



Dual-CMS Routing

How We Took Our New CMS Live
One URI at a Time



About me

- Systems Integration Architect @ UWG
- SysAdmin of both Windows and Linux
- Recovering SPARC Solaris Admin
- F5 Load Balancer, SAN, VMware, & Hyper-V Admin
- Coder
 - Ruby & PowerShell mostly
 - PHP-dabbler
- Automator of things with Puppet
- Linux & Open Source Advocate
- Data Center junkie
- Mac aficionado



The Challenge

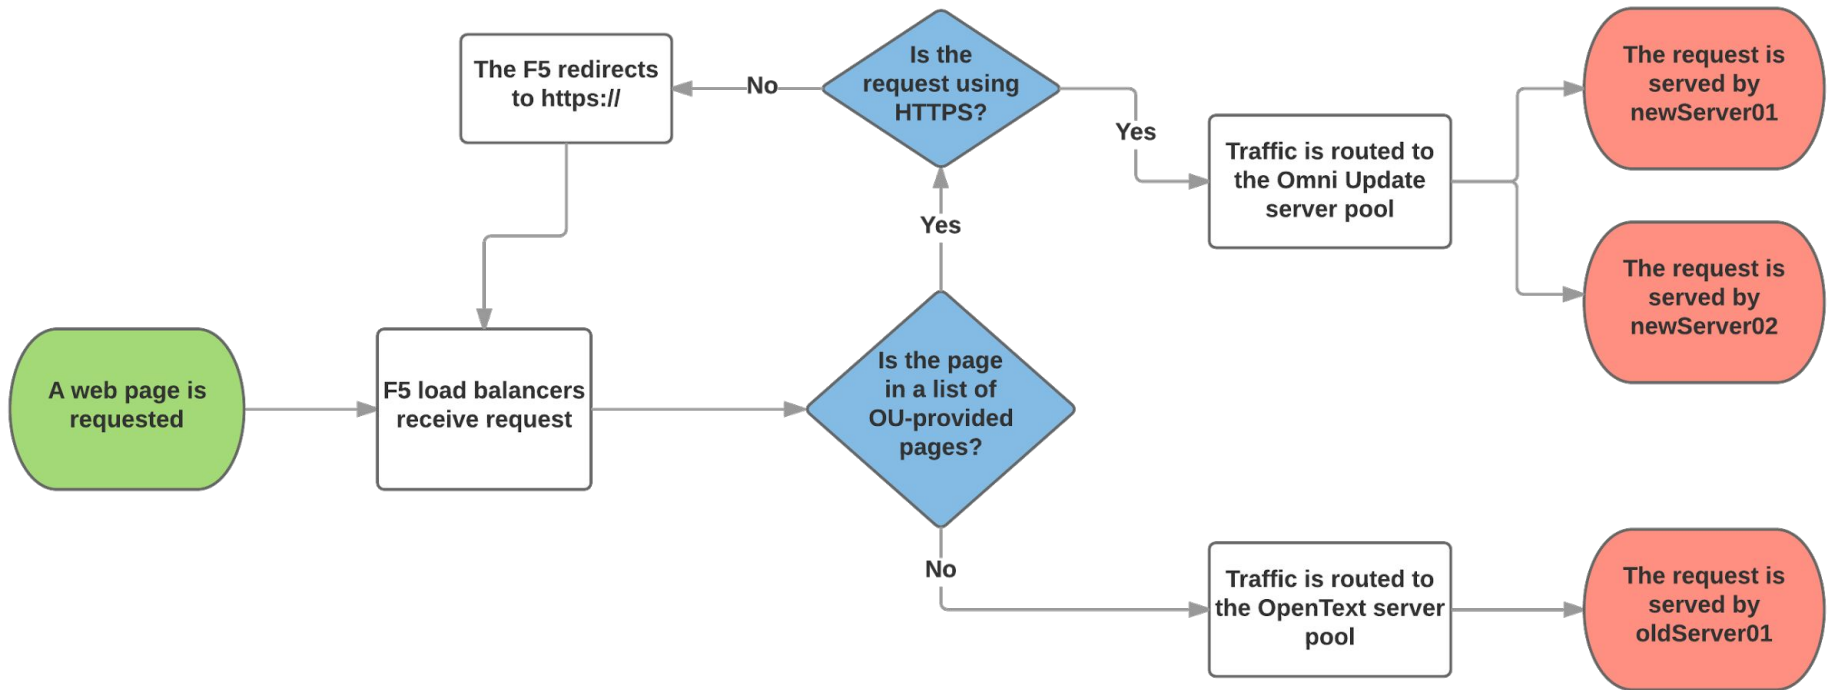
1. Move from OpenText to Omni Update
2. Keep current URL
3. Do the migration in phases
 - a. Phase 1 includes:
 - i. the front page of www.westga.edu
 - ii. ~ top 40 sites by traffic
 - b. Legacy subdomains / custom URL's

Ideas & Opportunities

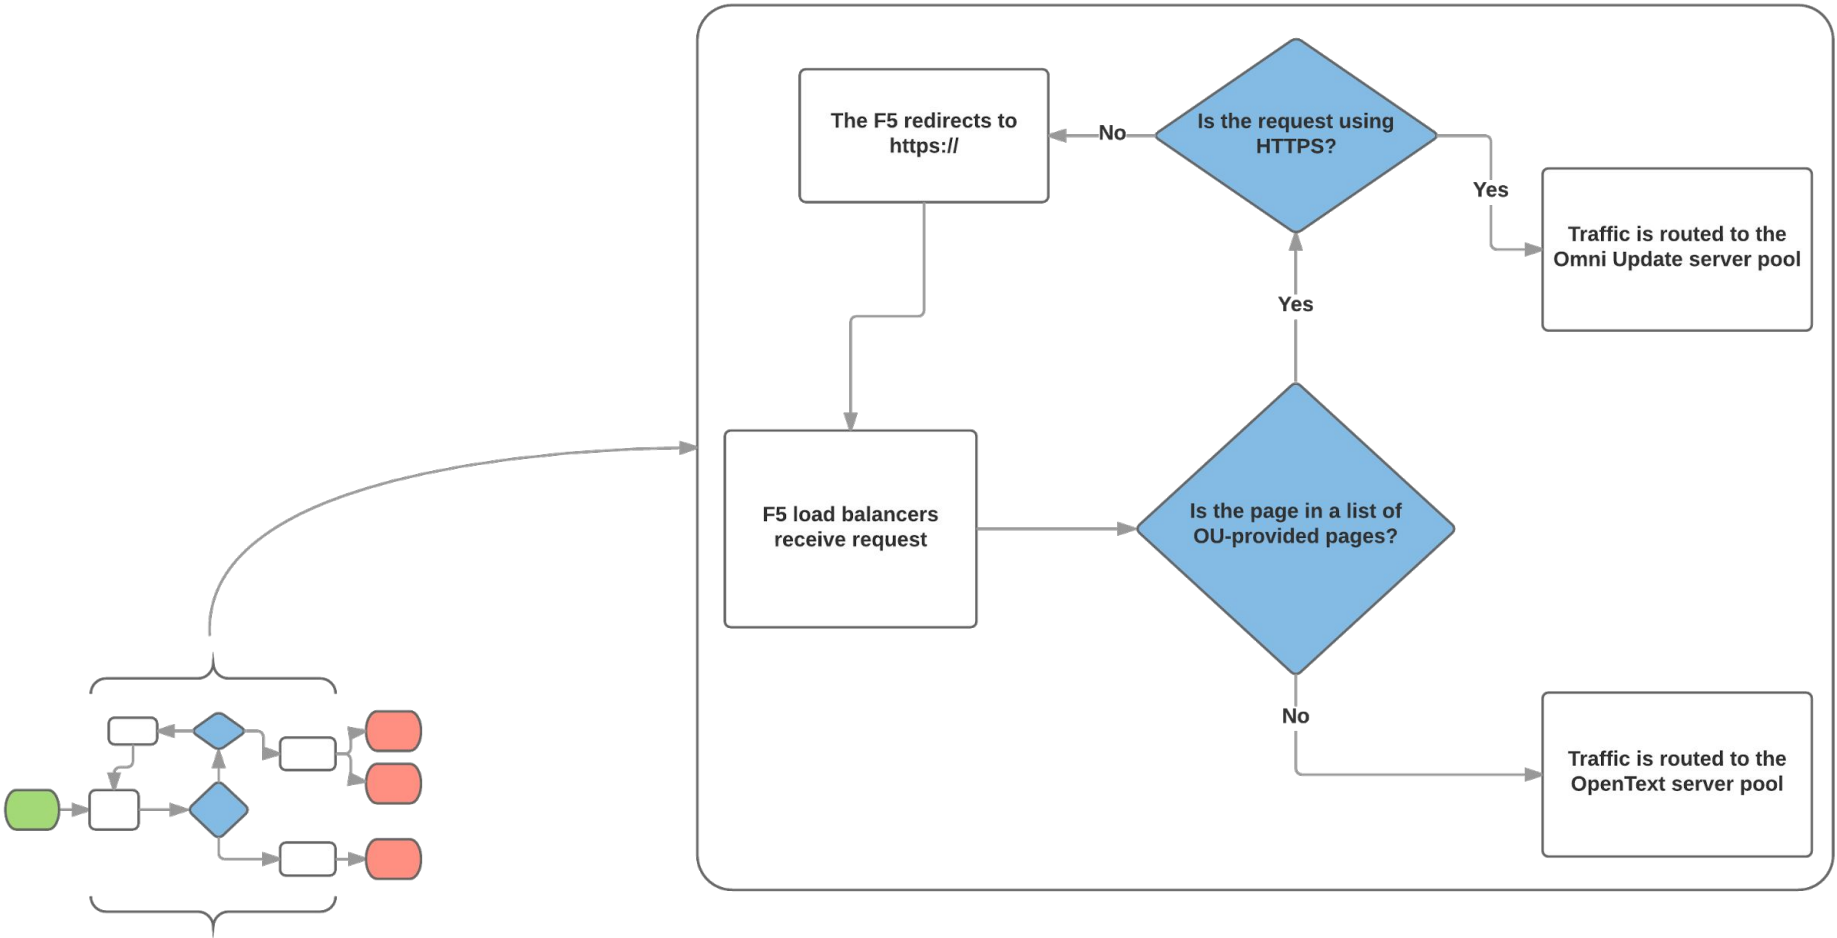
- Use a whitelist based on URI's
 - Site already fronted by F5's -- iRules for the win!
 - *Don't have an F5? Stay tuned... (NGINX to the rescue)*
- Implement high availability / redundancy
- Go 100% HTTPS
 - 2014 - <http://bit.ly/GoogleHTTPSBonus>
 - OpenText site's https version was busted



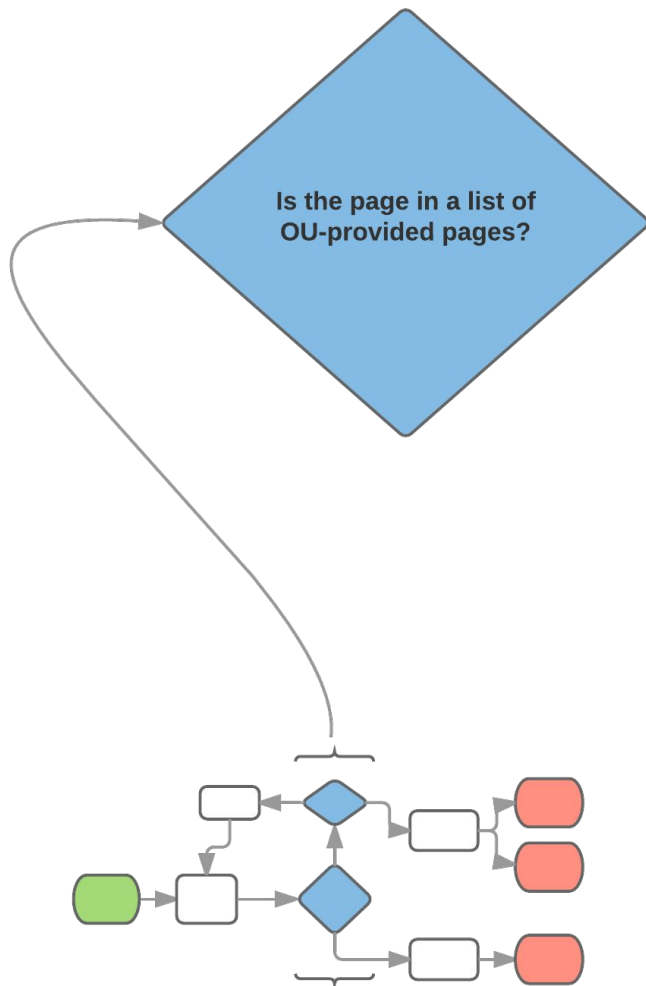
The Design



Route old to old, new to new



Whitelisting







1. Keep a list of migrated URI's
2. Check each request against the list
3. Send migrated requests to new CMS on HTTPS
4. Send everything else to old CMS
5. Facilitate rapid de-migration




Keep a list of migrated URI's


- Prod and Dev version of list
- Tracked in Git
- Edits via GitLab web interface
- Diffs sent to admins
- List deployment is copy / paste






Project in GitLab

Gene Liverman / omni-update-whitelists  This project Search   + 





Project Activity Repository Pipelines Graphs Issues 0 Merge Requests 0 


Files Commits Network Compare Branches Tags














master omni-update-whitelists / + 

Name	Last Commit > 6f5102d1 - 3 days ago DO NOT LAUNCH THIS YET!!! Wait unt...	History	Last Update
 README.md	Update README.md step to include alphabetical placement		4 months ago
 dev-whitelist	The decision and execution of moving academic calendar under academics h...		3 months ago
 f5-irule-uwg_dual_c...	Updated to account for aliases to the domain name		3 months ago
 f5-irule-uwg_dual_c...	Made www2 follow the rules for www.		2 months ago
 prd-whitelist	DO NOT LAUNCH THIS YET!!! Wait until Zach gives word on 10/7/2016		3 days ago



Git Changelog

Gene Liverman / omni-update-whitelists  This project Search   + 

Project Activity **Repository** Pipelines Graphs Issues 0 Merge Requests 0 

-  Zachary Lee authored 4 months ago  **cd7984cc** [Browse Files](#)
-  **Adding some missing entries from the development white list (index.php) and addi...** 
Zachary Lee authored 4 months ago  **28695db5** [Browse Files](#)
-  **Going to a less specific entry for "administration"**
Zachary Lee authored 4 months ago  **88b3d62a** [Browse Files](#)
-  **Added "about-uwg" page to the whitelist**
Zachary Lee authored 4 months ago  **ebbafef11** [Browse Files](#)
-  **Removing entries in favor of just having a catchall for /academics/**
Zachary Lee authored 4 months ago  **57e9b3ed** [Browse Files](#)
-  **Adding "assets-home" to the white list**
Zachary Lee authored 4 months ago  **ebdfaa61** [Browse Files](#)

24 Jun, 2016 4 commits

-  **Initial commit of dual cms prd iRule**
Jake Spain authored 4 months ago  **44a86cef** [Browse Files](#)

Diff Email

Zachary Lee pushed to branch master at [Gene Liverman / omni-update-whitelists](#)

Commits:

- [c57984cc](#)
by Zachary Lee at 2016-06-25T09:26:52-04:00
Adding a ton of root level .php files

1 changed file:

- [prd-whitelist](#)

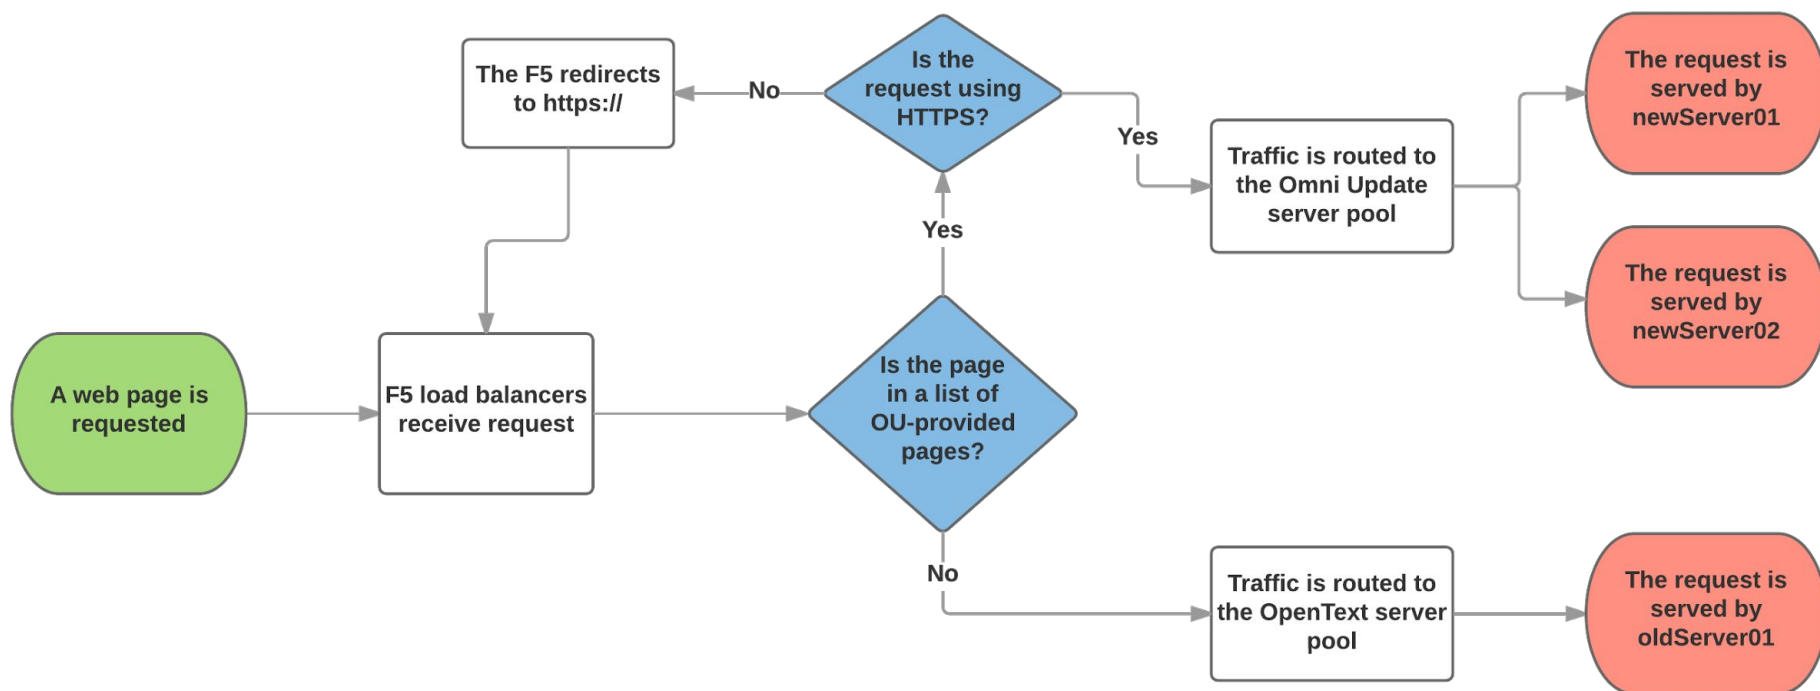
Changes:

- [prd-whitelist](#)

```
...    ... @@ -4,7 +4,9 @@
4      4      "/academics/" := "",
5      5      "/administration/" := "",
6      6      "/admissions/" := "",
7      7      "/apply.php" := "",
7      8      "/assets-home/" := "",
9      9      "/campus-life.php" := "",
8      10     "/campus-life/coliseum/" := "",
9      11     "/campus-life/csi/" := "",
10     12     "/campus-life/dinewest/" := "",
...    ... @@ -14,7 +16,10 @@
14     16     "/campus-life/wolvescard/" := "",
15     17     "/community/engagewest/" := "",
```

Check each request against the list

- Use iRule to process each request
- Use Data Group for the whitelist
- Set backend pool using iRule



Send everything else to old CMS

```
if { ([HTTP::host] equals "www.westga.edu") or  
      ([HTTP::host] equals "www2.westga.edu") } {
```

Requests to www / www2 not in whitelist

1. Determine if secure
2. Set appropriate pool for Open Text backend

```
} else {
```

Requests to some other domain

1. Determine if secure
2. Set appropriate pool for Open Text backend

```
}
```

Directory or File?

- Some content wasn't rendering if the trailing / was missing

```
set uri_parts [split [HTTP::uri] "/" ]
set end_of_uri [lindex $uri_parts end]
if { (not ("end_of_uri" equals "")) and
      (not ("end_of_uri" contains ".")) } {
  set modified_uri "[HTTP::uri]/"
  HTTP::uri $modified_uri
}
```

Facilitate rapid de-migration

Scenario: www.westga.edu/some_department gets migrated and an issue is found that didn't show up in testing.

- Pages on Open Text still exist
- Reverting migration equals deleting one line from a text file
- After undoing migration new pages still exist on Omni Update servers



Alternate Pathways

Port	Usage
80	Processed by iRule
443	Processed by iRule
8080	Sent directly to OpenText server
8443	Sent directly to first Omni Update server
9443	Sent directly to first Omni Update server

- Same on both dev and prod
- Used for testing and verification (esp. on load balanced pair)

Gotchas

- Assets such as CSS -- don't use same path on new server
 - Each URI can only be active on one server
- User-added iFrames + HTTPS conversion
 - Some providers don't offer HTTPS
- Browser caching
- Syncing content in a HA environment

Syncing content in a HA environment

1. Omni Update publishes to first server with SFTP
2. When service account logs out:
 - a. /sync-www.txt is updated with current time
 - b. Rsync is triggered to second server by pam_exec
3. Clean-up rsync is run by cron every 5 minutes

Current sync-status visible in browser:

- 1st server: <https://www.westga.edu:8443/sync-www.txt>
- 2nd server: <https://www.westga.edu:9443/sync-www.txt>



Triggering a sync

Sync by way of pam_exec:

```
/etc/pam.d/sshd:
```

```
# Use pam_exec.so to sync data between web servers  
session optional pam_exec.so /usr/local/bin/pam-sync-www.sh
```

Sync by way of root's crontab:

```
* /12 * * * * /usr/local/bin/sync-www.sh
```

Kicker Script Overview

`pam-sync-www.sh`:

```
( [ "${PAM_TYPE}" == "close_session" ] &&  
  [ "${PAM_USER}" == "${trigger_user}" ] ) || exit 0
```

- Only run script body on our service account's logout
- Exit success for every other scenario



Sync Script Overview

sync-www.sh:

1. On work when run by root
2. Uses lock file to keep from running multiple instances
3. Uses hostname to determine if being run on dev or prod
4. Updates sync-www.txt
5. Runs “rsync -avr --delete-after” of document root

Logging:

- All commands sent to logger
- Rsync output along with start and end times logged to file



Demo Time

www.westga.edu





[Apply](#) [Visit](#) [Give](#) [About UWG](#)



Why will you go West?

[Apply](#)

[Find out now >](#)

[Request more info >](#)

<https://www.westga.edu>

University System of Georgia
Annual Computing Conference @

ROCK EAGLE





[Apply](#) [Visit](#) [Give](#) [About UWG](#)



[Areas of Support](#)

[Services](#)

[Resources](#)

[About Us](#)



Information Technology Services

[Home](#) / [Information Technology Services](#)



<https://www.westga.edu/its>

Home

About Us

Departments

Employee links

Projects

Meet the Staff



GO WEST! GO WOLVES!

UWG Bookstore

Welcome to Campus Planning & Facilities!

Who we are...

<http://www.westga.edu/cpf/index.php>

No F5, no problem.

Got NGINX?

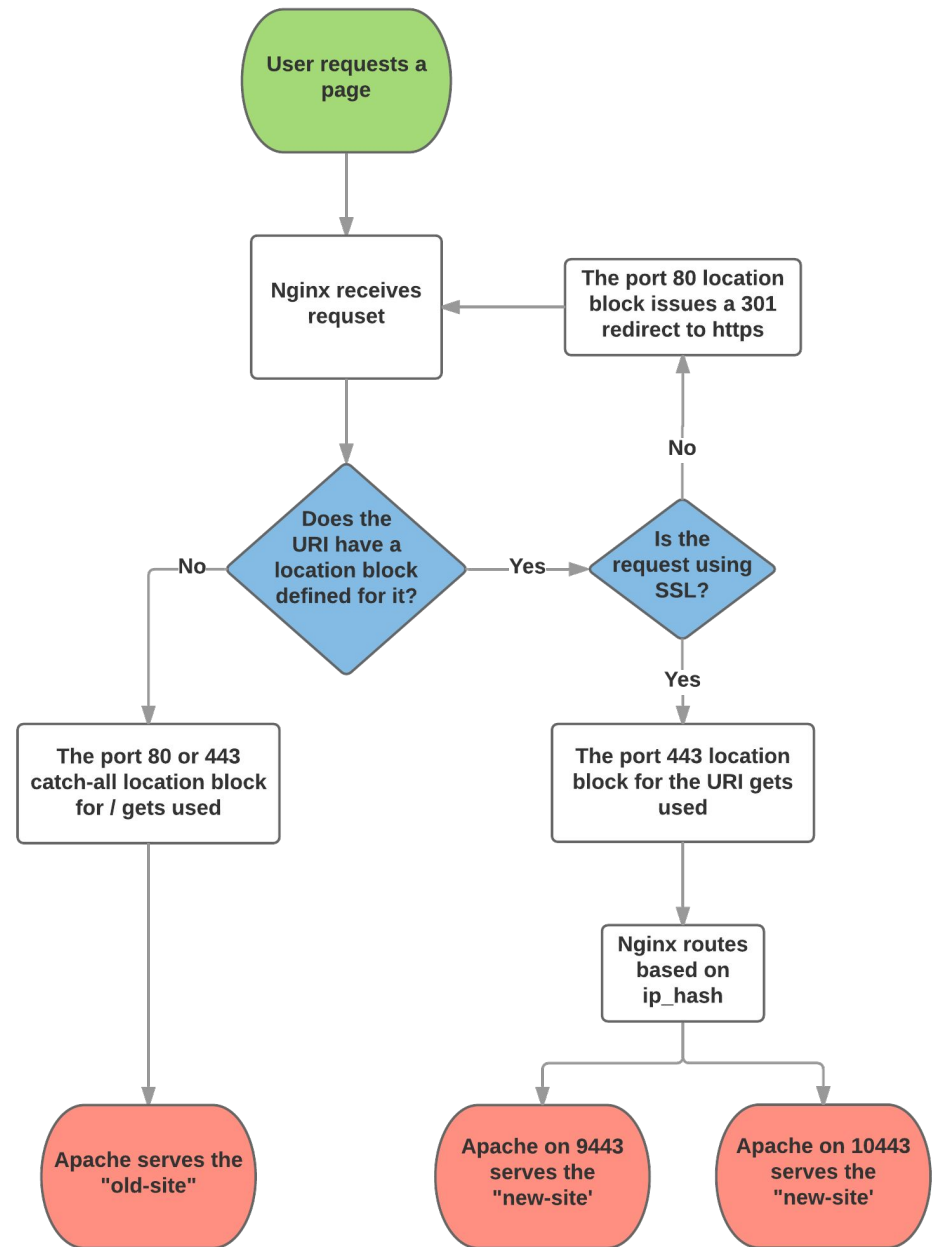


Nginx Setup

- Just a reverse proxy
- Doesn't serve any content
- Working demo on my GitHub

Tools Used

- Vagrant for demo
- One CentOS 7 VM
- Nginx instead of F5
- Apache vhosts for backends



Demo Time

<https://github.com/genebean/nginx-uri-proxy>



New Site in Apache

Current URI: /

Page	Status	Notes
<u>/</u>	Migrated	home page
<u>/index.php</u>	Migrated	Be sure to migrate your actual home page :)
<u>/part1</u>	Migrated	
<u>/part1/random/sub/section</u>		Inherits migrated status
<u>/part2</u>	Old Site	Inherits default location block
<u>/part2/special/page.php</u>	Migrated	Single page explicitly migrated
<u>/part2/special/hello.php</u>		Same directory but not migrated yet
<u>/part3</u>	Migrated	
<u>/not/migrated/yet</u>	Old Site	Also inherits default location block

When migrating content, don't forget about your CSS, JavaScript, and other paths that are included in your pages. I recommend making sure your new site and old site do not use the same path for includes... otherwise you will have problems here.

New Site in Apache

Current URI: /index.php

Page	Status	Notes
<u>/</u>	Migrated	home page
<u>/index.php</u>	Migrated	Be sure to migrate your actual home page :)
<u>/part1</u>	Migrated	
<u>/part1/random/sub/section</u>		Inherits migrated status
<u>/part2</u>	Old Site	Inherits default location block
<u>/part2/special/page.php</u>	Migrated	Single page explicitly migrated
<u>/part2/special/hello.php</u>		Same directory but not migrated yet
<u>/part3</u>	Migrated	
<u>/not/migrated/yet</u>	Old Site	Also inherits default location block

When migrating content, don't forget about your CSS, JavaScript, and other paths that are included in your pages. I recommend making sure your new site and old site do not use the same path for includes... otherwise you will have problems here.

Old Site in Apache

Current URI: /part2

Page	Status	Notes
<u>/</u>	Migrated	home page
<u>/index.php</u>	Migrated	Be sure to migrate your actual home page :)
<u>/part1</u>	Migrated	
<u>/part1/random/sub/section</u>		Inherits migrated status
<u>/part2</u>	Old Site	Inherits default location block
<u>/part2/special/page.php</u>	Migrated	Single page explicitly migrated
<u>/part2/special/hello.php</u>		Same directory but not migrated yet
<u>/part3</u>	Migrated	
<u>/not/migrated/yet</u>	Old Site	Also inherits default location block

When migrating content, don't forget about your CSS, JavaScript, and other paths that are included in your pages. I recommend making sure your new site and old site do not use the same path for includes... otherwise you will have problems here.



New Site in Apache

Current URI: /part3

Page	Status	Notes
<u>/</u>	Migrated	home page
<u>/index.php</u>	Migrated	Be sure to migrate your actual home page :)
<u>/part1</u>	Migrated	
<u>/part1/random/sub/section</u>		Inherits migrated status
<u>/part2</u>	Old Site	Inherits default location block
<u>/part2/special/page.php</u>	Migrated	Single page explicitly migrated
<u>/part2/special/hello.php</u>		Same directory but not migrated yet
<u>/part3</u>	Migrated	
<u>/not/migrated/yet</u>	Old Site	Also inherits default location block

When migrating content, don't forget about your CSS, JavaScript, and other paths that are included in your pages. I recommend making sure your new site and old site do not use the same path for includes... otherwise you will have problems here.



```
include          /etc/nginx/mime.types;
default_type    application/octet-stream;

# Load modular configuration files from the /etc/nginx/conf.d directory.
# See http://nginx.org/en/docs/nginx_core_module.html#include
# for more information.
include /etc/nginx/conf.d/*.conf;

upstream old_backend_http {
    server 127.0.0.1:8081;
}

upstream old_backend_https {
    server 127.0.0.1:8444;
}

upstream new_backend_https {
    ip_hash;
    server 127.0.0.1:9444;
    server 127.0.0.1:10444;
}
```

These define the backend server(s) used

```
location = / {
    # Order defined in Puppet: 001
    return 301 https://$host:443$request_uri;
}

location = /index.php {
    # Order defined in Puppet: 002
    return 301 https://$host:443$request_uri;
}

location /part1 {
    # Order defined in Puppet: 003
    return 301 https://$host:443$request_uri;
}

location = /part2/special/page.php {
    # Order defined in Puppet: 004
    return 301 https://$host:443$request_uri;
}

location /part3 {
    # Order defined in Puppet: 005
    return 301 https://$host:443$request_uri;
}

location / {
    # Order defined in Puppet: 999
    proxy_pass http://old_backend_http;
    proxy_set_header Host $host;
    proxy_set_header X-Real-IP $remote_addr;
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
}
```

These define the matches for http traffic

```
location /part1 {
    # Order defined in Puppet: 003
    proxy_pass https://new_backend_https;
    proxy_set_header Host          $host;
    proxy_set_header X-Real-IP     $remote_addr;
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
}

location = /part2/special/page.php {
    # Order defined in Puppet: 004
    proxy_pass https://new_backend_https;
    proxy_set_header Host          $host;
    proxy_set_header X-Real-IP     $remote_addr;
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
}

location /part3 {
    # Order defined in Puppet: 005
    proxy_pass https://new_backend_https;
    proxy_set_header Host          $host;
    proxy_set_header X-Real-IP     $remote_addr;
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
}

location / {
    # Order defined in Puppet: 999
    proxy_pass https://old_backend_https;
    proxy_set_header Host          $host;
    proxy_set_header X-Real-IP     $remote_addr;
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
}
```

These define the matches for https traffic

Questions?

Dual-CMS Routing

How We Took Our New CMS Live
One URI at a Time

Gene Liverman

gliverma@westga.edu

[@technicalissues](#)

github.com/genebean

linkedin.com/in/geneliverman