



Wikimedia Technical & Operational Infrastructure

A high level overview of Wikimedia
Operations

Presented by: Rob Halsell

Wikimania 2009 - Buenos Aires,
Argentina
August 26-28, 2009

Operations Personnel

- Pretty much everyone does some form of operations.
- 16 shell users, 10 of them with root.
- We are slowly moving into distinct development and operations roles, however we will probably never move into fully dev and ops roles only.

Datacenters

- Wikimedia currently actively occupies 4 different sites.
 - Two are in Amsterdam, Netherlands.
 - Two are in Tampa, FL, USA.
- New Primary DC soon in Ashburn, Virginia, USA.

DC Deployments

- The initial servers were in Tampa because that is where Jimmy lived.
- Tampa is not ideal for network transit or routing.
- Hurricanes!
- With new installs being mainly lights out manageable, on-site work is significantly reduced.

Tampa DC Overview

- Currently occupy two datacenters (PMTPA & SDTPA.)
- They are in the same building ;_;
- Each site has a router, and all traffic is routed out via PMTPA.
- SDTPA was deployed due to power and cooling restrictions in PMTPA.

Amsterdam DC Overview

- We have two locations in Amsterdam (KNAMS & ESAMS).
- KNAMS is strictly for peering and transit.
- ESAMS houses servers and associated kit.
- The majority of transit costs are donated via sponsorship and peering!
- We pay for some cross-connects and our fiber between KNAMS & ESAMS.

Daily Operations

- Operations is NOT software development.
- Site Requests: extensions, groups, configuration changes requiring shell access.
- Repair & Upgrade: broken hardware replacement.
- Deployment: new hardware allocation and setup.
- Tweaking: Adding new management tools, software...
- Inventory Management

Site Requests

ALL YOUR BUG
ARE BELONG
TO ME!



- Primarily handled in Bugzilla.
- bugzilla.wikimedia.org
- Sometimes handled by poking Rob in IRC and giving him the Bugzilla Ticket #.
- Changes such as logo, user groups, permissions for group assignments, extensions....
- Some things still get delayed, but overall more and more tickets are handled faster and faster.

Repairs

- Most new servers are either Dell or Sun
 - This makes ease of deployment and RMA repairs much easier.
- Repairs include any physical required access to the servers, such as hard disk replacement, memory upgrades, and the like.
- All new servers have remote management capabilities, so on-site hands are not required for day to day operations.

Deployment

- New deployments are FUN!
- They are also very time consuming to do it properly.
 - Designing takes time.
 - Implementing takes even more time.
 - We label everything!
- Regular Occurrence: Faster hardware for demands.


Tweaking



- This includes new software testing for operations use.
- Testing internal software development.
- Changing settings in things such as OTRS or Bugzilla.
- Updating blog software and the like...


Inventory Management

- We currently use Racktables to keep track of all our servers.
- We are also developing our own extensions to do this within Mediawiki. \o/
- If any extension developers want to work on this, we welcome it!
- Yearly audits require an accurate accounting for all hardware, software, and spending.



Cool Numbers

- Total # of servers:
- Requests per second at peak: 75,000 requests
- Comscore rates Wikipedia as #5 most visited site.



Presented by: Rob Halsell

Wikimania 2009 - Buenos Aires,
Argentina

August 26-28, 2009

Traditional Web Server

- Two major components: Apache and MySQL

Apache

MySQL

Wikimedia Web Server

- Squid > Apache > Memcached > MySQL



Presented by: Rob Halsell

Wikimania 2009 - Buenos Aires,
Argentina

August 26-28, 2009

Squid

- Caching Servers
- Currently deployed in both Tampa and Amsterdam
- Caching stores info so requested pages do not need to be recompiled by Apache.

Squid



Apache



Memcached



MySQL

Apache

- Does the heavy lifting.
- Currently deployed only in Tampa
- When a page changes, Apache pulls data from Memcached and MySQL to render the page for viewing.

Squid



Apache



Memcached



MySQL

Memcached

- Run on apache servers.
- Stores recently fetched database data in resident memory.
- No need to hit the database and pull lots of data if its in memory.

Squid



Apache



Memcached



MySQL

MySQL

- Run out of Tampa datacenters.
- Sites are broken into Database clusters.
 - s1 for English Wikipedia, s2 for the 19 other largest wikis, s3 & s4 for remaining 750+ wikis
- Each cluster has a master and multiple replication servers.

Squid



Apache



Memcached



MySQL

Software

- Ubuntu 8.04
- We run our own netboot/PXE server, as well as our own apt-repository.
- We create some custom packages for ease of installation (our apache and squid servers are very simple to reinstall)

More Software

- Apache 2.2
- PHP5
- Memcached 1.2.8
- Squid Version: 2.7
- Load balancing for squid and apache is handled on one of the 4 LVS servers. Each server runs pybal, which is written by our own Mark Bergsma.

MySQL

- There are multiple database roles within the cluster.
- Primary Databases contain revision data for every page and article of the projects.
- External Store databases contain the text data for every revision.
- Images are stored on an NFS server and apaches access it. Basically this stinks but we don't have a better way... yet. (Not really MvSQL related)

Neat Operations Tools



- Nagios
- Ganglia
- Torrus
- Racktables
- Mediawiki ;]

Presented by: Rob Halsell

Wikimania 2009 - Buenos Aires,
Argentina
August 26-28, 2009

Nagios

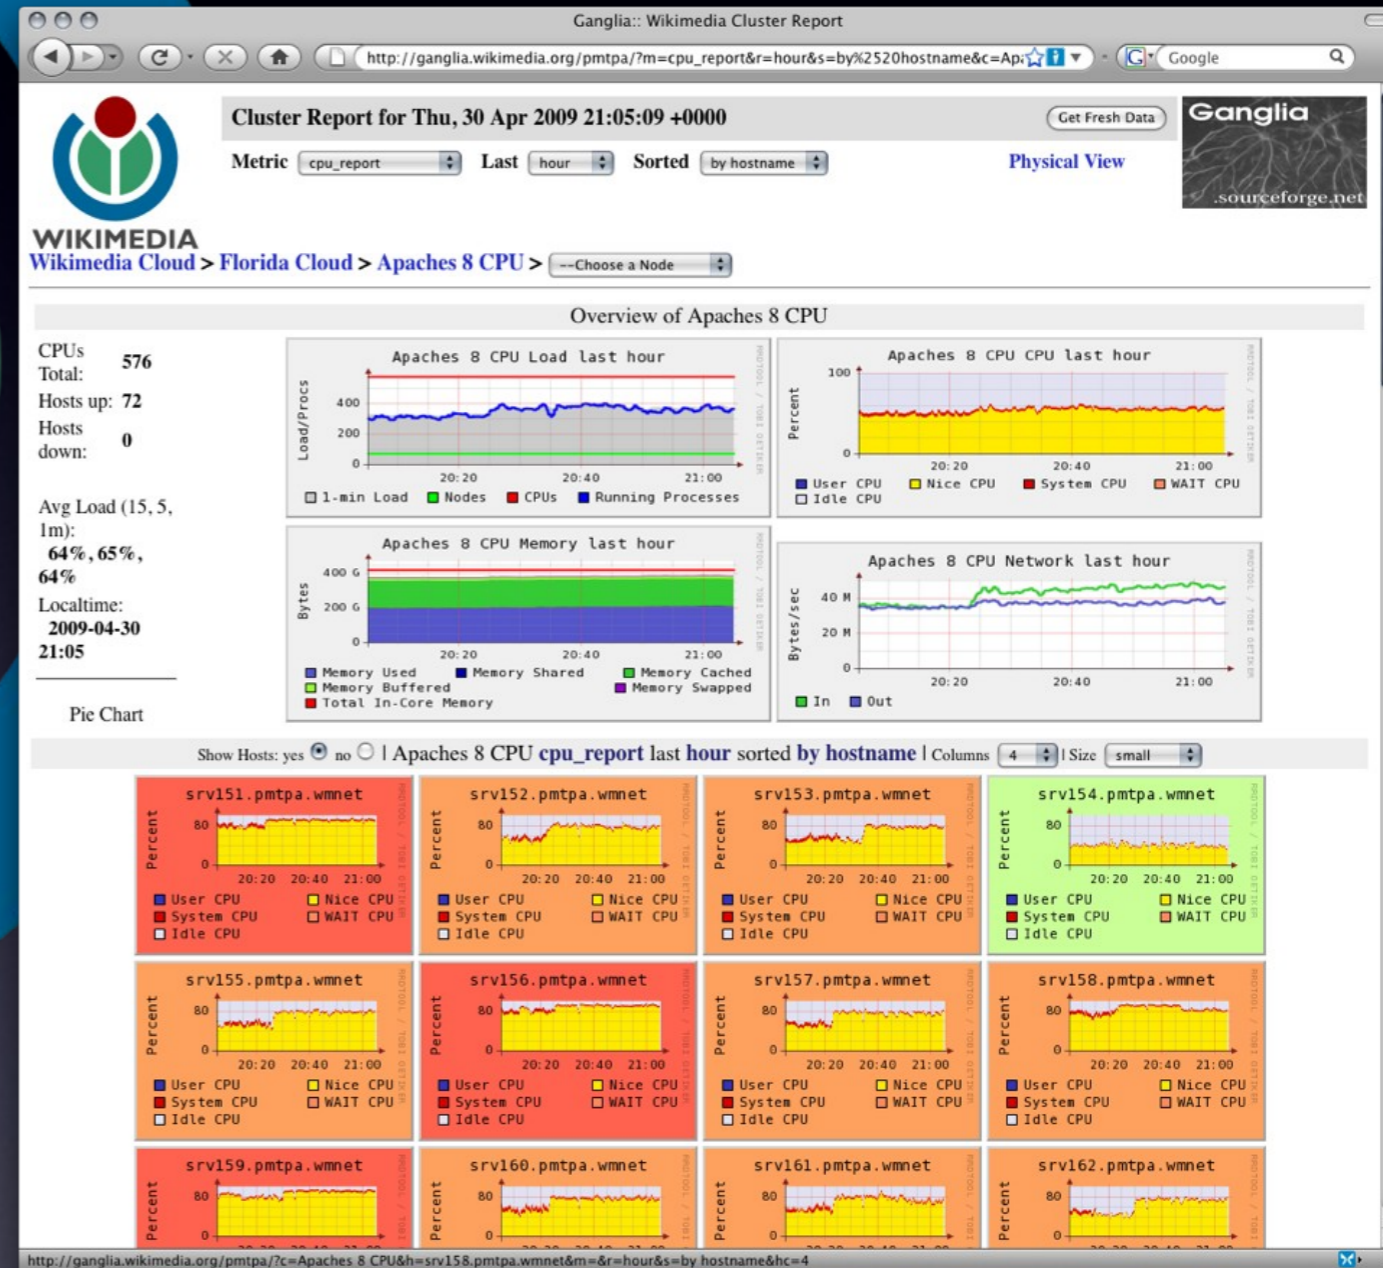
- Shows all servers, and the services they are running.
- Shows downtime of various services.
- nagios.wikimedia.org

The screenshot displays the Nagios web interface. At the top, it shows the current network status: Last Updated: Thu Apr 30 20:59:37 UTC 2009, Updated every 90 seconds, Nagios® - www.nagios.org, and Logged in as ?. Below this are three summary tables: Host Status Totals (Up: 399, Down: 9, Unreachable: 1, Pending: 5), Service Status Totals (Ok: 879, Warning: 4, Unknown: 0, Critical: 38, Pending: 0), and All Problems (10, 414). A 'Display Filters' section shows Host Status Types: Down, Host Properties: Any, Service Status Types: All Problems, and Service Properties: Any. The main section is 'Service Status Details For All Hosts', which contains a table with 14 rows of service entries, all marked as CRITICAL. The table columns are Host, Service, Status, Last Check, Duration, Attempt, and Status Information. The status information for most entries is 'CRITICAL - Socket timeout after 10 seconds', while some are 'Connection refused'.

Host	Service	Status	Last Check	Duration	Attempt	Status Information
svr02	SSH	CRITICAL	04-30-2009 20:56:55	36d 20h 34m 44s	1/3	CRITICAL - Socket timeout after 10 seconds
svr03	Apache	CRITICAL	04-30-2009 20:57:53	4d 21h 55m 0s	1/3	CRITICAL - Socket timeout after 10 seconds
	SSH status	CRITICAL	04-30-2009 20:56:54	4d 21h 55m 4s	1/3	CRITICAL - Socket timeout after 10 seconds
svr08	Lucene	CRITICAL	04-30-2009 20:56:55	6d 3h 39m 21s	1/6	Connection refused
svr07	Apache	CRITICAL	04-30-2009 20:56:55	10d 0h 56m 17s	1/3	CRITICAL - Socket timeout after 10 seconds
	SSH status	CRITICAL	04-30-2009 20:57:53	10d 0h 56m 17s	1/3	CRITICAL - Socket timeout after 10 seconds
svr09	Apache	CRITICAL	04-30-2009 20:57:53	10d 0h 54m 42s	1/3	CRITICAL - Socket timeout after 10 seconds
	SSH status	CRITICAL	04-30-2009 20:56:54	10d 0h 58m 46s	1/3	CRITICAL - Socket timeout after 10 seconds
svr118	Apache	CRITICAL	04-30-2009 20:56:54	9d 23h 29m 30s	1/3	CRITICAL - Socket timeout after 10 seconds
	SSH status	CRITICAL	04-30-2009 20:56:55	9d 23h 29m 34s	1/3	CRITICAL - Socket timeout after 10 seconds
svr131	Apache	CRITICAL	04-30-2009 20:57:53	1d 3h 46m 32s	1/3	CRITICAL - Socket timeout after 10 seconds
	SSH status	CRITICAL	04-30-2009 20:57:53	1d 3h 44m 44s	1/3	CRITICAL - Socket timeout after 10 seconds
storage1	Disk free	CRITICAL	04-30-2009 20:57:54	148d 1h 34m 22s	1/3	Connection refused by host
	rsyncd	CRITICAL	04-30-2009 20:57:54	148d 0h 51m 54s	1/3	Connection refused

Ganglia

- Distributed monitoring software, graphing, and reviewing tool for servers.
- Shows us how hard the servers are working.
- ganglia.wikimedia.org



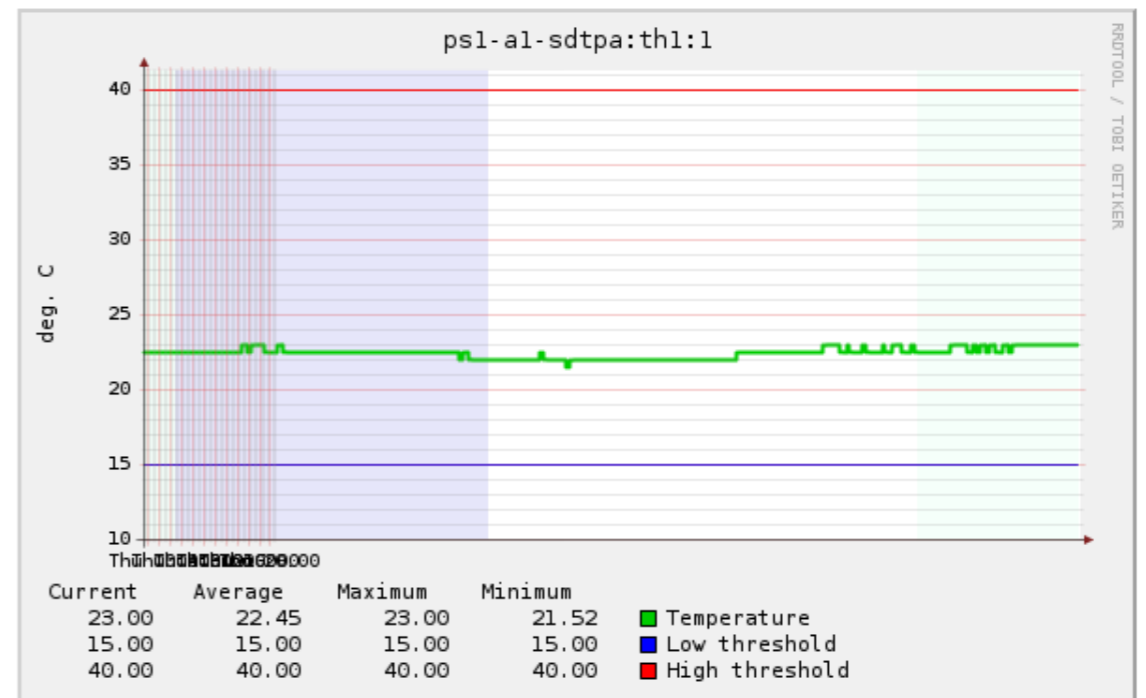
Wikimania 2009 - Buenos Aires,
Argentina
August 26-28, 2009

Presented by: Rob Halsell

Torrus

- SNMP monitoring and graphing software.
- We capture and display network usage, power and temperatures in our datacenters, squid performance, and other items.
- torrus.wikimedia.org

Last day graph



Racktables

- No URL needed, its locked and no one can login unless we employ you ;]
- Keeps track of every server, purchase price, date of purchase, warranty information, location, etc...

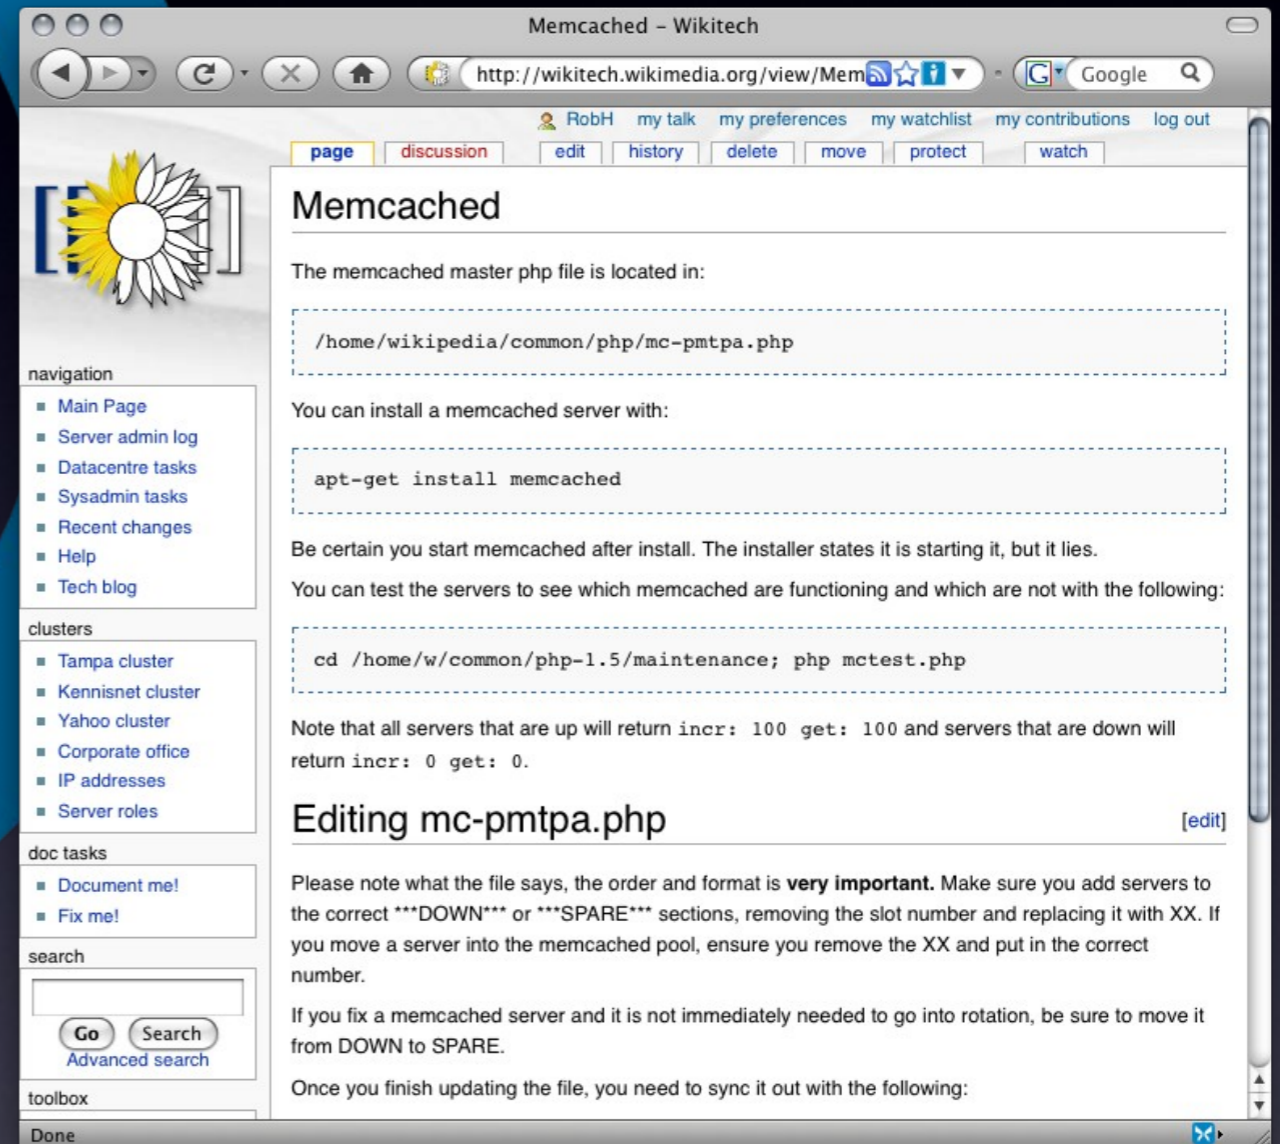
Rack diagram

sdtpa row A : A1 > > >

	Front	Interior	Back
42			
41			
40			
39			
38			
37			
36			
35			
34			
33			
32			
31			
30			
29		csw1-sdtpa	
28			
27			
26			
25			
24			
23			
22			
21			
20		msw1-sdtpa	
19			
18			
17		scs-a1-sdtpa	
16			
15			
14			
13		ms3	
12			
11			
10			
9			
8		ms4	
7			

Wikitech (MediaWiki)

- Mediawiki Installation that we use to document setup, maintenance, recovery, etc...
- Exists outside the cluster, so if the cluster goes down, we still have notes on how to fix it.
- wikitech.wikimedia.org



In Closing...

- Thanks for listening!
- Wikimania Rocks!
- Buenos Aires Rocks!
- Questions?

Presented by: Rob Halsell

Wikimania 2009 - Buenos Aires,
Argentina
August 26-28, 2009

Questions from Presentation

- What is the hit rate on the caching servers?
 - 98% on images, 90-95% on text.
- How much is WMF spending on new servers and associated kit this year?
- Not counting the new DC deployment, we will spend 1M USD just keeping up with growth trends (both use and software feature growth.)
- There were a number of others, I just do not recall them all, sorry =/