INFORMATION TECHNOLOGY SERVICES

CONNECTING CAMPUS

Drupal Architecture at the University of Iowa

08/09/2013 Bill Bacher | Unix Admin | University of Iowa

Agenda

- Overview
- Server Specifics
- Apache Configuration
- PHP Configuration
- MySQL Configuration
- Memcached Configuration
- Questions

Overview

- Typical Architecture:

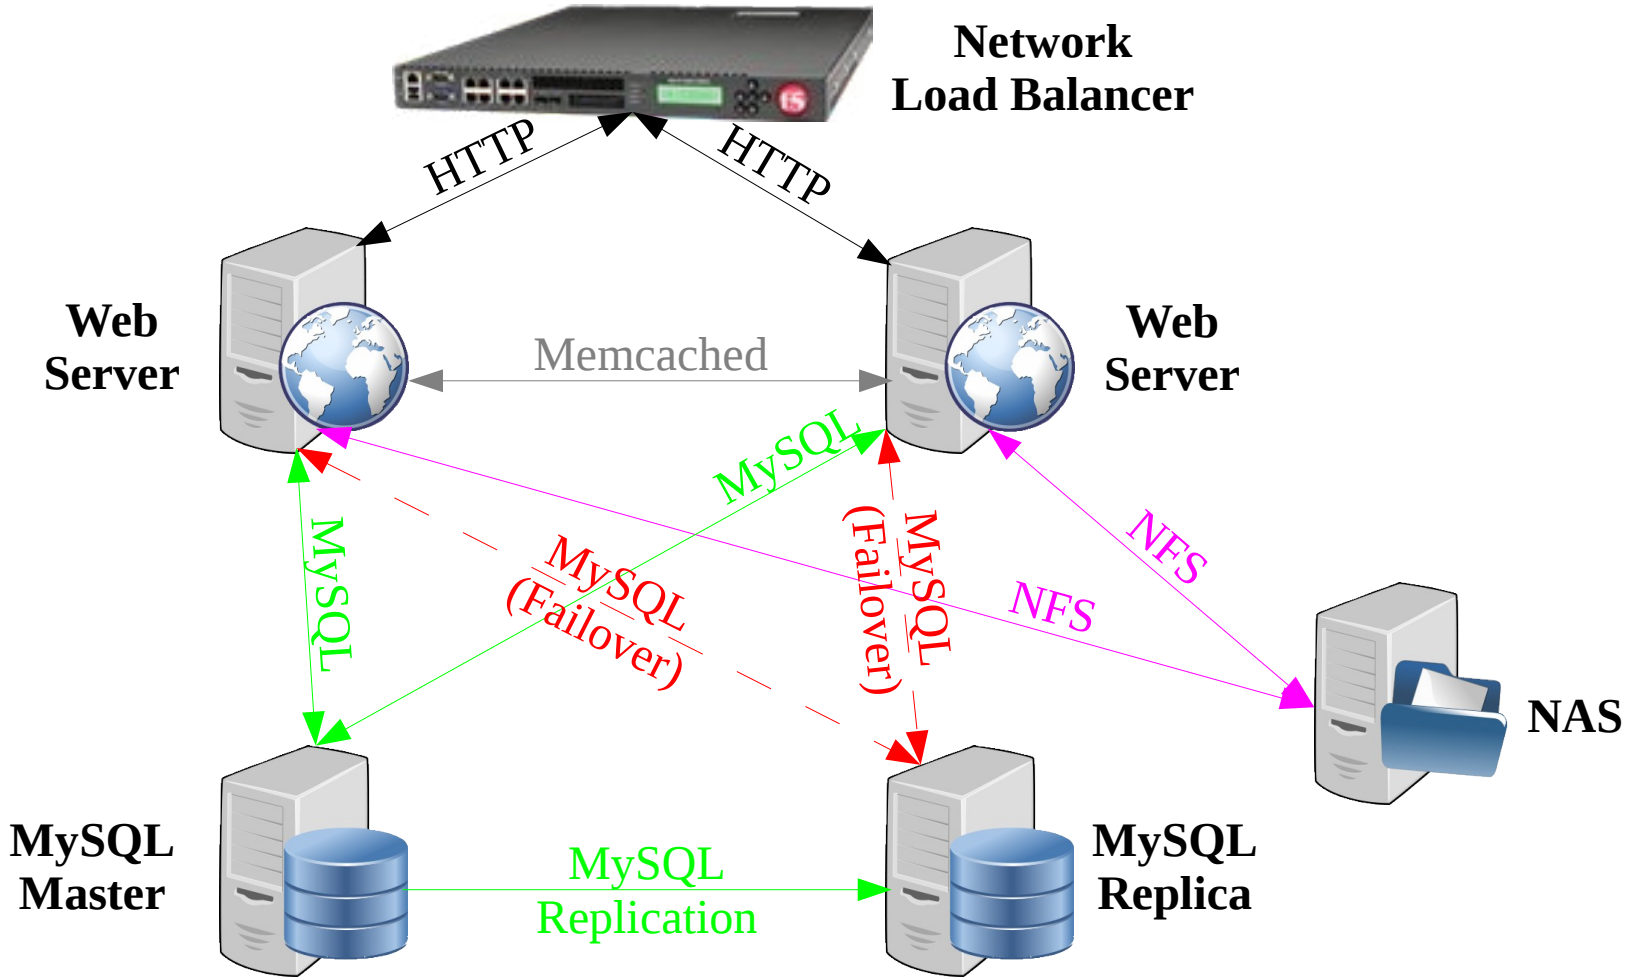

Overview

- 2 year effort
- 90 different sites on our 'Custom' Production Cluster
- Dedicated clusters for special classes of sites
  - org.uiowa.edu
  - course.uiowa.edu
  - studio.uiowa.edu
  - lab.uiowa.edu
  - Starting to move www.uiowa.edu to Drupal

Server Specifics

- All Servers are Virtual Machines
  - VMware ESXi High Availability Clusters
  - Affinity rules so no two servers in one Drupal cluster can reside on the same VMware host
  - Server images stored on 'snap' NAS Volumes
    - Regular snapshots of disk image stored for 'full server' backup
- Web Servers
    - 1 CPU, 8 GB RAM
- Database Servers
    - 1 CPU, 4 GB RAM
- Red Hat Enterprise Linux Version 6.4

Apache Configuration

- Red Hat standard Apache httpd 2.2.15
  - Prefork MPM
  - ModSecurity
    - Multiple custom rules to exclude Drupal functions
      - /admin shows up a lot
  - No SSL
    - Authentication via CAS module and University log-in service
  - Vhost definitions 'hide' Drupal from script kiddies

Apache Configuration

- Storage
  - In SHARED file space (NAS)
    - Apache Configuration
    - ModSecurity Configuration and Rules
    - Archived web logs
    - 'files' directory for each site
  - In LOCAL file space (VM)
    - Active web logs
    - Drupal

Sidetrack - Storage

- 'Local' vs 'Shared' Storage
  - Different?

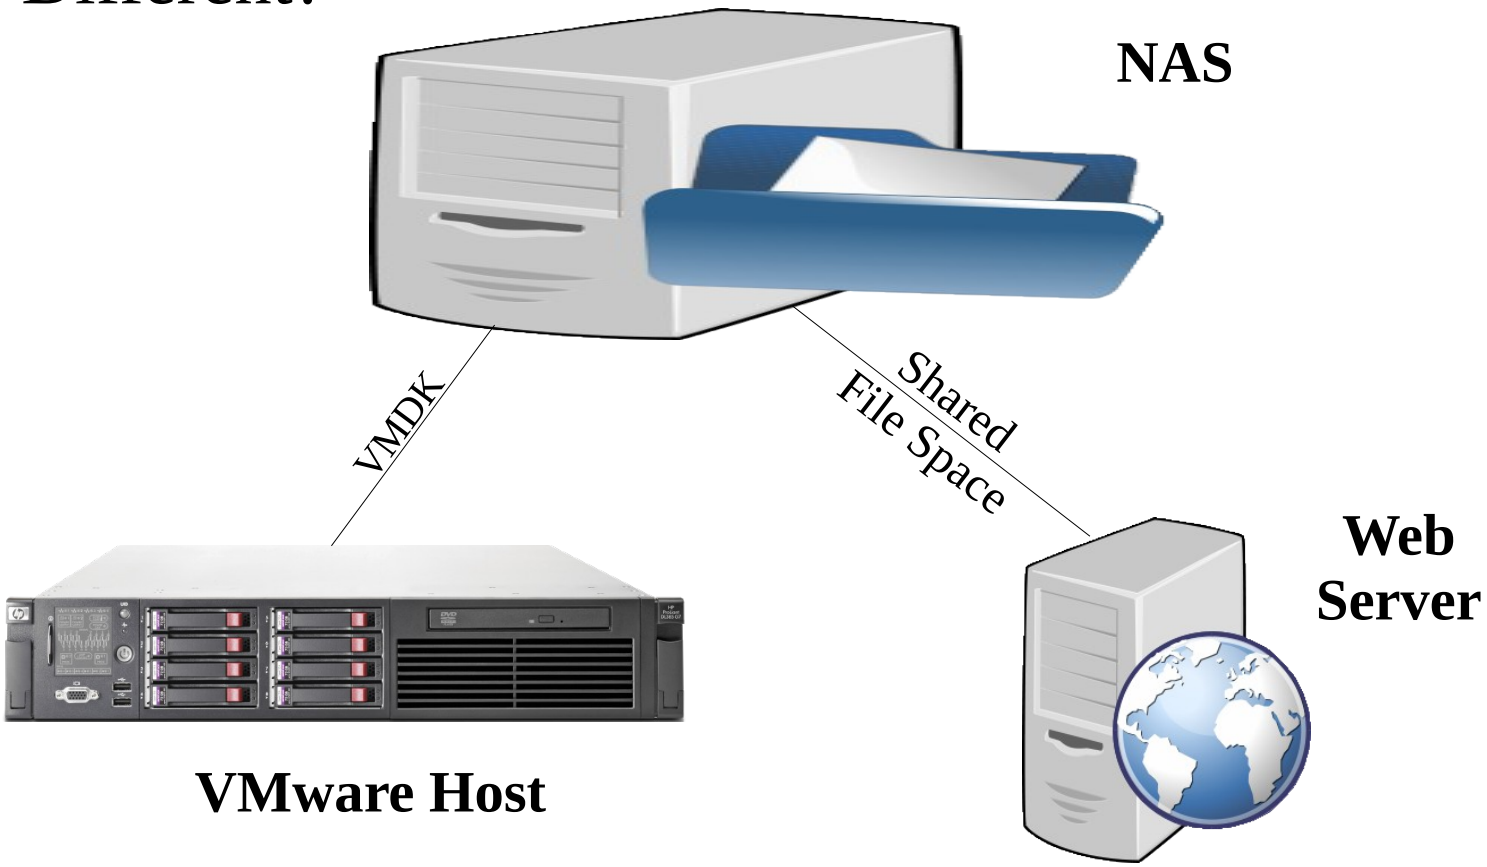

PHP Configuration

- Red Hat Standard PHP 5.3.3
  - Build our own php-mcrypt, php-mssql, php-oci8
  - EPEL for php-domxml-php4-php5, phpMyAdmin, php-pear-CAS, php-pear-DB
- Running Alternative PHP Cache (APC)
  - 256M of dedicated memory
    - 15M hits VS 3,400 misses (typical single day stats, single server)
      - Counts reset nightly with Apache graceful restart during log rotation

PHP Configuration

- PHP Configuration Directives
  - max_execution_time = 30
  - max_input_time = 60
  - memory_limit = 128M
  - post_max_size = 50M
  - upload_max_filesize = 50M
  - 'vhost' configuration typically sets 'AllowOverride' to All
- Note – phpMyAdmin doesn't work well load balanced
- Nothing particularly note worthy

MySQL Configuration

- Running Master/Slave replication
  - Web servers connect to Master server
    - All updates against the Master Database
    - Backups (mysqldumps) done on Replica server
- Running MySQL 5.5.30 from the 'Remi Collet' repository
  - Better performance (per Acquia)
  - Reliable replication (lost sync issues when running 5.1)
  - Red Hat may be providing MySQL 5.5 in RHEL6.5

MySQL Configuation

- InnoDB
  - InnoDB storage engine default when setting up a new Drupal site
  - (Red Hat) default – InnoDB stores all data in 'ibdata1' file
    - ibdata1 file **NEVER** gets smaller
      - Dropping tables or databases does not free up any disk space
        - Space will eventually be reused, if needed
      - Recommended solution is to
        - Dump all Databases
        - Stop MySQL
        - Delete ibdata1 file, ib_logfiles
        - Start MySQL
        - Import database dumps

MySQL Configuration

- InnoDB file/disk issues
  - Simpler Solution
    - Set the “innodb_file_per_table” in my.cnf
      - Creates a *.ibd file for each table
        - Dropping the table automatically deletes the .ibd file, reclaiming the disk space
      - Might take more disk space overall
        - Multiple *.ibd files compared to one ibdata1 file
      - Potentially more files open at any given time
      - No dumping/deleting/importing databases
      - Still end up with a 'smallish' ibdata1 file (50 – 100 MB)

MySQL Configuration

- Master my.cnf
  - ```
    # Replication Stuff
    log-bin=mysql-bin
    server-id=2
    expire_logs_days = 3
    sync_binlog=1

    # Performance stuff
    innodb_buffer_pool_size=3G
    innodb_log_buffer_size=4M
    innodb_flush_log_at_trx_commit=2

    # Manage ibdata1 files
    innodb_file_per_table
    ```

- Replica my.cnf
  - ```
    # Replication Stuff
    server-id=3
    relay-log=mysqld-relay-bin

    # Performance stuff
    innodb_buffer_pool_size=3G
    innodb_log_buffer_size=4M
    innodb_flush_log_at_trx_commit=2

    # Manage ibdata1 files
    innodb_file_per_table
    ```

MySQL Configuration

- Replication logs
  - Master Server logs transactions to mysql-bin.######
    - Replica follows this to keep track of what the master is doing
  - mysql-bin logs are rotated when
    - Size gets to 1GB
    - MySQL is restarted
  - MySQL removes old logs after 3 days (configurable)
- Logs are created based upon activity, managed based upon time

MySQL Configuration

- Between ibdata1 file and transaction logs, problems with disk filling up
  - When running MySQL 'Drupal Cache' not unusual to see 3-6 GB of transaction logs per day
- Script run out of cron every 4 hours
  - Checks replication status
  - Checks log master is writing to, replica is following
  - Uses MySQL command to remove any older logs from master

Memcache

- 2nd part of disk space solution is use memcached rather than MySQL based Drupal Cache
  - No more 'cache' traffic in MySQL transaction logs
  - 3x faster page loads
- Allocate 4 GB per server
  - Using 2.5 GB per Server
- Locked down with iptables
- /etc/sysconfig/memcached:

```
PORT="11211"
USER="memcached"
MAXCONN="1024"
CACHESIZE="4096"
OPTIONS=""
```

Memcache

- settings.php:
  - ```
    $conf['cache_backends'][] =
    'sites/all/modules/contrib/memcache/memcache.inc';
    $conf['cache_default_class'] = 'MemCacheDrupal';
    $conf['memcache_key_prefix'] = 'UID';
    $conf['memcache_servers'] = array(
      '123.456.78.90:11211' => 'default',
      '123.456.78.91:11211' => 'default'
    );
    $conf['memcache_bins'] = array(
      'cache' => 'default',
    );
    ```
  - Need '***UID***' in a multi-site environment
  - memcached servers need to be listed in the same order on all web servers
    ```

Summary

- This is our current state, not necessarily the best setup
- Everything will be reviewed when (if?) we move to RHEL7/Drupal 8
  - Benefits versus Cost of multiple web servers
  - Benefits versus Cost of MySQL Replication

Questions

???

bill-bacher@uiowa.edu