

An overview of the Drupal infrastructure and plans for future growth

prepared by the Drupal Association for the Drupal community

Recommendations

More system administration resources, tuning
12GB Database RAM growth
42GB Web server RAM growth
Virtualization may support more dedicated services

Drupal.org old infrastructure services

Problems:

Web servers not efficiently used

Database Bottleneck

Performance, failures cause downtime, availability problems

Network
Round Robin DNS
Virtual IP alternates requests between D1 & D2

Network uses round robin division of load. Does not check if web server died. Not efficient use of web serving capacity.

Drupal 1
Apache Web Server
PHP 4.4
APC PHP Cache
NFS client files dir
Rsync php drupal.org,.be

Drupal 2
Apache Web Server
PHP 4.4
APC PHP Cache
NFS client files dir
Rsync php drupal.org,.be
groups.d.o, scratch.d.o

Drupal 3
CVS
Mail
PHP 5.2.2

Load is not evenly distributed across servers. Not using maximum RAM capacity of servers.

Drupal DB
MySQL 4.0.x
Network File System (NFS)

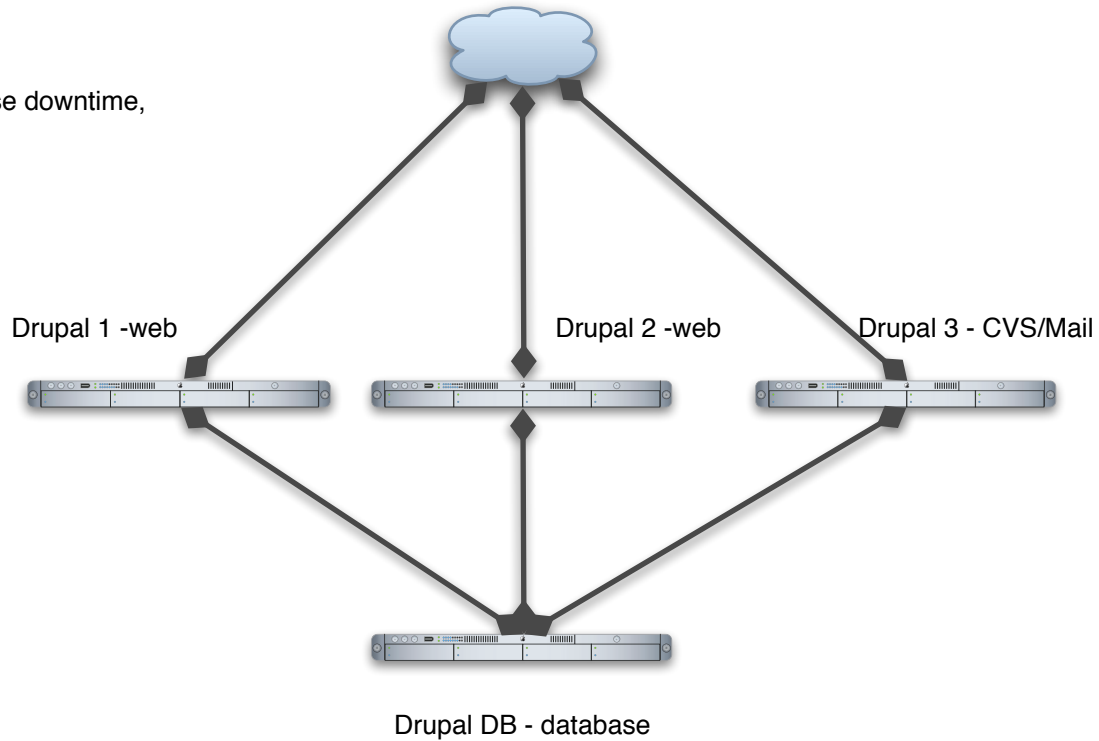
Database is bottleneck in architecture. Need more RAM for MySQL. Could use second db slave for back-ups, high availability failover.

OSUOSL
FTP

Drupal.org Old Infrastructure

Problems:

- Web servers not efficiently used
- RAM upgrades needed
- Database Bottleneck
- Will not scale 250% /year
- Performance, failures cause downtime, availability problems



Firewall uses round robin division of load. Does not check if web server died. Not efficient use of web serving capacity.

Load is not evenly distributed across servers. Not using maximum RAM capacity of servers.

Database is bottleneck in architecture. Need more RAM for MySQL. Could use second db slave for back-ups, high availability failover.

Drupal 1
Dell Poweredge 1850
RAM: 2 x 1 GB DDR-2 400 (6 slots total)
Storage: 2 x 73 GB U230 SCSI drives
(RAID1)
Processor: 2 x Intel Xeon
Motherboard:
Dual power supplies

Drupal 2
Dell Poweredge 1850
RAM: 2 x 1 GB DDR-2 400 (6 slots total)
Storage: 2 x 73 GB U230 SCSI drives
(RAID1)
Processor: 2 x Intel Xeon
Motherboard:
Dual power supplies

Drupal 3
Dell Poweredge 1850
RAM: 2 x 1 GB DDR-2 400 (6 slots total)
Storage: 2 x 73 GB U230 SCSI drives
(RAID1)
Processor: 2 x Intel Xeon
Motherboard:
Dual power supplies

http://www.dell.com/content/products/productdetails.aspx/pedge_1850

Dell's purchased in fundraising drive
<http://drupal.org/node/26707>

Drupal Database
SUN FIRE V20Z
RAM: 4 x 1 GB DDR-2 400 RAM
Storage: 2 x 73 GB U230 SCSI drives (RAID1)
Processor: 2 x AMD Opteron
Motherboard:

Donated by Sun
Microsystems

<http://www.sun.com/servers/entry/v20z/specs.jsp>

Drupal.org new infrastructure services

Improvements:

Load balancers more efficient, HA

Drupal DBs HA, data integrity

Load Balancer 1
IPVS
Weighted Least-Connection
Heart beat

Load Balancer 2
IPVS
Weighted Least-Connection
Heart beat

Load balancers evenly distribute work load and can survive a hardware failure.

Drupal 1
Apache Web Server
PHP 4.4
APC PHP Cache
Squid Cache
Network File Server(NFS)
NFS webroot

Drupal 2
Apache Web Server
PHP 4.4
APC PHP Cache
Squid Cache
NFS webroot

Drupal 3
Apache Web Server
PHP 4.4
APC PHP Cache
Squid Cache
NFS webroot
CVS

Load is distributed evenly. More Drupal web properties can be supported.

Drupal DB1
MySQL 4.0.x
MySQL master

Drupal DB2
MySQL 4.0.x
MySQL slave
MySQL backups

Database replication allows for High Availability, hardware failure, backups, less downtime for maintenance.

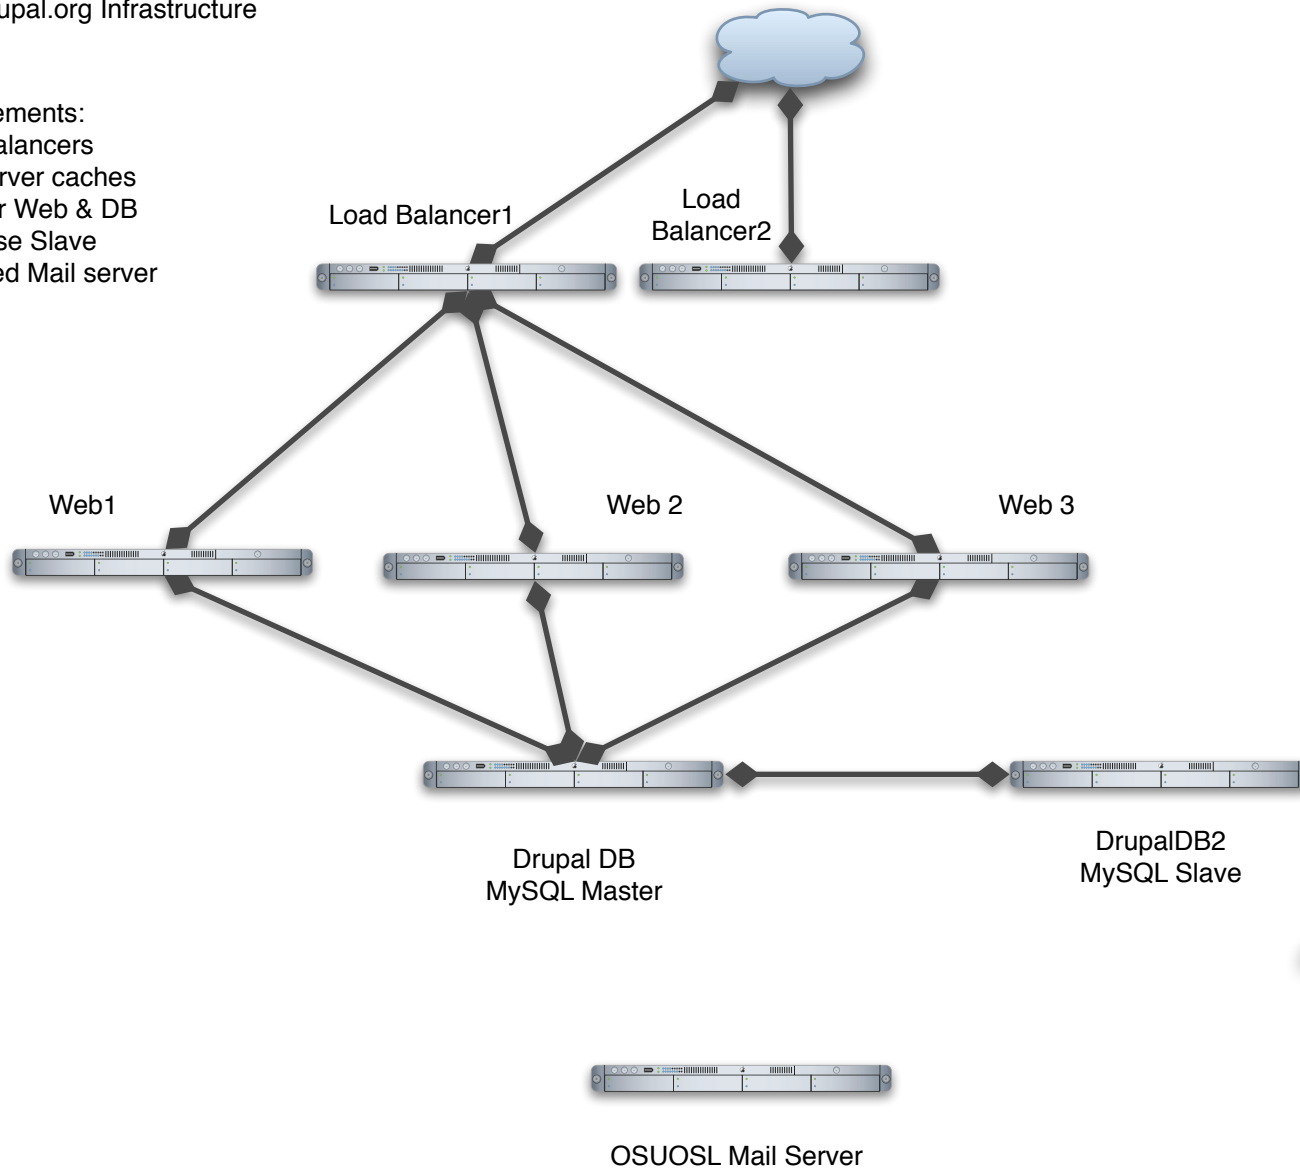
OSUOSL Mail
Mail
Drupal lists

OSUOSL FTP
FTP

Managed mail reduces service needed to be run. Free's hardware for web serving.

Drupal.org Infrastructure

- Improvements:
- Load Balancers
- Web server caches
- RAM for Web & DB
- Database Slave
- Managed Mail server



Load balancers provide higher availability and more efficient use of hardware, hardware growth.

Three web nodes allow for more web page serving capacity, more memory sharing, re-use of old DB server

More powerful DB server allows for more data growth, faster searching. Master-slave allows for higher availability if we have a hardware failure. Read only database could help scale.

Managed mail server reduces volunteer admin overhead, frees up server resources

Load Balancer 1
Proliant DL140
RAM:
Storage:
Processor:
Motherboard:

Load Balancer 2
Proliant DL140
RAM:
Storage:
Processor:
Motherboard:

Load balancers
donated by
OSUOSL

Drupal 1
Dell Poweredge 1850
RAM: 4 x 1 GB DDR-2 400 (6 slots total)
Storage: 2 x 73 GB U230 SCSI drives
(RAID1)
Processor: 2 x Intel Xeon
Motherboard:
Dual power supplies

Drupal 2
Dell Poweredge 1850
RAM: 4 x 1 GB DDR-2 400 (6 slots total)
Storage: 2 x 73 GB U230 SCSI drives
(RAID1)
Processor: 2 x Intel Xeon
Motherboard:
Dual power supplies

Drupal 3
Dell Poweredge 1850
RAM: 4 x 1 GB DDR-2 400 (6 slots total)
Storage: 2 x 73 GB U230 SCSI drives
(RAID1)
Processor: 2 x Intel Xeon
Motherboard:
Dual power supplies

Dell's purchased in fundraising drive
<http://drupal.org/node/26707>

http://www.dell.com/content/products/productdetails.aspx/pedge_1850

Drupal Database Master
SUN FIRE V20Z
RAM: 8 GB DDR-2 400 RAM
Storage: 2 x 73 GB U230 SCSI drives (RAID1)
Processor: 2 x AMD Opteron
Motherboard:

Donated by Sun
Microsystems

Drupal Database Slave
Proliant DL145 G2
RAM:
Storage:
Processor:
Motherboard:

Loaned by
OSUOSL

<http://www.sun.com/servers/entry/v20z/specs.jsp>

Mail services to be
donated by OSUOSL

Drupal.org new infrastructure services

Improvements:

Load balancers more efficient, HA
Drupal DBs HA, data integrity

Load Balancer 1
IPVS
Weighted Least-Connection
Heart beat

Load Balancer 2
IPVS
Weighted Least-Connection
Heart beat
CVS

Load balancers evenly distribute work load and can survive a hardware failure.

Drupal 1
Apache Web Server
PHP 4.4
APC PHP Cache
Squid Cache
NFS server
NFS web root
MemCache Pool*

Drupal 2
Apache Web Server
PHP 4.4
APC PHP Cache
Squid Cache
NFS web root
MemCache Pool*

Drupal 3
Apache Web Server
PHP 4.4
APC PHP Cache
Squid Cache
NFS web root
MemCache Pool*

Drupal N
Apache Web Server
PHP 4.4
APC PHP Cache
Squid Cache
NFS client web root
MemCache Pool*

Load is distributed evenly. More Drupal web properties can be supported.

Drupal DB1
MySQL 4.0.x
MySQL master

Drupal DB2
MySQL 4.0.x
MySQL slave
MySQL backups

Drupal DB3
MySQL 4.0.x
MySQL slave
MySQL read only*

Drupal DBN
MySQL 4.0.x
MySQL slave
MySQL read only*

Database replication allows for High Availability, hardware failure, backups, less downtime for maintenance.

Drupal Development
Virtual DLAMP instances
Unit Testing
Performance Testing
Security level services
Logging/Analytics

Drupal CVS
NFS web root
Drupal Distributions
Drupal Tools

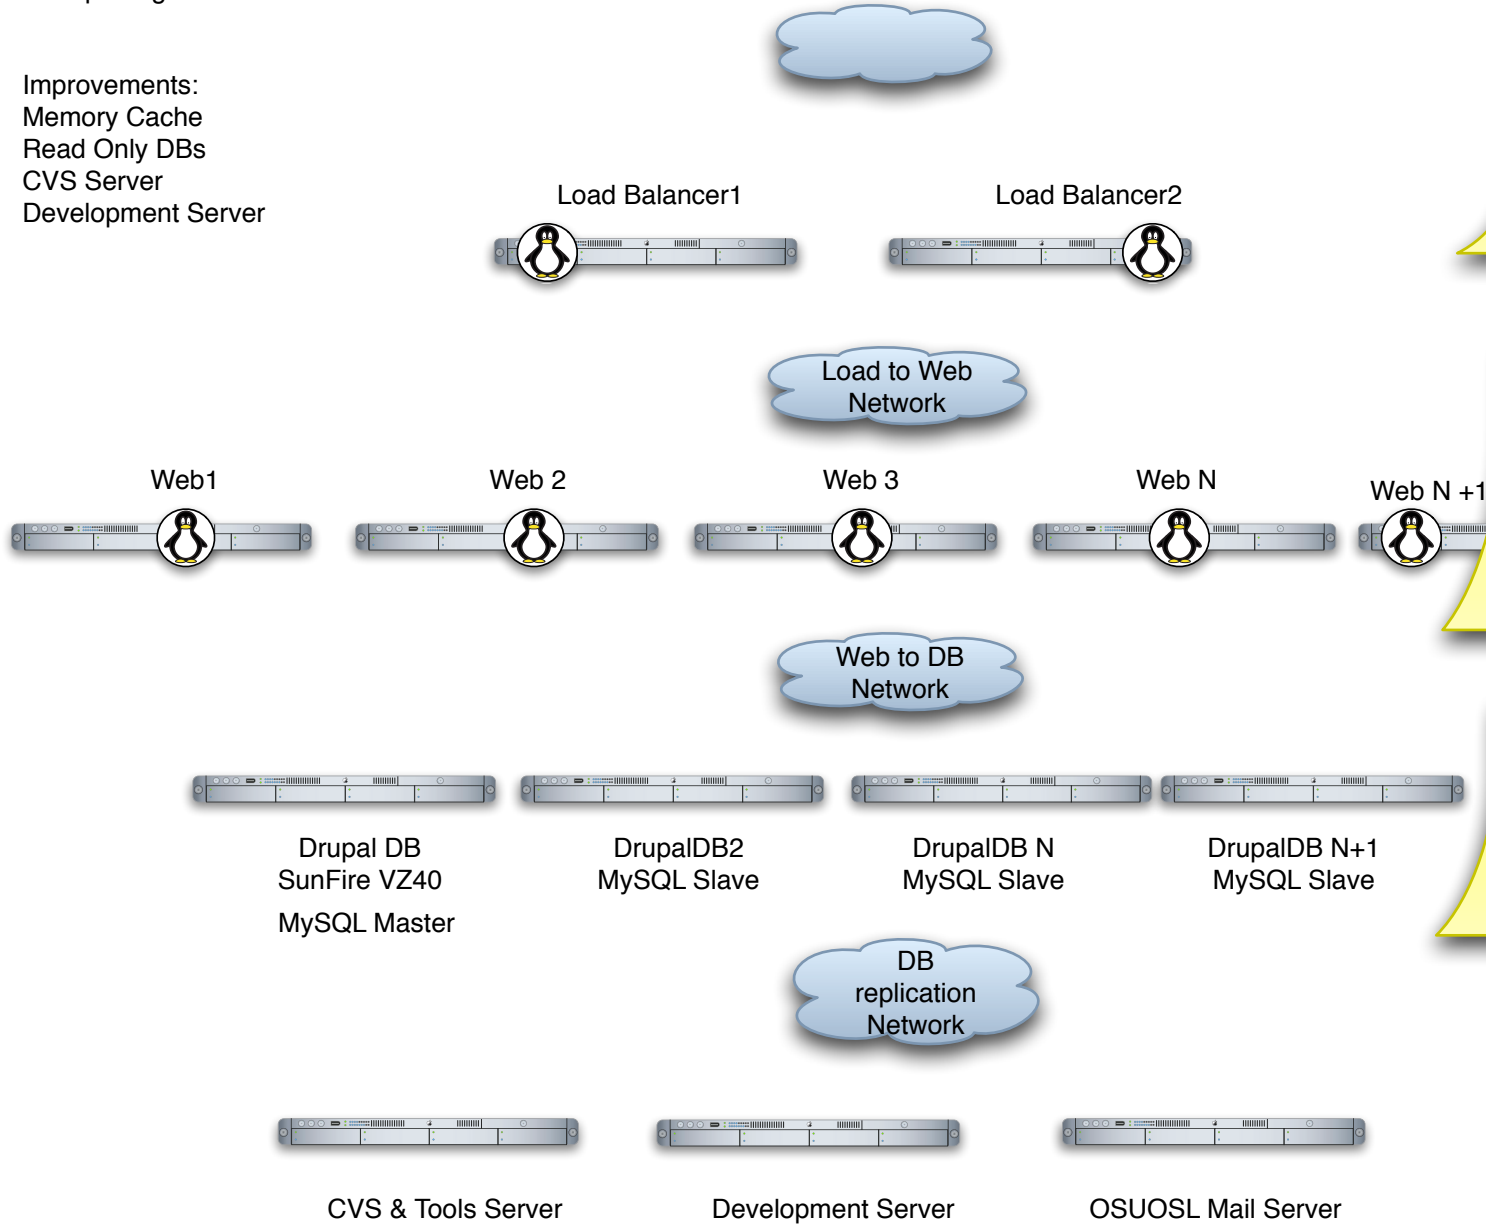
OSUOSL Mail
Mail
Drupal lists

Dedicated servers for unit, developer, security, performance testing. CVS server free's D3 for web serving.

*In order to scale databases horizontally with read only nodes, authenticated sessions must be moved to MemCache

Drupal.org future infrastructure

- Improvements:
- Memory Cache
- Read Only DBs
- CVS Server
- Development Server



High Availability and efficient web serving

Load balancers allow horizontal web server scaling to alive web servers. Cache may reduced load on web servers serving pages from memory. Memcache allows memory pooling so more content can be served from memory.

Most web requests just want information and don't need to update* or delete information. By directing read only requests to read only databases we can scale databases horizontally.

* Authenticated users require a session update which could be moved to memory with memcache

Web Servers
Support up to 16GB RAM each

Additional Web Servers

Drupal Association
fundraising is necessary
and coordination with
OSUOSL

Database master
Support up to 16GB RAM

MySQL database slave
Additional slaves servers

Dedicated CVS Server

Dedicated Development Server

Software development,
software solutions, and
virtualization services
may make future
hardware growth
unnecessary