

2000



- Sydney hosts the Olympic games
- Concord crashes in Paris
- Y2K problem
- UCS launches the Managed Web Server

MWS 1: Solaris 7, Apache 1.3, chroots & loopbacks, static content only, single server

2008



Beijing hosts the Olympic games
Almost nothing else happens
Except UCS launches revised/improved MWS

MWS 2: Solaris, Apache 2, PHP, MySQL, Solaris Zones and ZFS, manual dual-site resilience
~230 sites, ~320 hostnames

WHY A MANAGED WEB SERVICE?

- Eliminate badly managed web servers
- Reduce duplicated staff effort
- VMs <—> MWS <—> [Falcon, Drupal]

NEW MWS DESIGN PRINCIPLES

- Preserve most of the functionality of the current MWS
- Provide an environment similar to that provided by commercial web hosting providers
- Reduce the amount of Unix knowledge required to successfully manage a MWS site
- Delegate as much configuration control as possible directly to MWS site administrators

"MANAGED"

We provide

- Hardware
 - Dedicated virtual server (1 GB RAM, 20GB SSD disk (10GB usable))
- Physical hosting
 - Automatic geographically-diverse redundancy
 - WDC and RNB
- Operating system (Debian Jessie)
- Software
 - Apache 2.4
 - MySQL
 - PHP
- Backups
- YOU provide and MANAGE everything else

SO, WHAT DO YOU GET?

Mention some or all of these

- Delegated management - Web-based control panel
 - All users need UIS accounts
 - New sites on request - **demo**
- Public IPv4 and IPv6 addresses, unlimited name-based vHosts
- ssh-based shell access, scp, sFTP, etc.
 - Independent, non-shared accounts
 - No plain Telnet or FTP
 - ssh public key authn, SSHFP records
- Apache, PHP, Python, MySQL, file system permissions, Cron jobs, email
- IP-register integration
- UIS Passwords
- Snapshots and backups; rollback
- https support

https://panel.mws3.csx.cam.ac.uk/

UNIVERSITY OF CAMBRIDGE
Managed Web Service server: mws3test Welcome Jon Warlick (381) (2x30)

MWS server "mws3test" control panel

Main Settings Server Settings Web Sites Authorised users My other MWS servers Your server is ready

- Server settings
- Account settings
- Manage authorised users and groups
- System analytics
- Billing settings

Server info

IPv4: 221.111.08.10	IPv6: 2001:430:312:B::A6:10
hostname: mws-40760.mws3.csx.cam.ac.uk	defaul website: mws3test.cam.ac.uk

List of SSH host key fingerprints

192A: 42:44:44:24:8F:46:6B:05:05:10:15:0F:54:79:77:82	05A: 42:47:41:42:84:8B:05:57:34:90:26:42:82:42:47:79
EC25A: 5:12:7B:13:2B:10:15:17:6F:84:2C:84:64:62:84:24:19	EC2519: 29:94:04:79:42:9F:48:84:71:11:4d:5c:2:1:07:68:68

The Managed Web Service is provided by the University Information Services.

UNIVERSITY OF CAMBRIDGE © 2014 University of Cambridge University A-Z Contact the University Accessibility Division of Information Terms and conditions	Study at Cambridge Undergraduate Graduate Continuing education Executive and professional education Courses by education	About the University How the University and Colleges work Visiting the University Maps News Jobs Giving to Cambridge Global Cambridge	Research at Cambridge News Features Discussion Funding on... About research at Cambridge
--	--	---	--

```
jw35 — jw35@mws-priv-7: /var/www/default — ssh — 80x24
jw35@mmenth:~$ ssh mws-40760.mws3.csx.cam.ac.uk
Enter passphrase for key '/home/jw35/.ssh/id_rsa':

* This VM is part of the UIS Managed Web Service.
  MWS server: mws3test
  VM role: Production
  VM host name: mws-priv-7.mws3.csx.cam.ac.uk
  Service name: mws-40760.mws3.csx.cam.ac.uk

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Fri May 20 14:33:31 2016 from 2001:630:212:110:488d:5bac:c727:8f60
jw35@mws-priv-7:~$ cd /var/www/default
jw35@mws-priv-7:/var/www/default$ ls
README.MWS3  adminidir  cgi-bin  docroot  log  tls  tls.key
jw35@mws-priv-7:/var/www/default$
```



Extensive documentation

<http://mws-help.uis.cam.ac.uk/>

The screenshot shows the 'Managed Web Service Help & Support' website. The header includes the University of Cambridge logo and navigation links for 'Study at Cambridge', 'About the University', and 'Research at Cambridge'. The main content area is titled 'Managed Web Service Help & Support' and includes a navigation menu with options like 'Home', 'Introduction', 'Using provided software', 'Using third-party software', and 'Billing and administration'. The 'Introduction' section is currently selected and displays text about the Managed Web Service (MWS) and a list of links for further information.

WORKS WITH

- Static sites
- PHP sites
- Python sites
- Wordpress
- Drupal
- ...and many more
- ...but you have to look after them all

Probably works with...

- No root/sudo access
- £100/year

- Automatic payment management
- Requires purchase order
- Runs from the day you create the server
- Initial month to sort out payment

TRANSITION

- Old sites need to go somewhere
- Old hardware
 - Difficult to manage
 - Getting too old for comfort
 - Causing operational problems
 - Can't deploy IPv6 multicast on data centre network
- Payment discount
- Transition advice on the web site

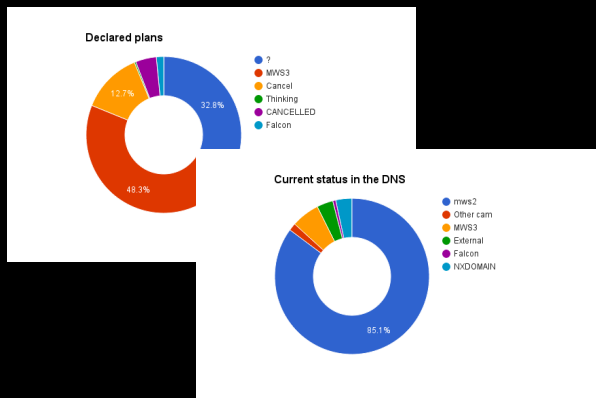
TALK TO US BY THE END
OF JUNE

We've been trying to talk to current users, but not everyone seems to want to talk to us.

If you don't, your sites will be switched off first week July.

IF YOU DO, YOU HAVE TO THE END
OF SEPTEMBER TO TRANSITION

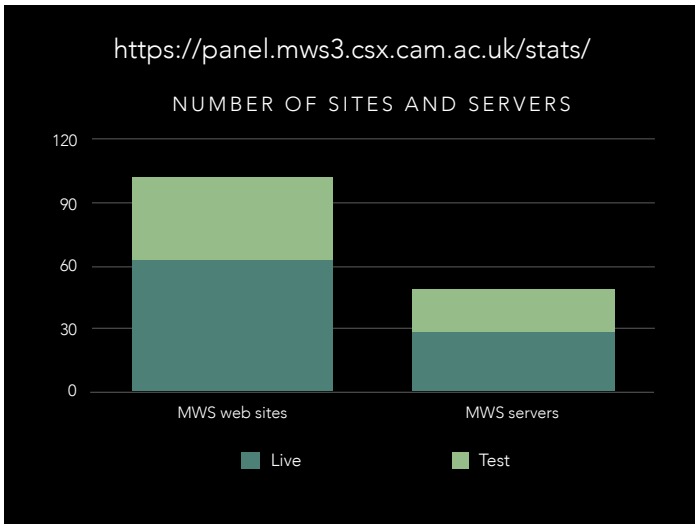
<https://mnementh.csi.cam.ac.uk/mws3-progress.html>



- Mismatch between plans and progress

These are the sites that we don't yet know about

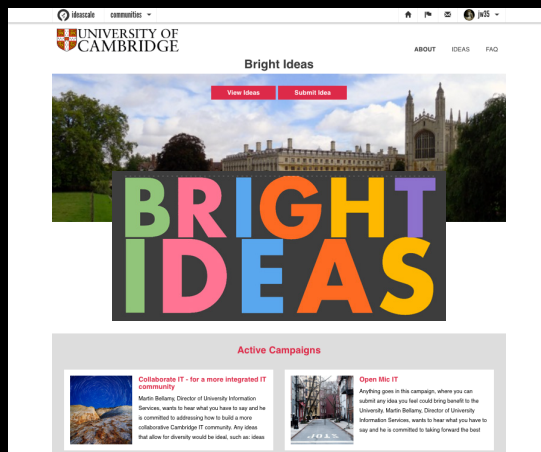
intranet.ames.cam.ac.uk
www.unite.cam.ac.uk
www.archanth.cam.ac.uk
www.astronomicalimages.group.cam.ac.uk
www.biotech.cam.ac.uk
www.biot.cam.ac.uk
www.mrcisi.cam.ac.uk
www.cambridgecancercentre.org.uk
www.cancer.cam.ac.uk
www.cag.group.cam.ac.uk
www.cifr.cam.ac.uk
mvs.cimr.cam.ac.uk
financehub-old.csl.cam.ac.uk
new-www.clarehall.cam.ac.uk
www.dev.clarehall.cam.ac.uk
www.test.clarehall.cam.ac.uk
www.demo.clarehall.cam.ac.uk
www.cis.cam.ac.uk
www.cmi.cam.ac.uk
cms.homerton.cam.ac.uk
www.cnig.group.cam.ac.uk
www.comodo.group.cam.ac.uk
www.conservation.cam.ac.uk
www.physics.corpus.cam.ac.uk
www.cria.polis.cam.ac.uk
www.cria.arch.cam.ac.uk
www.ipsc.cruk.cam.ac.uk
www.lightmicroscopy.cruk.cam.ac.uk
www.ctc.cam.ac.uk
www.qi.damtp.cam.ac.uk
www.dtg.group.cam.ac.uk
www.ct2014.dpmms.cam.ac.uk
www.econsoc.hist.cam.ac.uk
www.edeg.org
www.edwinstowe.group.cam.ac.uk
www.eeci.cam.ac.uk
development.eeci.cam.ac.uk
cit.eng.cam.ac.uk
legisjg.mrc.epid.cam.ac.uk
www.gates.scholarships.cam.ac.uk
www.genet.ac.uk
henge.bioanth.cam.ac.uk
christianization.hist.cam.ac.uk
www.econpublic.hps.cam.ac.uk
www.jobs.hps.cam.ac.uk
www.lara.psychol.cam.ac.uk
lipkow.sybiol.cam.ac.uk
www.many-core.group.cam.ac.uk
www.genetics.mcdonald.cam.ac.uk
new-www.mediaplayer.group.cam.ac.uk
mega.bioanth.cam.ac.uk
www.barlett.psychol.cam.ac.uk
www.ims.cam.ac.uk
www.aim.msm.cam.ac.uk
www.t1.aim.msm.cam.ac.uk
www.mv.crash.cam.ac.uk
mvt.psychol.cam.ac.uk
mvs.psychol.cam.ac.uk
old-www.mz2015.psychol.cam.ac.uk
www.camcompog.psychol.cam.ac.uk
www.mzs.cvx.cam.ac.uk
www.northasianborders.net
www.p-ace.group.cam.ac.uk
www.parenchyma.archanth.cam.ac.uk
www.pave.arch.cam.ac.uk
new-noggin.pdn.cam.ac.uk
www.peerupport.cam.ac.uk
www.phagedisplay.org
www.gsfc.phy.cam.ac.uk
www.qm.phy.cam.ac.uk
www.sd.phy.cam.ac.uk
mws.polis.cam.ac.uk
old-www.cirmena.polis.cam.ac.uk
www.prime.bioanth.cam.ac.uk
www.purpleish.phy.cam.ac.uk
www.pwf.csi.cam.ac.uk
www.rock-art.group.cam.ac.uk
manual.recoup.educ.cam.ac.uk
www.manual.recoup.educ.cam.ac.uk
www.wyse-applications.socanth.cam.ac.uk
old-www.socanth.cam.ac.uk
www.scopic.ac.uk
www.shf.group.cam.ac.uk
www.sjhs.mml.cam.ac.uk
mws.sociology.cam.ac.uk
www2.synbio.cam.ac.uk
ad-consortium.synbiol.cam.ac.uk
logic.synbiol.cam.ac.uk
www-mws.synbiol.cam.ac.uk
ralser.synbiol.cam.ac.uk
trailblazing.group.cam.ac.uk
www.trailblazing.group.cam.ac.uk
www.techdem.crash.cam.ac.uk
transgenerational.zoo.cam.ac.uk
www.viscog.psychol.cam.ac.uk
www.westcam.safety.group.cam.ac.uk
www2.zoo.cam.ac.uk



THE FUTURE

- Current work coming to an end
- Backlog of 40-50 items that won't get done this time, for example
 - PostgreSQL
 - At-cost upgrades
 - Remote access to databases
 - GUI database admin tools pre-installed
 - Support for LetsEncrypt
 - Nginx/Tomcat/...
 - Varnish cache


<http://cambridgeuniversity.ideascaleapp.eu/>



The screenshot shows the 'Bright Ideas' campaign page for the University of Cambridge. The page features a header with the university logo and navigation links for 'ABOUT', 'IDEAS', and 'FAQ'. Below the header is a large banner image of a Cambridge building with the text 'BRIGHT IDEAS' overlaid in large, colorful, block letters. Two red buttons labeled 'View Ideas' and 'Submit Idea' are positioned above the banner. Below the banner, there is a section titled 'Active Campaigns' with two sub-sections: 'Collaborate IT - for a more integrated IT community' and 'Open Mic IT'. Each sub-section includes a small image and a brief description of the campaign.

Plan is to put them into a Bright Ideas campaign
...and ask you all for other ideas
...and ask you all to vote

ARCHITECTURE



Pink: External services

We manage all of it with ansible

Application Deployment + Configuration Management + Continuous Delivery

Dynamic Inventory

Raven Webauth

Lookup

AUTHORISATION (LDAP_{SSH} BASED)

Authorised users and groups

You can authorise other users as administrators or ssh-only users. Administrators will have access to all the features of the control panel and can connect to the server via SSH. SSH-only users will only have access to the server via SSH. They won't be able to access the web panel.

You can search for other users using the text field below by typing their name or CRSid. It will let you autocomplete by selecting their entry from the drop down list.

You can also authorise users or administrators using [Lookup groups](#).

Administrators:

SSH-only users:

Administrators lookup groups:

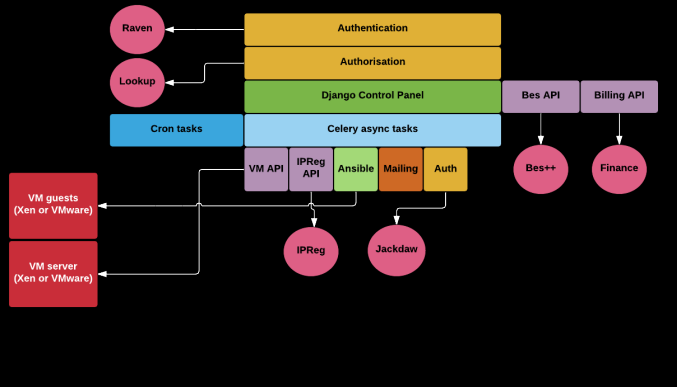
SSH-only lookup groups:

Submit >

The list of users authorised using Lookup groups are refreshed every 24 hours. If you want to refresh it now, you can use the following button.

Force update >

ARCHITECTURE



Jackdaw for UID of user

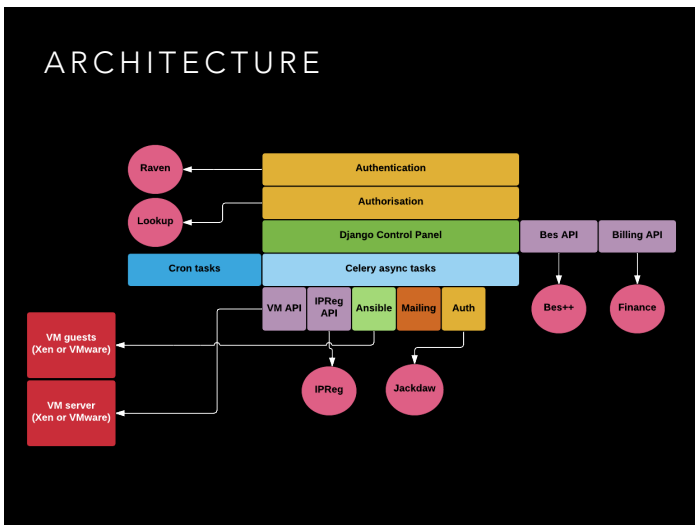
User is installed in the VM (Using Ansible)

UID (important for shared file storage) taken from Jackdaw (User central database)

Periodic task to refresh installed users (in VMs) authorised via LDAP groups

SSH public key uploaded to the web panel

ARCHITECTURE



IP Register

Preallocated IP addresses

*.cam.ac.uk hostnames available for users (API)


Service/Host addresses

SSHFP records and DNSSEC

Preallocated: we can have instant access instead of waiting for DNS servers updates.

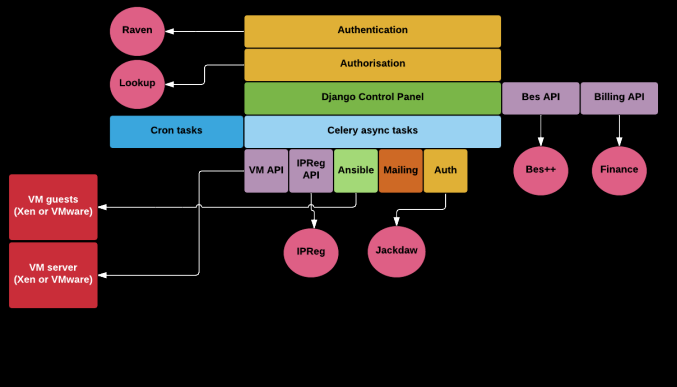
IP4&6 addresses and hostnames

ARCHITECTURE



Bes++ (central inventory)
JSON file with information about all hosts:
Location, IP, hostname, VM properties
Used for other purposes like monitoring

ARCHITECTURE



Celery + Redis
Some API calls
Background processes
Cron jobs (update machine with last changes in lookup / system packages / backups and snapshots)
Message queue storage
Mutex
Different VM APIs (compatible)



The service is designed to be scalable and has a small cluster of AMQP message brokers (which receive metrics) and a small cluster of carbon/graphite machines to store and graph data..


statsd & collectd

cluster AMQP message brokers

cluster carbon/graphite (storage)

PRIVILEGE SEPARATION

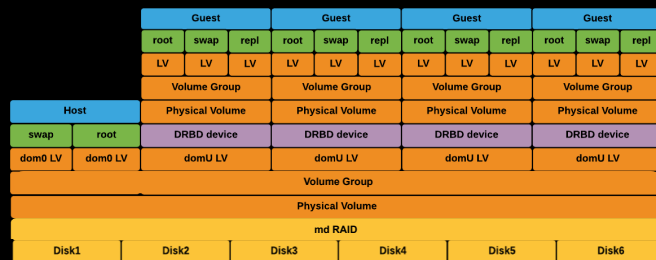
< Hexapodia as the key insight >



CLIENT STORAGE

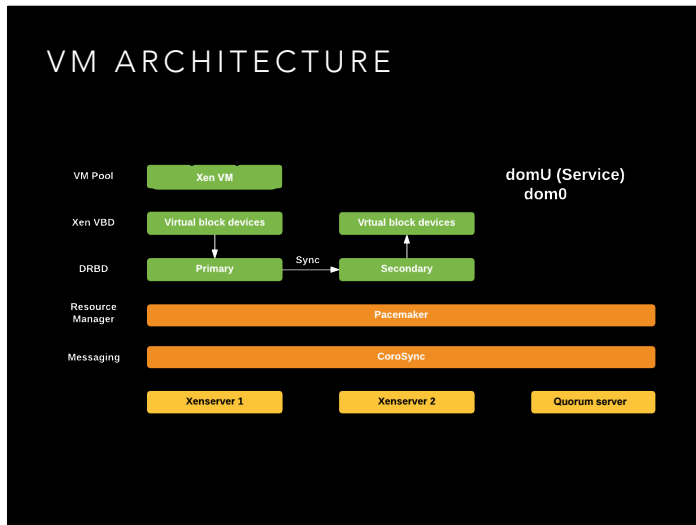
- 1 LVM PV (25G)
 - root LV (4.5G)
 - swap LV (0.5G)
 - thin pool (20G)
 - user data LV (10G) /replicated
 - snapshots

UNDERLYING STORAGE



A DRBD device is generated for each one of the VMs at dom0 level. These devices have each one a domU logical volume as their storage. On top of the DRBD solution, a physical volume is created. On top of that physical volume three logical volumes are created, one for swap, one for root and another for replicated. Guests (domU) only see the DRBD as their storage, therefore only a physical volume with three logical volumes.

VM ARCHITECTURE



Storage replication is done at a higher level in dom0 instead of in domU. Clusters are formed by two xen servers and a dummy VM as a 3-node cluster with each xen guest being a "resource" in the cluster. Pacemaker tend to divide these "resources" up between the two Xen hosts (but they can be moved from one to the other as necessary) domU VMs are treated as pacemaker resources and therefore are moved from one xen server to another in case of failure.

