer
Load balancing according to ? algorithm
Decrypted
Port 81
Web Servers
Port 80
Varnish
Caches responses from apache based on URL requested
Requests passed to apache instance on localhost if not found in cache or not cacheable (dynamic content/cookies)
Port 8443
Port 8080
Apache + mod PHP
Pretty standard setup, except for the nonstandard ports, and some environment variables set to tell PHP that requests on 8443 should be treated as if they were SSL
Couchbase cluster
Caches Drupal cache tables, PHP sessions. Essentially a high speed in memory cache like memcached.
DB/NFS Servers + SAN
MySQL Active/ Passive cluster
Solr Search backend for Drupal. Not clustered (DB2 only)
NFS Active/Passive cluster Shared storage between web servers at /mnt/nfs

Figure 2: Anatomy of an HTTP Request

Basic device information

10.2 Firewall

The firewall is a standard issue Rackspace hardware firewall, a Cisco ASA 5520 with 300Mbps throughput.

All configuration for this device is performed by Rackspace support, who are contacted through the ticketing system.

10.3 Load balancer

The Load Balancer is a standard issue Rackspace hardware load balancer. Unknown designation.

All configuration for this device is performed by Rackspace support, who are contacted through the ticketing system.

10.4 Linux servers

There are 4 Linux servers in the cluster: 2 Web Servers, and 2 DB/NFS servers forming an active/passive cluster.

Cluster wide server configuration notes

10.5 Rackspace standard setup

Rackspace have their own standard server setup. Some points to note about the default setup are:

- Red Hat Enterprise Linux
- Yum software update configured to run every night
- Apache, PHP, Perl, Python installed and configured according to Rackspace standards. Apache runs at startup on all boxes.
- MySQL installed on DB servers as per Rackspace standard
- Password based ssh access as standard (no root login), su to root allowed for all users, ActionAid UK have full root access to all machines
- No iptables rules. All network rules implemented on the hardware firewall
- VSFTPD (FTP daemon) installed and configured to allow login by most non-system users
- SELinux disabled, enabling is unsupported and not recommended
- Managed backup for both filesystems and MySQL databases
- Rackwatch monitoring
- Server hostnames set to *.actionaid.org (should have been actionaid.org.uk)

10.6 Changes to Rackspace standard setup

As part of the setup process ActionAid UK have made, or instructed Rackspace to make, several changes to the standard setup in order to increase security and facilitate our use of the cluster. These changes are outlined briefly here, so there’s a list in one place. Those changes that involve software configuration changes are covered in more detail in section 2.

- Password based ssh login disabled. All ssh access must now be via key based login
- Use of su binary restricted to users in wheel group by configuration.
- All users in wheel group granted full sudo rights “%wheel ALL=(ALL) ALL”

- Use of FTP (VSFTPD) locked down to SSL connections only, login only allowed for users listed
- Host names for the two web (frontend) servers changed to *.actionaid.org.uk. Host names for the DB servers not changed at Rackspace’s request
- As a result of above hostname change, idmapd.conf on backend servers changed to read Domain = actionaid.org.uk to ensure NFS permissions work correctly
- The following software packages installed on web servers: php-pecl-memcached php-mcrypt php-pecl-apc varnish drupal6-drush git couchbase-server-enterprise
- Apache configuration changes on web servers: Many and varied. Standard ports changed to 8080, 8443 and 443, configuration changed.
- Varnish set to listen on ports 80,81
- Mysql configuration changes to my.conf
- Couchbase clustered across the two frontend web servers
- NFS mounted on all frontend servers. Linked to from vhost directories with symlink

10.7 Web servers

Web1 and web2 are a pair of redundant web servers, and have been configured pretty much identically.

The following software has been installed on each server:

- RHEL 6
- Apache 2
- PHP 5.6.29
- Varnish 3
- Couchbase server 1.8

Couchbase server is clustered across the two web servers and set to auto failover. The two servers share an NFS mount for common content. Other than that, the two web servers operate independently.

Web server 1
Device Name: web1
Device Spec: DELL PowerEdge R710
Single Socket Quad Core Intel Xeon L5520 2.26GHz
2x300GB SAS 15K RPM Drive in RAID1
Red Hat Enterprise Linux 6

Web server 2
Device Name: web2
Device Spec: DELL PowerEdge R710
Single Socket Quad Core Intel Xeon L5520 2.26GHz
2x300GB SAS 15K RPM Drive in RAID1
Red Hat Enterprise Linux 6

10.8 Backend servers

There are two backend servers, db1 and db2. They share the duties of NFS and MySQL server in an active/passive configuration, with db1 as the primary database server, and db2 as the primary NFS server.

Failover is achieved by means of a SAN and the use of virtual IP addresses. If one of the servers fails the other mounts the appropriate SAN volume, and assumes the IP address of the service being provided.

The Solr (lucene based search) server is hosted on Db2. Inside a Tomcat container on port 8080. This service is not redundant.

Backend server 1
Primary MySQL server
Device Name: db1
Device Spec: DELL PowerEdge R710
Single Socket Six Core Intel Xeon E5645 2.4GHz
4x300GB SAS 15K RPM Drive RAID10
Red Hat Enterprise Linux 6

Backend server 2
Primary NFS server Sole Solr server. Apache Tomcat, Solr.
Device Name: db2
Device Spec: DELL PowerEdge R710
Single Socket Six Core Intel Xeon E5645 2.4GHz
4x300GB SAS 15K RPM Drive RAID10
Red Hat Enterprise Linux 6

10.9 Software configuration

This section gives details of how important elements of the software stack have been configured, particularly where such configuration differs from standard configurations.

10.9.1 Apache

Obviously, the web server is a pretty important element in any web application deployment. The Apache configuration on the Rackspace cluster has been changed from the default in significant ways, which may cause problems if you're not fully aware of them.

Firstly, apache runs on 3 ports. Port 8080 serves standard HTTP requests (Varnish cache listens on port 80). Port 443 serves HTTPS requests as normal. Port 8443 serves HTTPS requests that have already been decrypted at the load balancer. Apache does not listen on port 80.

Apache config has been modified from RHEL standard to look for virtual host configuration files. All new Virtual Hosts should be set up there, named [primary-domain-name].conf.

To account for the various non-standard ports Apache is using, the UseCanonicalName directive has been set to On. It is therefore vitally important that all virtual hosts set the ServerName directive to reflect the port clients will be coming in on. Strictly speaking it's perfectly possible to set up Apache to run without this, but it's likely that setting the port at the apache config level will eliminate some sources of error.

e.g. for a vhost running on port 8080, receiving requests coming into the LB on port 80
ServerName [primary-domain-name]:80

For a vhost running on port 8443, receiving requests coming from the LB on port 443
ServerName [primary-domain-name]:443
SetEnv HTTPS on
SetEnv SERVER_PORT 443

For a vhost running on port 443, serving standard SSL requests
ServerName [primary-domain-name]:443

Additional changes from standard Rackspace/RHEL apache config:

- KeepAlive set to on (this should be more efficient with Varnish)
- Timeout set to 180 seconds
- Listen 80 remove
- Listen 8080, 8443 added
- NameVirtualHost *:8080, NameVirtualHost *:8443, NameVirtualHost *:443 added
- Default vhosts just serve an HTML file linking users to www.actionaid.org.uk

10.9.2 Varnish

Varnish is installed on all frontend web servers and listens on ports 80, 81. The security updates are pulled in by the standard yum update scripts.

On the ActionAid UK cluster, Varnish has been set up with a somewhat complex vcl. This is what it's supposed to do:

- There are two backends. Default, for requests on port 80 and fakessl for requests on port 81. Each is cached separately for obvious reasons.
- Varnish adds some request headers for Apache, indicating whether the request was http or https. Specifically set X-Forwarded-Port and X-Forwarded-Proto
- Anything that isn't a GET or HEAD request is passed straight through to Apache (we don't want to cache POST)
- Any Google Analytics URL parameters or cookies are stripped off the request (our server doesn't need to see these)
- Some generic cleanup is performed on the URL
- Accept-Encoding is normalized (and stripped completely for a list of non-compressible file types)
- We then pull in different vhost specific configuration files based on the host header. If we don't recognise the host requested the request is simply passed to Apache without caching
- Drupal specific vcl removes cookies for all document/image/javascript requests and excludes certain paths from caching
- Specific cookies that Drupal does not need to see are removed for example those set by PPC advertising scripts). It is expected that this list will need to be updated from time to time

- We hash on url, backend, host and encoding
- Varnish adds an x-cache header specifying “cached” or “uncached” to responses and unsets some php/apache headers
- On error, an error vcl is loaded with an error message

10.9.3 PHP

PHP is installed in both CLI and mod-php flavours on all frontend servers.
Some notes on important/changed configuration settings for PHP and APC follow.

- Safe mode, magic quotes, register globals and any other PHP misfeatures disabled, as per defaults.
- APC enabled, SHM size 256M
- APC.stat set to 0. This means that code updates will not take effect until the APC cache is flushed.
- Max_execution_time: 120
- Memory_limit: 512MB
- error_reporting = E_ALL & ~E_DEPRECATED, display_errors = Off, log_errors = On
- upload_max_filesize = 32M

10.9.4 Couchbase

Couchbase is installed on all frontend web servers.

There’s a single default bucket. 2GB per server. The two servers are set up as a cluster, auto failover is enabled with a timeout of 60 seconds.

10.9.5 MySQL

Served from the backend active/passive cluster.

A short list of changes made to the MySQL configuration:

- Character-set-server utf8
- Query cache size increased to 64M, would increase further, but apparently too much can hurt performance too
- Tmp-table-size increased to 64M
- Max-allowed-packet increased to 32M
- Key-buffer-size increased to 256M, which is likely bigger than the indexes will ever get, but we have enough RAM
- InnoDB cache increased considerably for Drupal, this is probably a key setting which could go higher
  innodb-buffer-pool-size = 8192M
  innodb-log-file-size = 512M
  innodb-log-buffer-size = 16M
  Note that changing innodb log size requires deleting all innodb logs. Changing this and blindly restarting mysql will bring down the cluster.
- innodb-flush-method = O_DIRECT

- This bypasses OS level caching for InnoDB tables, which should be more efficient, given the large amount of RAM already allocated for that purpose.
- innodb-buffer-pool-instances = 4
- innodb-thread-concurrency = 12 (hex core with SMP)
- IMPORTANT: Max-connections set to 512, this prevents errors on the frontend servers at high concurrent loads.

10.9.6 NFS

NFS is handled by an active/passive cluster set up in a very similar way to the MySQL cluster. Actual storage is mounted at /san/nfs, by default on Db2.

Administration of the NFS cluster is handled by clusvcadm.

10.9.7 Solr

Apache Solr is a search system based around the widely used Lucene library. It’s written in Java, therefore we run it inside Tomcat. It’s the only part of this system that has absolutely no redundancy. There’s no way to do an active/passive arrangement in the same way. So a single instance actually seems like the most reliable option here. The instance is installed on db2. The Solr application code is contained in the archive apache-solr-3.6.1.war. This is the file that will need to be updated if you wish to update Solr. This is not from a repository, the .war is simply downloaded from http://lucene.apache.org/solr/downloads.html (use wget). Tomcat configuration is at /etc/tomcat6/Catalina/localhost/solr.xml. There is currently no access control.

11 Appendix 2: Rackspace staging cluster

This section outlines the basic technical details of the ActionAid UK staging servers. These servers have been set up on the Rackspace cloud, using the smallest practical instances.

All of these instances are unmanaged in order to save on costs. The AAUK web developer/technical lead are expected to perform all system administration on these servers.

This document is intended to be read alongside the documentation for the main Rackspace cluster, items are detailed only where the staging setup differs from the live setup.

Basic hardware and software details

All 4 staging servers are running the same operating system (Red Hat Enterprise Linux 6) as the live servers. Most of the software setup from the live cluster is replicated exactly, with the notable exception of the active/passive clustering on db1/db2 for MySQL and NFS.

All 4 staging instances are next generation Rackspace cloud servers running RHEL 6. They are currently running the minimum possible hardware spec (512MB RAM/20GB disk). The exact setup is detailed in the hardware diagram overleaf.

Basic hardware diagram

Figure 1: Basic Hardware Diagram

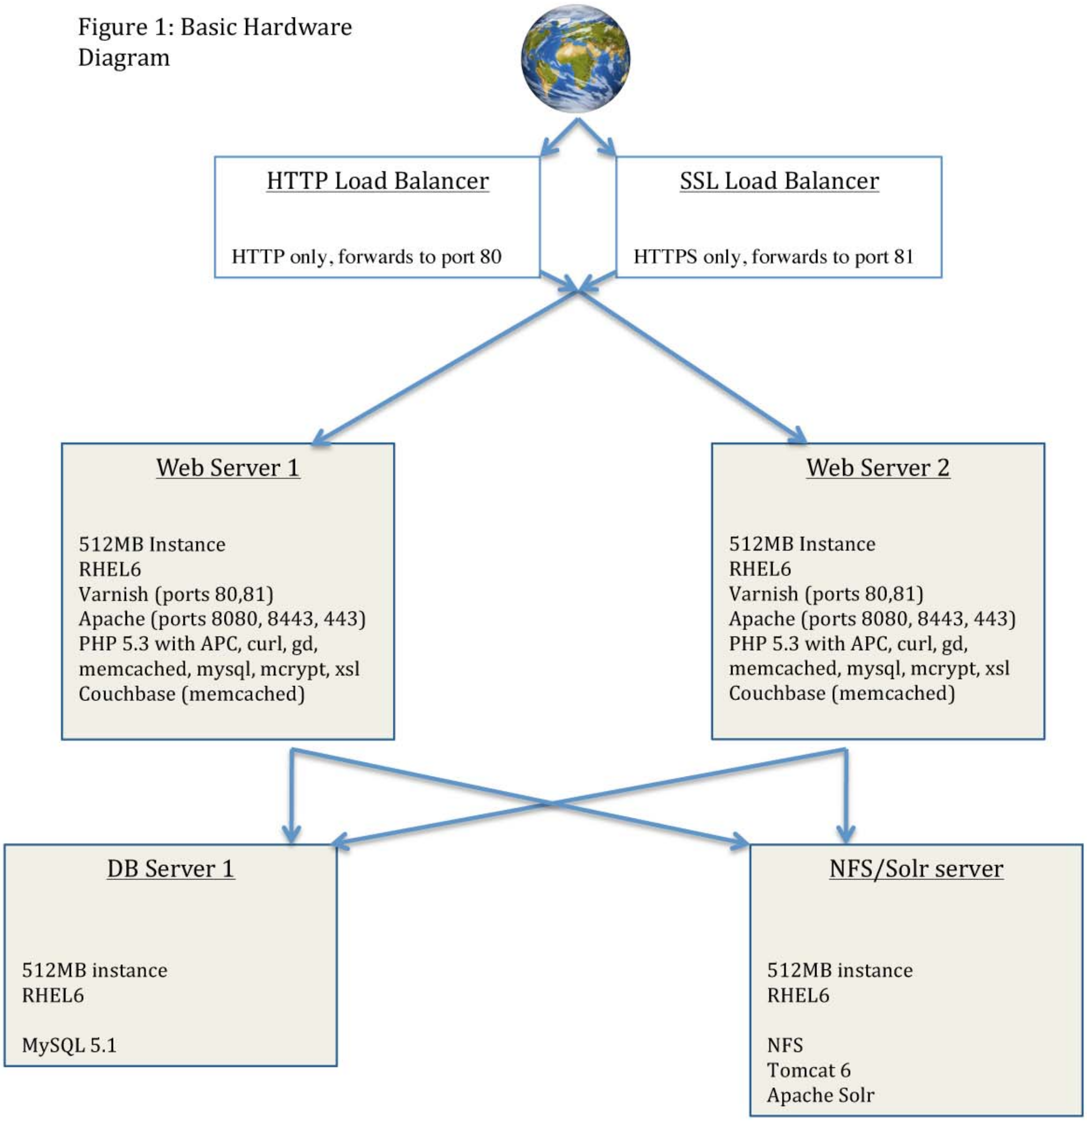

Differences from live setup:

- No active/passive clustering of MySQL and NFS servers. This could theoretically be accomplished on CentOS using the same tools as on the live system, but acquiring virtual IP addresses might be tricky. If you happen to have a lot of time on your hands, consider implementing this.

Differences from live setup:

- No active/passive clustering of MySQL and NFS servers. This could theoretically be accomplished on CentOS using the same tools as on the live system, but acquiring virtual IP addresses might be tricky.
- Less RAM
- Firewalling implemented using Shorewall (iptables frontend) on individual servers, rather than a single dedicated hardware firewall.
- Rackspace "cloud" load balancers function slightly differently to the hardware LB used on the live site, but they should appear the same from the web application's perspective.
- Slightly different software packages.
- The main software packages (Apache, PHP, Couchbase, MySQL, etc) are all installed and configured as per the live system. Rackspace have their own setup for their dedicated servers though and it's not practical to replicate every single installed package and its configuration, so some minor differences in software setup may exist.
- Different domain names and paths to virtual hosts named after those domain names. Domain names changed in Varnish and Apache configs

Firewall

The Rackspace cloud doesn't provide for hardware firewalling. Firewalling on the staging servers has therefore been implemented using Shorewall (version 4.5), which is a frontend for the standard Linux iptables firewall.

The staging servers have three IP addresses, one public IPv4, one public IPv6 and one IPv4 internal to the Rackspace network. Internal traffic is not billed for, any traffic over the public interface is charged for per GB.

The internal Rackspace network is shared by a great number of cloud servers, and should not be considered secure.

Shorewall has been set up to recognise 4 zones (networks in which different groups of firewall rules apply), these are

- net - The Internet
- rack - The Rackspace internal network
- admin - AAUK IP addresses allowed access to privileged ports
- aasta - The 4 staging servers

Varnish/Apache changes from live

All that really needs to be changed from the live environment is the hostnames and paths to the document root for configuration files.

Staging server configuration files

These are stored in the git repository.

12 Appendix 3: Rackspace production and staging hosting summary

Production

Devices

Device Name	Nickname	Type	Region
web1	web1	DELL PowerEdge R710 Linux	LON3
web2	web2	DELL PowerEdge R710 Linux	LON3
lbal1	lbal1	Load-Balancer	LON3
db1	db1	DELL PowerEdge R710 Linux	LON3
db2	db2	DELL PowerEdge R710 Linux	LON3
san1	san1	Managed Storage	LON3
SQLClus1	SQLClus1	Virtual Server Cluster	LON3
NFSClus1	NFSClus1	Virtual Server Cluster	LON3
fw2	fw2	Firewall - Cisco ASA	LON3
montr1	montr1	Custom Monitoring	LON3

Lba1

Brocade ADX 1000 Series	Brocade ADX 1000 Series, Brocade Options: Basic
IP address quantity	Number of IPs, # IPs: 1
Load-Balancer Required	Load-Balancer Required
Monitoring	Rackspace Monitoring
Data Center	London, UK (LON3)

San1

Managed Storage Required	Managed Storage Required
Managed storage space	Silver Shared SAN, SAN Storage in GB: 450
Data Center	London, UK (LON3)

Web1 & Web2

CommVault	Managed Backup Agent - CommVault, Application License Provider: Rackspace
Dell Memory	24 GB DELL RAM, GB Memory: 24
Dell Servers	DELL PowerEdge R710 Linux Required
Dell Servers	Single Socket Quad Core Intel Xeon L5520 2.27GHz, #Processors: 1, #Cores per Proc: 4
Hard Drive	300GB SAS 15K RPM Drive, HDD RPM: 15000, GB Hard Drive: 300
Hard Drive	300GB SAS 15K RPM Drive, HDD RPM: 15000, GB Hard Drive: 300
Hard Drive Size	3.5 in. Hard Drives
IP address quantity	Number of IPs, # IPs: 1
Included Bandwidth	Included Bandwidth, GB Bandwidth: 0

Linux OS	Red Hat Enterprise Linux 6
Managed Backup Retention	2 weeks retention - Onsite
Monitoring	Rackspace Monitoring
Non-Managed	Non-Managed
OOSBP	Out of Schedule Backup Permitted
Point Release	RHEL 6 Base - Opt-in (nightly sync with Red Hat)
RAID Configuration	RAID 1
RAID Group 1	RAID Group 1, Drives: Hard Drive 2, Drives: Hard Drive 1, RAID Type: RAID 1
Red Hat Single Socket License	Red Hat Single Socket License
Support	Fanatical Support
Advanced networking	1000Mb Port
Backup type	Weekly Full + Daily Differential
Data Center	London, UK (LON3)

Fw2

Firewalls	Cisco ASA 5520
Firewalls Required	Firewall - Cisco ASA Required
IP address quantity	Number of IPs, # IPs: 1
Monitoring	Rackspace Monitoring
Net Device Upgrade Required	Net Device Upgrade Required
RackConnect	RackConnect v2
Advanced networking	1000Mb Port
Data Center	London, UK (LON3)

Db1 & Db2 specification

CommVault	Managed Backup Agent - CommVault, Application License Provider: Rackspace
Dell Memory	24 GB DELL RAM, GB Memory: 24
Dell Servers	DELL PowerEdge R710 Linux Required
Dell Servers	Single Socket Six Core Intel Xeon E5645 2.4GHz, #Processors: 1, #Cores per Proc: 6
Hard Drive	300GB SAS 15K RPM Drive, HDD RPM: 15000, GB Hard Drive: 300
Hard Drive	300GB SAS 15K RPM Drive, HDD RPM: 15000, GB Hard Drive: 300
Hard Drive	300GB SAS 15K RPM Drive, HDD RPM: 15000, GB Hard Drive: 300
Hard Drive	300GB SAS 15K RPM Drive, HDD RPM: 15000, GB Hard Drive: 300
Hard Drive Size	3.5 in. Hard Drives
IP address quantity	Number of IPs, # IPs: 1
Included Bandwidth	Included Bandwidth, GB Bandwidth: 0
Linux OS	Red Hat Enterprise Linux 6
Managed Backup Retention	2 weeks retention - Onsite
Monitoring	Rackspace Monitoring
Non-Managed	Non-Managed
OOSBP	Out of Schedule Backup Permitted
Point Release	RHEL 6 Base - Opt-in (nightly sync with Red Hat)
RAID Configuration	RAID 10
RAID Group 1	RAID Group 1, Drives: Hard Drive 3, Drives: Hard Drive 2, Drives: Hard Drive 1, Drives: Hard Drive 4, RAID Type: RAID 10
Red Hat Single Socket License	Red Hat Single Socket License
Support	Fanatical Support
Advanced networking	1000Mb Bonded NICs (RH Cluster, DRAC on ExNet)
Backup type	Weekly Full + Daily Differential
Managed storage	SAN HBA - Clustered Servers (Dual HBAs)
Data Center	London, UK (LON3)

Staging

Cloud servers

web1 & web2

RAM	1GB
CPU	1 vCPU
System Disk	40 GB
Network	120 Mb / s

db1 & db2

RAM	512MB
CPU	1 vCPU
System Disk	20 GB
Network	80 Mb / s

Cloud load balancers

Name	Address	Protocol:Port
http_staging	Same virtual server	HTTP:80
staging		HTTP:81
stagingsupport-http	Same virtual server	HTTP:80
stagingsupport-https		HTTP:8443

Databases

db.actionaid.org

Type	Percona 5.6
RAM	1 GB
Disk	2.9GB of 20GB used

13 Appendix 4: File space usage on production

Website sizes

Drupal (actionaid.org.uk) 2GB + 32GB for file storage

Magento (support.actionaid.org uk) 2.9GB + 84MB for file storage

Web1 on production

```
Filesystem              Size  Used Avail Use%

                        271G  189G   69G  74%
tmpfs                    12G   44K   12G   1%
/dev/sda1               239M   64M  163M  29%
                        2.0G  4.3M  1.9G   1%
                        345G   33G  295G  10%
```

Web2 on production

```
Filesystem              Size  Used Avail Use%

                        271G  221G   36G  87%
tmpfs                    12G   44K   12G   1%
/dev/sda1               239M   64M  162M  29%
                        2.0G  3.9M  1.9G   1%
                        345G   33G  295G  10%
```

Db1 on production

```
Filesystem              Size  Used Avail Use%
                        545G   86G  432G  17%
tmpfs                    12G   30M   12G   1%
/dev/sda1               239M  143M   83M  64%
                        2.0G  3.1M  1.9G   1%
                         99G  5.9G   88G   7%
                        345G   33G  295G  10%
```

Db2 on production

```
Filesystem              Size  Used Avail Use%

                        545G  7.9G  510G   2%
tmpfs                    12G   26M   12G   1%
/dev/sda1               239M  143M   83M  64%
                        2.0G  3.1M  1.9G   1%
```