

Running an OpenStreetMap cache server in Asia

Steven Shiau, Ceasar Sun, Thomas Tsai

NCHC, Taiwan

Q3, 2015




Outline

- OSM Cache server
 - Why?
 - NCHC's capacity
 - Hardware, network
 - OS, services
 - Maintenance
- Q&A



OpenStreetMap

- A free, editable map of the whole world.
- Built by volunteers largely from scratch and released with an open-content license.
- Allows free access to map images and all of the underlying map data.



* Image source: <http://www.openstreetmap.org>

Why? - Background

- OSC 2013 Aug Kansai@Kyoto
- **No any cache server in Asia at that time**
 - Long loading time
- Daniel Kastl from Georepublic asked, and mentioned:
 - “Universities here have a terrible **administrative overhead** with lots of formal requirements. Even community members working at universities seem to try to avoid the paperwork.”
 - “Data center providers we talked to are mostly "scared" about the **data traffic**. In general internet speed in Japan is super fast, and traffic is unlimited for private users. But it seems the mix of "power-users" and "low-traffic" users, which makes "unlimited traffic" possible. After talking to data center providers it seemed to me, that internet traffic in fact is quite expensive in Japan. Hardware costs were not really an issue for them.”

オープンソースの「今」を伝える

オープンソースカンファレンス
2013 Kansai@Kyoto

Background – network bandwidth

- The traffic is distributed by tile.openstreetmap.org using **GeoDNS** to pick the "local" server. In partnership with the cache provider we (OSM sysadmins) decide which countries are best served by a particular server. See:
<http://dns.openstreetmap.org/tile.openstreetmap.org.html> for current setup.
- Traffic:
 - Using **May 2012** statistics, Japan uses around **312 Kilobytes/s** (inbound+outbound) of tile traffic (averaged over 7 days)
 - Peak will be around double that, low being around half.
 - The tile rendering server is based in the UK (AS786), normal cache byte hit ratio is around 80%.
 - The servers are constantly monitored, traffic is automatically redistributed if a greater than 5min outage occurs.
 - Expected growth rate is around 3% per month.

* Quoted from OSM systemadm team, provided by Daniel Kastl

Background – network bandwidth in Sep/2013

- Traffic Estimates per country averaged over 24 hours during week:
 - Bangladesh 3.71 KBytes/s outbound
 - Cambodia 4.2 KBytes/s outbound
 - China 169.63 KBytes/s outbound
 - Hong Kong 32.14 KBytes/s outbound
 - India 322.86 KBytes/s outbound
 - Indonesia 72.06 KBytes/s outbound
 - Japan 208.28 KBytes/s outbound
 - Laos 2.31 KBytes/s outbound
 - Malaysia 23.30 KBytes/s outbound
 - Myanmar 3.50 KBytes/s outbound
 - Nepal 7.07 KBytes/s outbound
 - North Korea 0.02 KBytes/s outbound
 - Pakistan 17.01 KBytes/s outbound
 - Philippines 149.40 KBytes/s outbound
 - Singapore 50.94 KBytes/s outbound
 - South Korea 68.06 KBytes/s outbound
 - Taiwan 63.60 KBytes/s outbound
 - Vietnam 42.95 KBytes/s outbound
 - Total: 1241.129 KBytes/s outbound
- Inbound is approximately 10% of outbound.

* Quoted from OSM systemadm team, Grant Slater

OSM Tile CDN

- Tile CDN (Content Delivery Network)
- Cache isn't a file mirror, it is a proxy + caching setup
 - **Not a files mirror only**. Therefore the mechanism is more complicated than an open source/free software mirror site.
 - **Two months** in communication with OSM sysadmin team (Grant Slater).
 - We spent **4 more months** to find the solution to follow the administration policy at NCHC, prepare and setup the machine.



National Center for High-Performance Computing



1991
Officially Founded

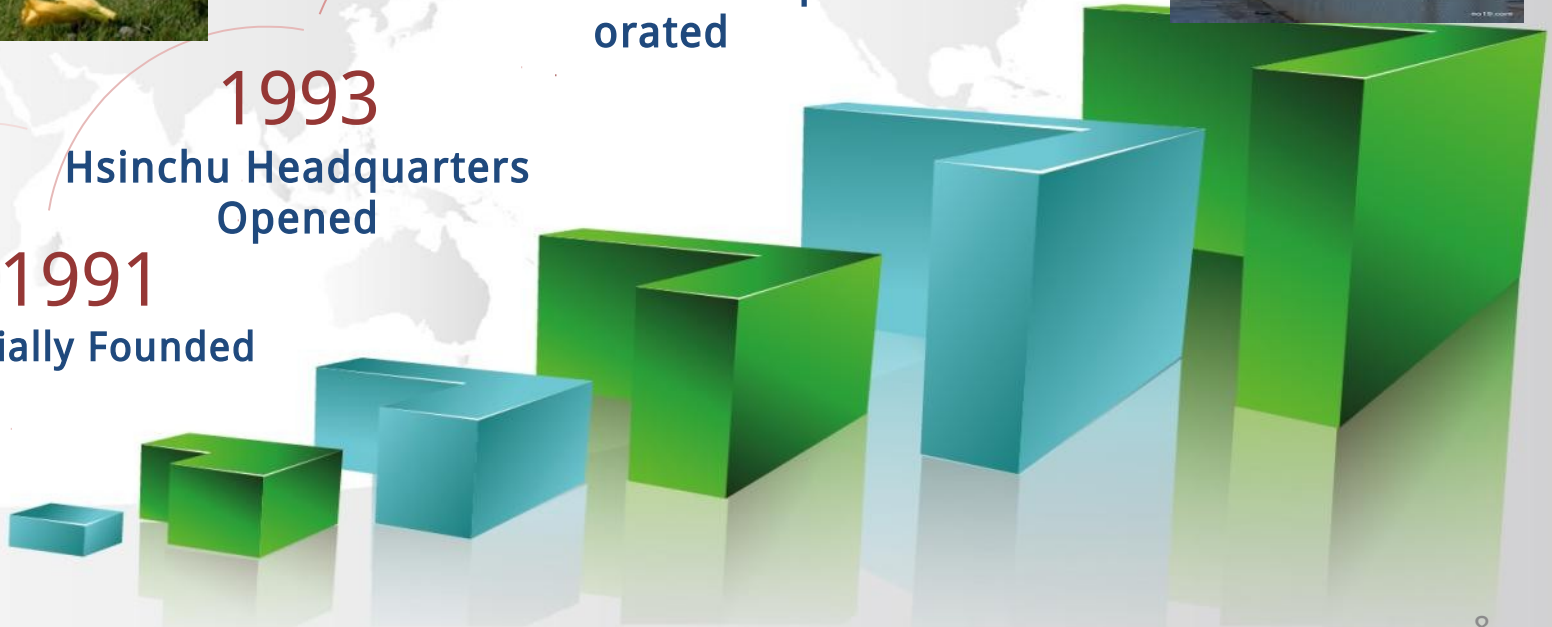
1993
Hsinchu Headquarters
Opened

2003
Became Incorporated

2005
Tainan Office
Opened

2008
Taichung Office
Opened

1988
Started
Planning



HPC Services

- Open to academic, research, and Industrial users
- Supporting 700+ research projects per year



- ALPS system – most recent supercomputer built in 2011

- R_{max} 177 TFLOPS sustained, 442.00 MFLOPS/W

- 25,600 Cores • 73,728 GB Memory • 1,074 TB Disk
- Jun. 2011: Top500 Ranking: No. 42 / Green500 Ranking: No. 25

御風者
WINDRIDER

Advanced Large-scale Parallel Supercluster (ALPS)




IBM Cluster 1350 / 19.91TF



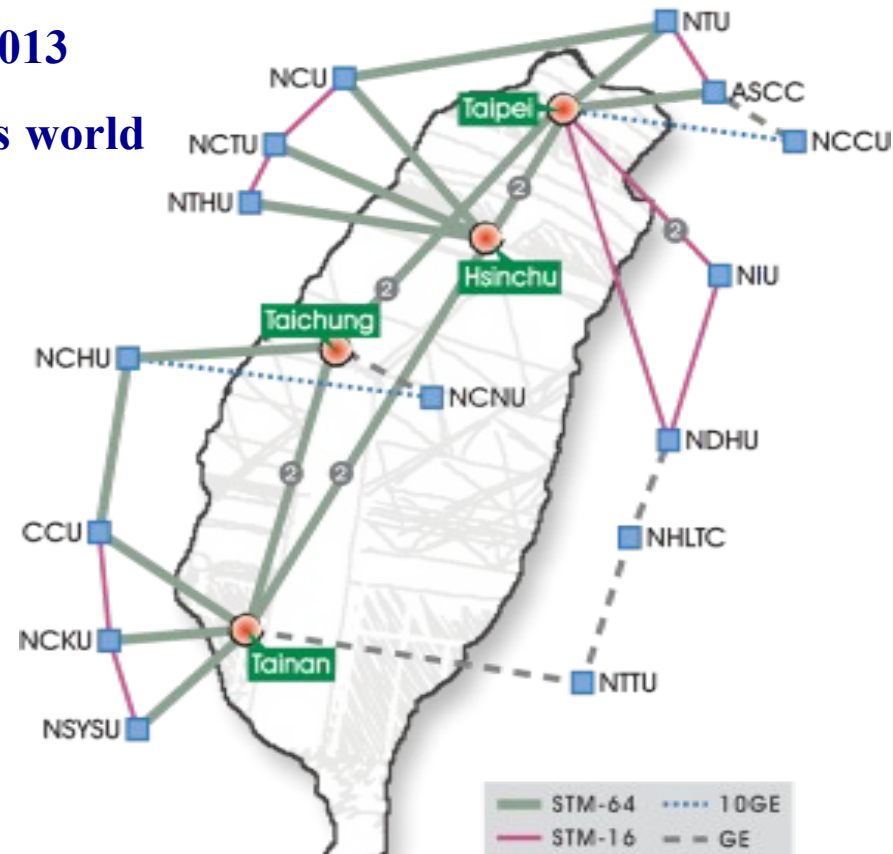
Formosa 5/ 89.9 TF

NCHC Total Computing Capacity



Research and Education Network

- Providing research network, education network (TANet), and optical lightpath services with **20 Gbps backbone**
- Working toward 100Gbps backbone from 2013
- Peering with 35 IPv4 and 24 IPv6 networks worldwide with **5Gbps connection**
- Network availability rate up to 99.991%
- Dynamic circuit provisioning enabled



Storage Services

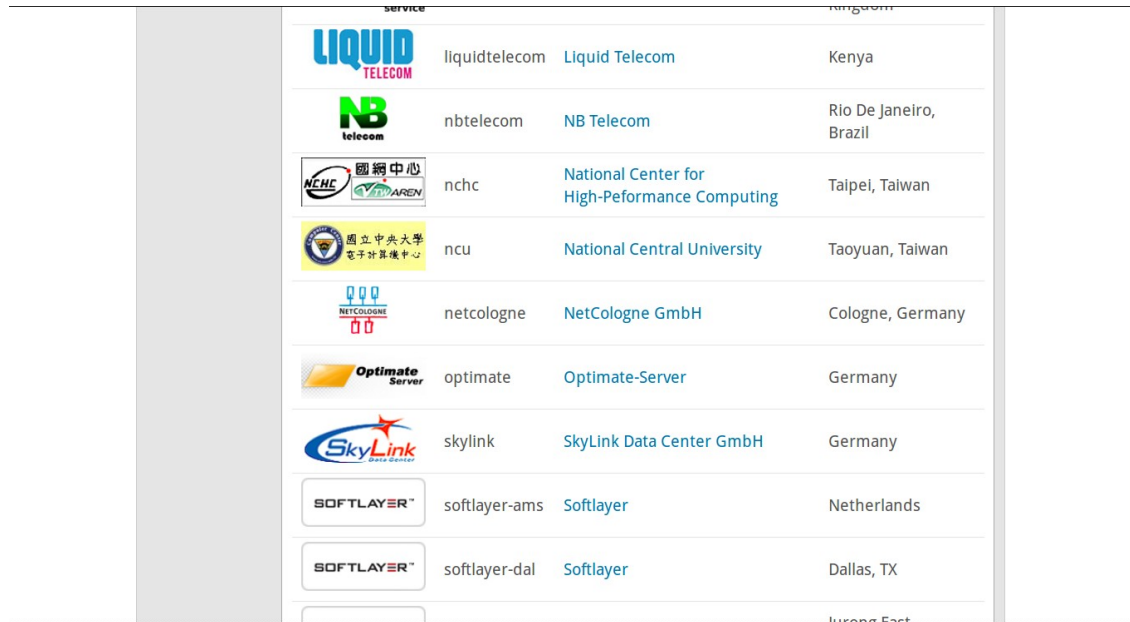





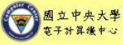





Storage Capacity

- Three-site, 3-tier backup
- Total 5.4 PB Capacity
- Supports 30+ projects from academia and research institutes
- Deploys disk and tape facilities in Hsinchu, Taichung, and Tainan; Interconnected via TWAREN and Storage Area Network (SAN)

Some mirrors@NCHC

- **Sourceforge** mirror site from 2005
- Major GNU/Linux distribution and **OpenSource/Free Software mirrors**: <http://free.nchc.org.tw>
 - CentOS, Debian, Fedora, Gentoo, Linux Mint, Ubuntu, OpenSuSE, VLC, Firefox...



service	mirror	location
	liquidtelecom Liquid Telecom	Kenya
	nbtelecom NB Telecom	Rio De Janeiro, Brazil
	nchc National Center for High-Performance Computing	Taipei, Taiwan
	ncu National Central University	Taoyuan, Taiwan
	netcologne NetCologne GmbH	Cologne, Germany
	optimate Optimate-Server	Germany
	skylink SkyLink Data Center GmbH	Germany
	softlayer-ams Softlayer	Netherlands
	softlayer-dal Softlayer	Dallas, TX
		Huang East

Source: <http://sourceforge.net/p/forge/documentation/Mirrors/>

About us

- Free Software Lab, NCHC, Taiwan
- Developers of the free software DRBL, Clonezilla Partclone, DRBL-Winroll, and more...
- Steven is also the maintainer of GParted live CD



Taiwan image source: wikipedia.org

How?

- Procedure to setup a cache server
 - 1) Install a server running Ubuntu 12.04 (AMD64)
 - 2) Create an account for OSM systemadm
 - 3) OSM systemadm logs in and setup management setup (Chef) which installs +configures everything needed.
 - 4) OSM systemadm moves a little traffic for first week and feed back to cache server owner.
- Cache server owner will have login access to machine at all times. Any **shutdowns or disconnects** will automatically be detected by OSM system and the server will automatically be removed from the pool.

* Quoted from OSM systemadm team, Grant Slater

Basic requirements

- http://wiki.openstreetmap.org/wiki/Servers/Tile_CDN
 - Basic regional tile delivery server requirements:
 - 16 GB RAM (at least; better 32 GB);
 - **Fast network connection with high usage or unlimited traffic**; (Traffic is directed by GeoDNS)
 - **Full root/sudo access** (Remote Management beneficial eg: HP Integrated Lights-Out);
 - Ubuntu 14.04 LTS 64-bit (AMD64);
 - Storage of at least 146GB excluding OS. (10kRPM disk or better preferred)



Tile server@NCHC

Longma 龍馬

- Hardware
 - CPU: Intel Xeon CPU E5-2620 v2 @ 2.10GHz, 6 cores
 - RAM: 32 GB
 - Hard drives: 160 GB SATA disk and 400 GB SATA disk
 - Two Gigabits Ethernet cards
- OS
 - Ubuntu 12.04 LTS (2014/02-2014/12)
 - Ubuntu 14.04 LTS (2014/12-Now)
- Proxy server: Squid 2.7.STABLE9
 - OSM systemadm planned to upgrade to varnish 3.x



Source: <http://design.ubuntu.com>; <http://www.squid-cache.org>

1st cache server in Aisa

- On **Feb/19/2014**, the cache server was ready, and some test traffic was redirected
- It was until **Jan/02/2015** OSM.org announced it on the blog:
- <https://blog.openstreetmap.org/2015/01/02/four-new-tile-servers/>



OpenStreetMap blog

LANGUAGES

 English

BLOGROLL

[Blogs.OpenStreetMap.org](https://blogs.openstreetmap.org/)

RECENT POSTS

[Fundraising drive 2015](#)

Four New Tile Servers

Have you noticed faster tiles lately? Browsing the map on openstreetmap.org should now be even more responsive. Three new servers, started providing tiles over the last 2 weeks, joining a server which started earlier in the year.

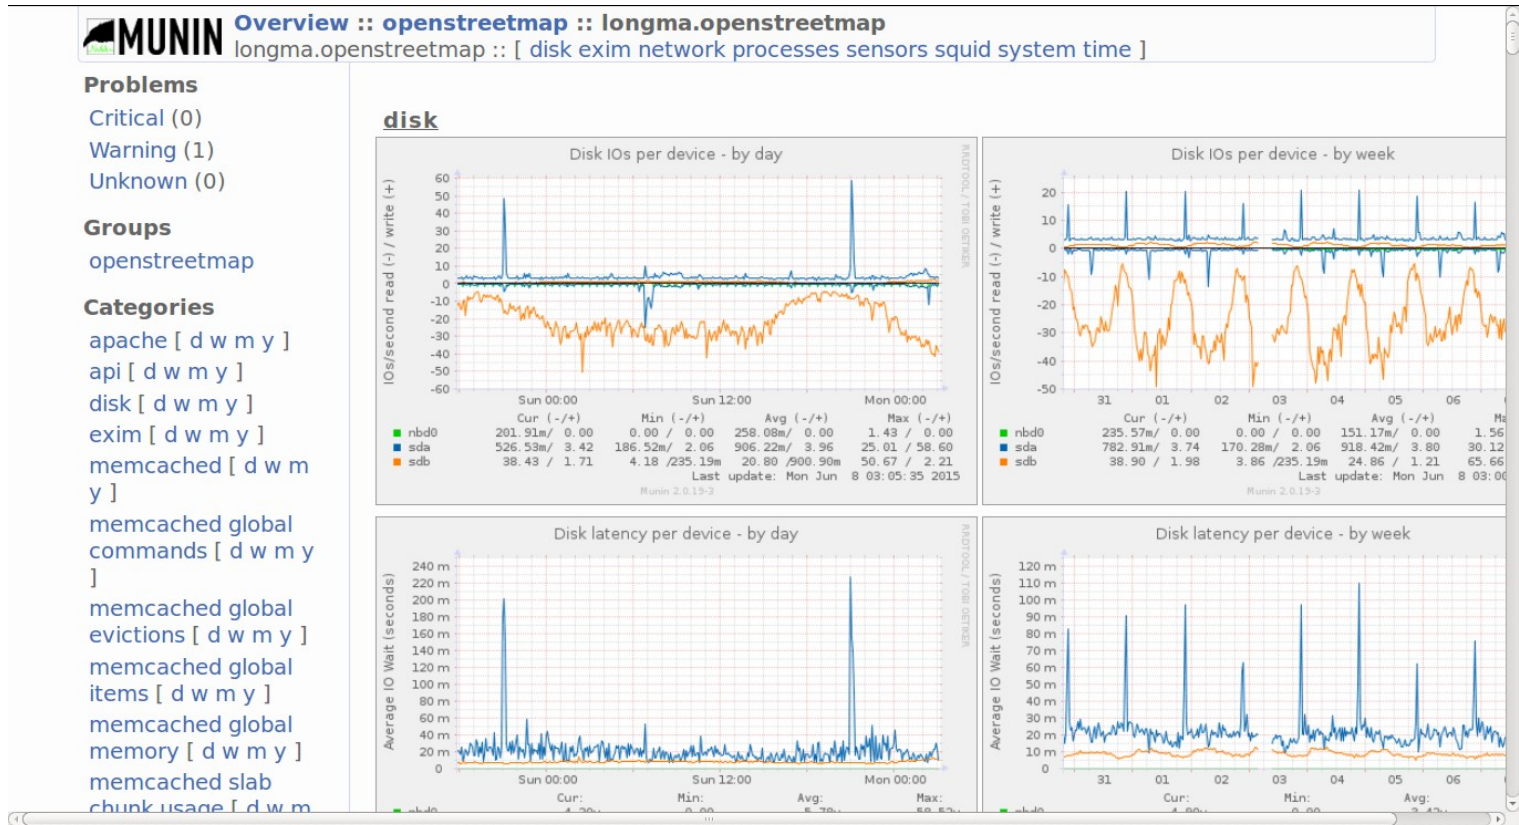
- Tile server [saphira](#), located in London UK kindly hosted by [Jump Networks](#).
- Tile server [viserion](#), located in Pula Croatia, kindly hosted by [CARNet](#).
- Tile server [stormfly-02](#), Located in Corvallis USA, kindly hosted by [OSUOSL](#).
- Tile server [longma](#), Located in Hsinchu Taiwan, kindly hosted by [NCHC](#).



Map tiles are delivered to users based on their [GeoDNS location](#). The OpenStreetMap tile content delivery network (CDN) now supports [EDNS-client-subnet](#) to improve locating the closest region tile cache


System Monitoring

- Munin
- <http://munin.osm.org/openstreetmap/longma.openstreetmap/index.html>



Contact with OSM Systemadm team

- Email
 - operations@osmfoundation.org
- Jabber
- IRC:
 - #osm-dev on oftc network
 - Also available via <http://irc.OpenStreetMap.org>



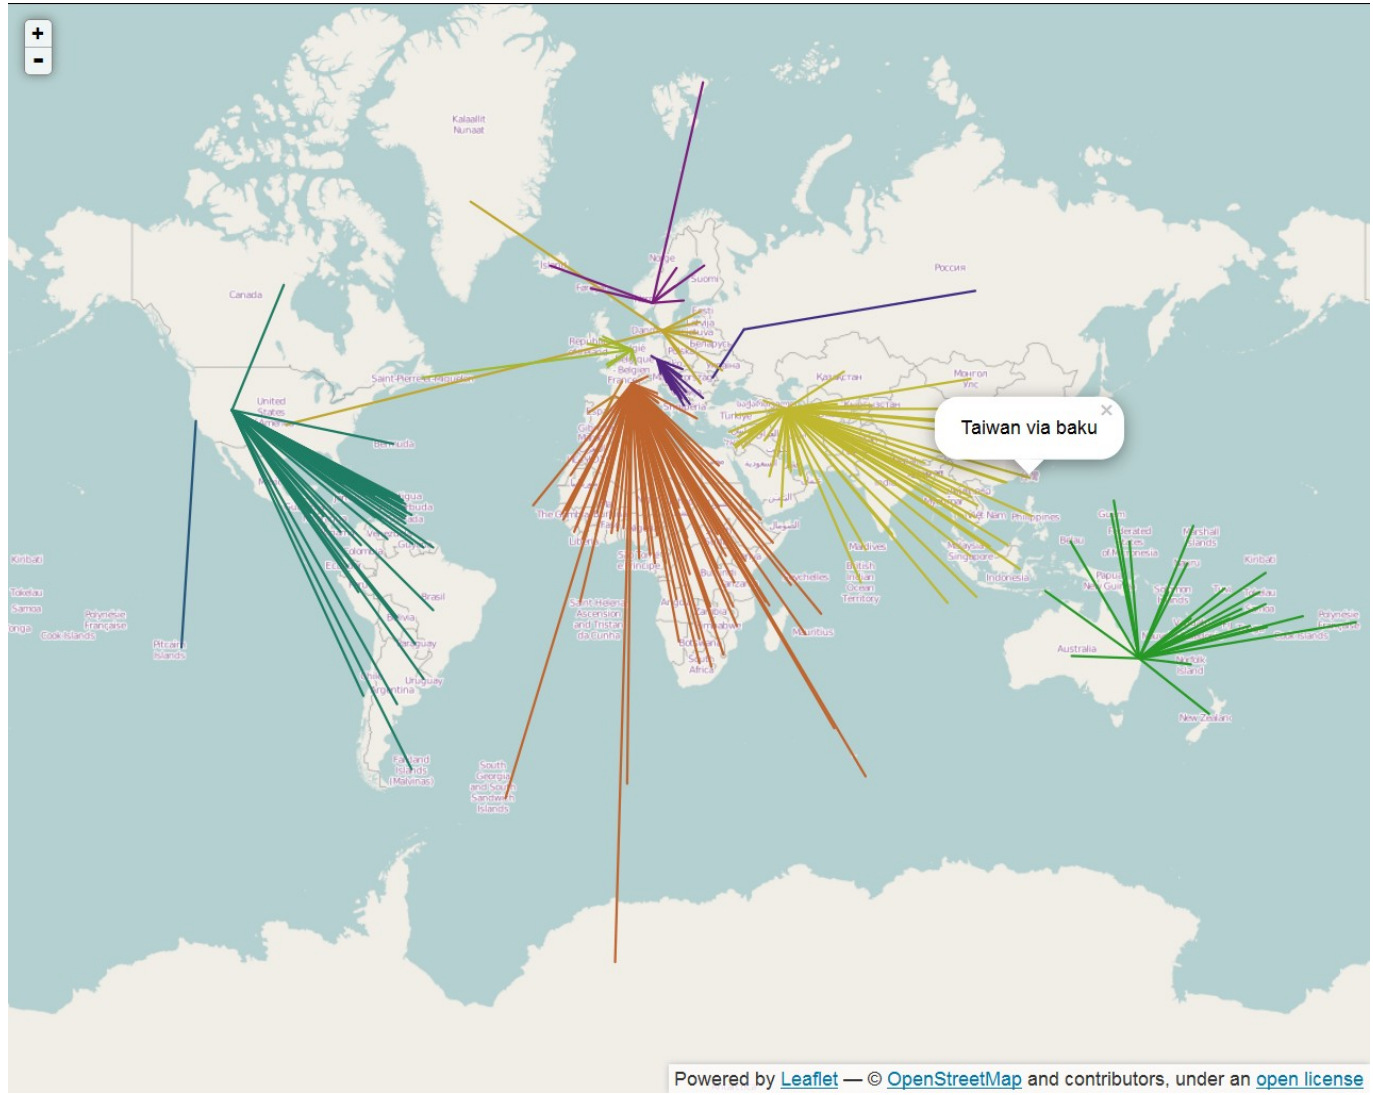
Source: (1) <http://en.wikipedia.org/> (2)<http://jabber.org> (3) <http://www.oftc.net>

Having problems?

1. Try to fix the issue by ourselves first
2. **Reboot the tile server**
3. Ask OSM systemadm to solve the issue remotely



Before



Source: <http://dns.openstreetmap.org/tile.openstreetmap.org.html> on 2014/Jan

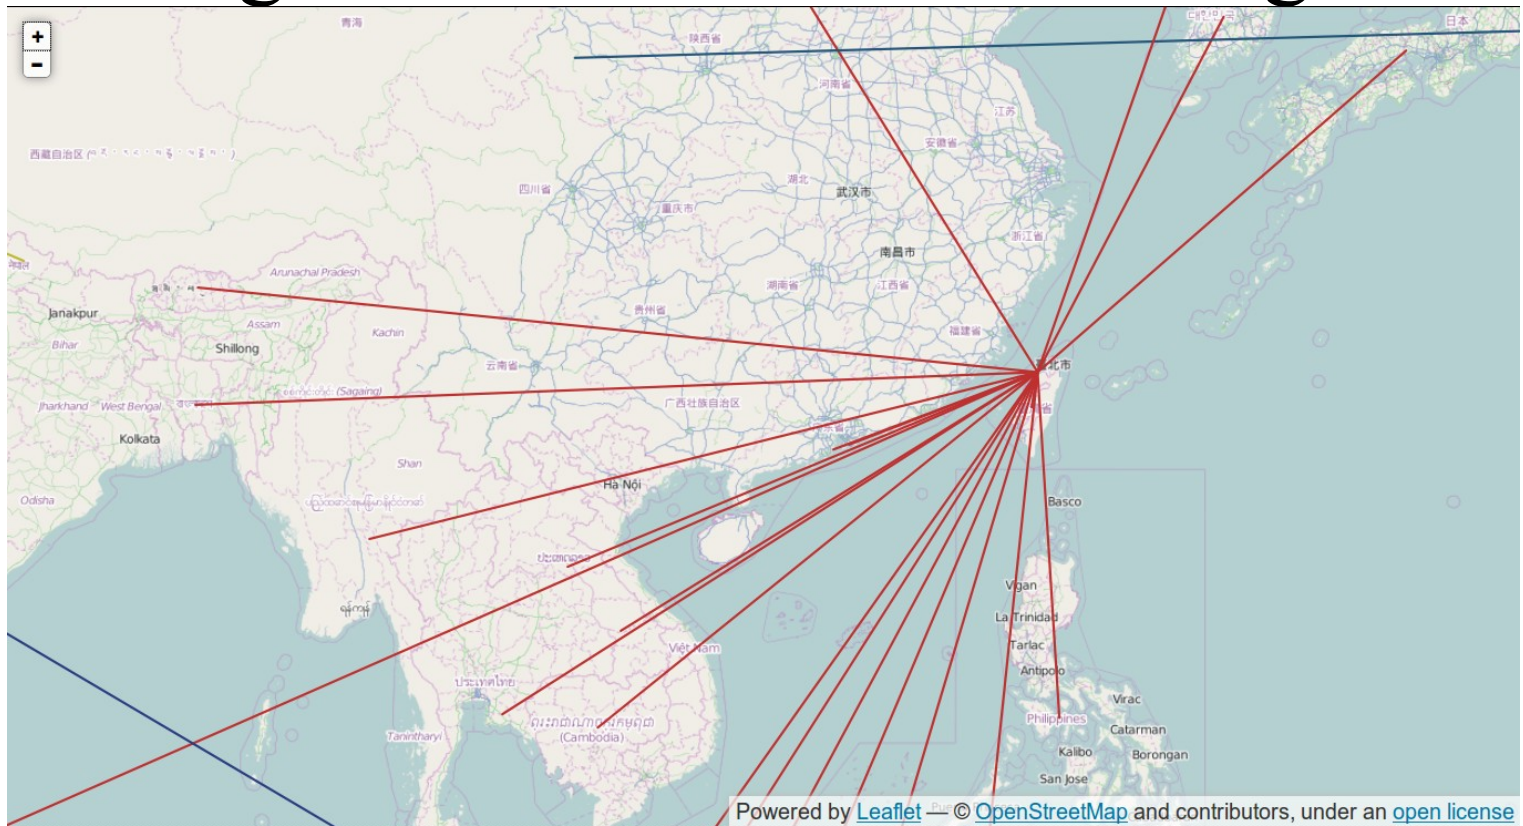
After



Powered by [Leaflet](#) — © [OpenStreetMap](#) and contributors, under an [open license](#)

Source: <http://dns.openstreetmap.org/tile.openstreetmap.org.html> on 2015/May


Longma serves 22 countries/regions




Taiwan, Japan, South Korea, North Korea, Mongolia, Bhutan,
Bengal, Myanmar, Laos, Sri Lanka, Vietnam, Thailand,
Macao, Hong Kong, Cambodia, Singapore,
Cocos (Keeling) islands, Malaysia, Christmas island, Brunei,
Indonesia, Philippines

Source: <http://dns.openstreetmap.org/tile.openstreetmap.org.html> on 2015/May

Stats: System loading




Stats: CPU usage




6 cores + hyper threading, so there are 12 CPUs in Longma

Stats: Memory usage



Stats: Network traffic




Average is about 8.7 Mbps (out) / 2.8 Mbps (in)

~0.17% (out)/0.06% (in) of NCHC's total bandwidth


Throughput: 2.8 TB/month (out) / 0.9 TB/month (in)

*There was a firewall configuration issue at NCHC in early Jan 2015


Stats: Squid cache




Stats: Squid client requests



Stats: Squid traffic

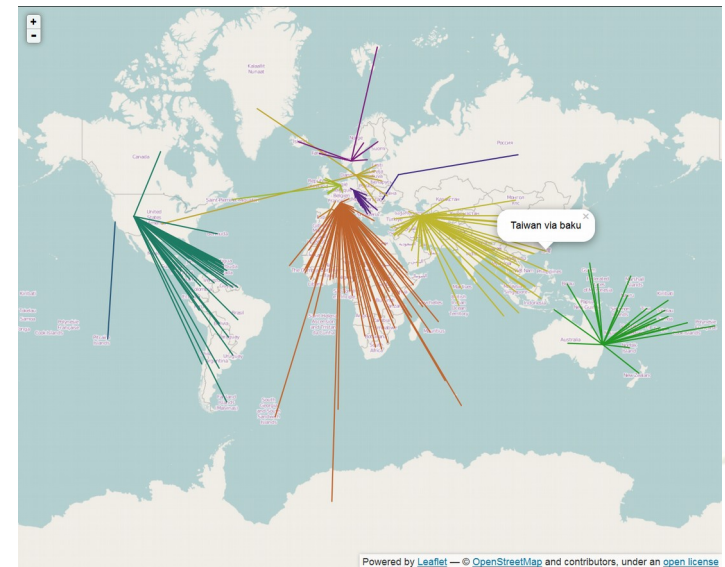
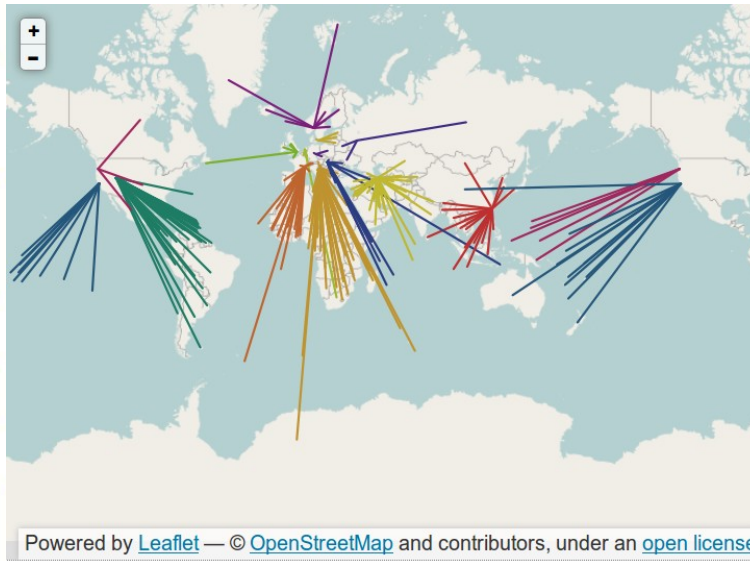


Stats: System uptime



But we are still alone in Asia...

- Asian users need more cache servers
- Each other as a redundant server



Conclusions

- OpenStreetMap systemadm team has a **very good mechanism** to setup and monitor the tile server. Therefore the efforts we spend on the system maintenance is minimum.
- More cache servers are needed in Asia. The **redundant mechanism** has to be established.



Acknowledgement

- This work is sponsored by MOST (Ministry of Science and Technology), Taiwan



Reference

- OpenStreetMap: <http://www.openstreetmap.org>
- OpenStreetMap Taiwan: <http://openstreetmap.tw>
- Academia Sinica: <http://www.sinica.edu.tw>
- NCHC: <http://www.nchc.org.tw>
- OSM Munin: <http://openstreetmap.tw>
- DRBL: <http://drbl.org>
- Clonezilla: <http://clonezilla.org>
- Gparted: <http://gparted.org>



Questions ?

