

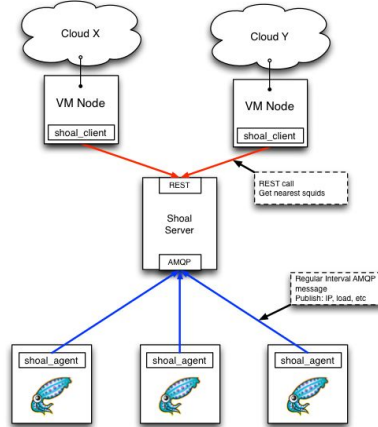
Shoal

a squid cache publishing and advertising tool

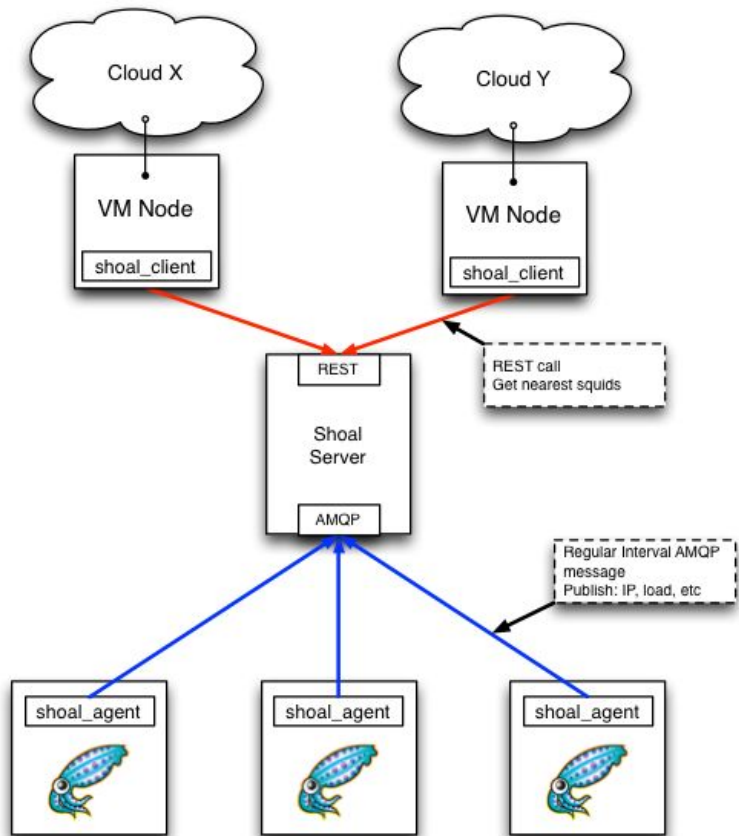
Marcus Ebert (mebert@uvic.ca)

Colson Driemel, Da Meng, Randy Sobie

University of Victoria, Canada



What is Shoal?



- Squid discovery system
 - finds closest (best) squid
- 3 components
 - shoal-server
 - collects information and matches squids to client requests
 - shoal-agent
 - runs on each squid and advertises squid to shoal-server
 - shoal-client
 - queries server for closest squid
- developed in the HEP-RC group of the University of Victoria (UVic)[1]

Shoal at UVic

- HEP-RC group does Grid computing based on [dynamic resource allocation on distributed clouds](#) [2]
 - currently supporting Atlas and Belle-II, testing Dune
- Grid jobs need software from CVMFS repositories
- Grid jobs don't run close to the geographic location of the site or close to each other
 - each Grid job could be on any of the available clouds
 - utilizing clouds in Canada, USA, UK, Germany, Austria, Australia currently
 - support also commercial clouds
- fixed site-squid not a solution

Shoal at UVic

- HEP-RC group does Grid computing based on [dynamic resource allocation on distributed clouds](#) [2]
 - currently supporting Atlas and Belle-II, testing Dune
- Grid jobs need software from CVMFS repositories
- Grid jobs don't run close to the geographic location of the site or close to each other
 - each Grid job could be on any of the available clouds
 - utilizing clouds in Canada, USA, UK, Germany, Austria, Australia currently
 - support also commercial clouds
- fixed site-squid not a solution

That is what shoal was developed for.

Shoal before V1.0.0

- for about 10 years in production use
- based on Python 2
- shoal knows 3 states for a squid
 - Global: squid is accessible from anywhere
 - Private: squid is only accessible from within the same network
 - Local Access Only: squid is accessible from anywhere but shouldn't be used outside of their own network

Shoal before V1.0.0

- shoal-client can reconfigure the local CVMFS system to use specific squids
 - also has FRONTIER option
- server's network matching based on old domain db (commercial, mmdb)
- server's geographic calculations based on city IP db (commercial, mmdb)
- server verifies squid access before squids can be used for client requests
- server can take load of squid and geographic location into account when giving out squids to a client
 - list given out based on hostnames
- shoal-agent measures network load and sends registration to server periodically
 - success of whole system depends heavily on correct shoal-agent configuration
 - sysadmin needs to configure system (firewalls), squid (access pattern), and shoal-agent independently

Shoal before V1.0.0

List of Active Squids

23 active in the last 180 seconds

#	Hostname	Public IP	Private IP	Bytes Out	City	Country	Latitude	Longitude	Last Received	Alive	Clock Sync	Verified	Access Level
1	vobox.spacescience.ro	85.120.46.18	172.20.0.18	6.3583984375 kB/s		Romania	45.9968	24.997	3s	90h49m42s	✓	✗	Global
2	115.146.87.74	115.146.87.74		0 kB/s		Australia	-33.494	143.2104	4s	90h49m36s	✓	✗	Private
3	issaf.spacescience.ro	46.243.114.226	10.99.99.1	0.4990234375 kB/s		Romania	45.9968	24.997	6s	90h49m43s	✓	✗	Global
4	ca-mey-frontier-6d2f31843a.cern.ch	137.138.150.196		29486 kB/s	Geneva	Switzerland	46.2054	6.1459	7s	90h49m55s	✓	✗	Local Access Only
5	ca-mey-frontier-0097bb1c18.cern.ch	137.138.123.72		38551 kB/s	Geneva	Switzerland	46.2054	6.1459	7s	90h49m54s	✓	✗	Local Access Only
6	ca-mey-frontier-bigfin-c6a7fff898.cern.ch	188.184.98.18		24708 kB/s	Geneva	Switzerland	46.2054	6.1459	7s	90h49m30s	✓	✗	Local Access Only
7	atlassq.uibk.ac.at	138.232.146.212	10.10.0.12	1090.11816406 kB/s	Innsbruck	Austria	47.2637	11.4016	10s	90h49m57s	✓	✗	Local Access Only
8	squid-grid.uaic.ro	85.122.31.70		588.745117188 kB/s	Iasi	Romania	47.1662	27.5988	10s	90h49m38s	✓	✗	Global
9	sampapx02.if.usp.br	200.17.30.146	10.10.15.42	1.197265625 kB/s		Brazil	-22.8305	-43.2192	10s	58h56m34s	✓	✗	Global
10	sampaproxy.if.usp.br	200.17.30.145	10.10.15.41	1.4541015625 kB/s		Brazil	-22.8305	-43.2192	10s	58h57m46s	✓	✗	Global
11	marsq03.in2p3.fr	134.158.20.28		6.703125 kB/s		France	48.8582	2.3387	13s	90h49m43s	✓	✗	Global
12	marsq02.in2p3.fr	134.158.20.33		2.9658203125 kB/s		France	48.8582	2.3387	13s	90h49m43s	✓	✗	Global
13	192.168.2.170	129.215.193.72	192.168.2.170	600 kB/s	Edinburgh	United Kingdom	55.951	-3.1972	15s	90h49m34s	✓	✗	Private
14	squid.novalocal	129.215.193.72	192.168.2.170	3265.68066406 kB/s	Edinburgh	United Kingdom	55.951	-3.1972	16s	90h49m36s	✓	✗	Global
15	206.12.91.244	206.12.91.244	10.39.27.67	1736 kB/s	Vancouver	Canada	49.282	-123.1103	17s	90h49m36s	✓	✓	Global
16	js-17-31.jetstream-cloud.org	129.114.17.31	172.24.83.3	0.0 kB/s	Austin	United States	30.3773	-97.71	18s	90h49m51s	✓	✗	Global
17	vm53.in.tier2.hep.manchester.ac.uk	195.194.109.200		3403.76074219 kB/s	Manchester	United Kingdom	53.4663	-2.1342	21s	90h49m45s	✓	✗	Global
18	vm55.in.tier2.hep.manchester.ac.uk	195.194.109.199		3739.07421875 kB/s	Manchester	United Kingdom	53.4663	-2.1342	21s	90h49m45s	✓	✗	Global
19	ca-mey-frontier-8b6fd389cb.cern.ch	137.138.63.73		8150 kB/s	Geneva	Switzerland	46.2054	6.1459	22s	90h49m59s	✓	✗	Local Access Only
20	ca-mey-frontier-485577af8b.cern.ch	188.184.31.232		6422 kB/s	Geneva	Switzerland	46.2054	6.1459	26s	90h49m50s	✓	✗	Local Access Only
21	ca-mey-frontier-fac0d42b9c.cern.ch	188.184.28.244		0 kB/s	Geneva	Switzerland	46.2054	6.1459	29s	90h49m29s	✓	✗	Local Access Only
22	alien.spacescience.ro	85.120.46.17	172.18.0.1	4282.96289062 kB/s		Romania	45.9968	24.997	29s	90h49m58s	✓	✗	Global
23	js-17-28.jetstream-cloud.org	129.114.17.28	172.24.83.16	0.0 kB/s	Austin	United States	30.3773	-97.71	30s	90h49m51s	✓	✗	Global

Shoal V1.0.x

- redevelopment based on [Python 3](#) and fixing issues that came up with the old system
 - shoal-agent and shoal-client still work with Python 2 too
- large code improvements result in [much better performance](#)
- most properties [autoconfigured](#) now
- [more squid testing](#) for better matching and results
- new development will follow *Major.Minor.Micro* numbering

shoal-agent V1.0.x

- installs via [pip/pip3](#) and [post-install script](#) for basic configuration
- only needs very minimal configuration now
 - shoal-server to be used, admin email for contact in case of issues,...
 - all options have working defaults already set
 - no need to have easily available properties put into a config file
- [most properties auto-configured at runtime](#)
 - squid port, public IP, private IP, network interface, network interface link speed,
- does no longer reports access states to the server (Global, Private, Local Access Only)
- [add broadcast feature](#)
 - sends announcement broadcast in local network periodically to be discovered by a squid without the need to contact the shoal server
- tests squid access before reporting to server or sending broadcast
 - [only reports when squid works](#)
 - otherwise email to admin (once per 24h)

shoal-server V1.0.x [3]

- large code improvements and much better performance now
- tests each squid and **determines access level**
 - Private when access blocked or Global when access is possible
 - everything IP based
 - **does not rely on hostname anymore**
- uses DNS information to determine if a squid is in the same domain as the client
 - sets flag in the information sent to the client when the squid is in the same domain
 - private squids are only given out to clients in the same domain
 - **no need for a domain db anymore**
- ranks squids by distance and load for each request and gives out list to client
 - still uses IP to City mmdb
 - instead of using maxmind mmdb, use db-ip mmdb now
 - load and distance contribute as weight to the ranking (configurable)
- wpad file can also be retrieved

shoal-client V1.0.x

- installs via [pip/pip3](#)
 - rpm available but not in a repository currently
- [CERNVM has the new shoal-client preinstalled](#)
 - also had the old version preinstalled
- requests $n+5$ squids from the server, ordered by distance (and load)
 - 5 extra squids requested to be able to return n squids in case some do not work
- [all squids retrieved are tested for access and access time measured](#)
 - public and private IPs are tested for squids in the same domain as client
 - list reordered by access time and locality
 - squid received via broadcast is preferred one
- [local squids accessible on private IPs preferred](#), public IPs sorted by access time
 - also takes into account defaults configured within the system where the client runs
 - `/etc/cvmfs/default.local`
- reconfigures cvmfs and still has FRONTIER option

Shoal V1.0.x

List of Active Squids

26 active in the last 180 seconds

#	Hostname	Public IP	Private IP	Bytes Out	City	Country	Latitude	Longitude	Last Received	Alive	Clock Sync	Access Level
1	206-12-92-244.cloud.computecanada.ca	206.12.91.244	10.39.27.67	15734.0 kB/s	Victoria (Harris Green)	Canada	48.4244	-123.36	1s	60h47m3s	✓	Global
2	atlassq.uibk.ac.at	138.232.146.212	10.10.0.12	11.903 kB/s	Innsbruck (Hötting)	Austria	47.265	11.3429	1s	60h46m49s	✓	Private
3	squid.indiacms.res.in	144.16.111.37	192.168.2.37	0.543 kB/s	New Delhi	India	28.5866	77.2395	3s	12h50m44s	✓	Private
4	vm-115-146-93-82.rc.cloud.unimelb.edu.au	115.146.93.82		0.0 kB/s	Carlton	Australia	-37.802	144.959	4s	19h7m54s	✓	Global
5	elephant103.heprc.uvic.ca	206.12.154.103	10.1.3.22	0.0 kB/s	Victoria (Harris Green)	Canada	48.4244	-123.36	6s	60h47m7s	✓	Private
6	ca-mey-frontier-bigfin-c6a7fff898.cern.ch	188.184.98.18		37373.0 kB/s	Geneva	Switzerland	46.234	6.05295	6s	60h46m52s	✓	Private
7	ca-mey-frontier-8b6fd389cb.cern.ch	137.138.63.73		17819.0 kB/s	Meyrin	Switzerland	46.2296	6.0526	10s	60h47m4s	✓	Private
8	ca-mey-frontier-6d2f31843a.cern.ch	137.138.150.196		5585.0 kB/s	Meyrin	Switzerland	46.2296	6.0526	10s	60h47m5s	✓	Private
9	heptaur.ucy.ac.cy	194.42.36.20		0.811 kB/s	Nicosia	Cyprus	35.1856	33.3823	12s	60h46m48s	-143.3s	Private
10	js-17-28.jetstream-cloud.org	129.114.17.28	172.24.83.16	0.072 kB/s	Austin (North Burnet–Gateway)	United States	30.3912	-97.7218	16s	60h47m6s	✓	Private
11	ca-mey-frontier-angel-07c4f695ef.cern.ch	188.184.94.32		27607.0 kB/s	Geneva	Switzerland	46.234	6.05295	16s	60h46m51s	✓	Private
12	ca-mey-frontier-485577af8b.cern.ch	188.184.31.232		4363.0 kB/s	Geneva	Switzerland	46.234	6.05295	17s	60h46m46s	✓	Private
13	issaf.spacescience.ro	46.243.114.226	10.99.99.1	36.011 kB/s	Măgurele	Romania	45.1022	26.0373	18s	60h46m42s	✓	Private
14	js-17-31.jetstream-cloud.org	129.114.17.31	172.24.83.3	0.0 kB/s	Austin (North Burnet–Gateway)	United States	30.3912	-97.7218	19s	60h47m6s	✓	Private
15	vobox.spacescience.ro	85.120.46.18	172.20.0.18	1.27 kB/s	Măgurele	Romania	44.3614	26.0382	19s	60h46m42s	✓	Private
16	squid-grid.uaic.ro	85.122.31.70		0.0 kB/s	Iasi	Romania	47.1585	27.6014	21s	60h46m47s	✓	Private
17	marsq03.in2p3.fr	134.158.20.28		7.153 kB/s	Villeurbanne	France	45.7719	4.89017	22s	60h46m51s	✓	Private
18	marsq02.in2p3.fr	134.158.20.33		2.796 kB/s	Villeurbanne	France	45.7719	4.89017	22s	60h46m51s	✓	Private
19	alien.spacescience.ro	85.120.46.17	172.18.0.1	2147.514 kB/s	Măgurele	Romania	44.3614	26.0382	22s	60h46m40s	✓	Private
20	ca-mey-frontier-fac0d42b9c.cern.ch	188.184.28.244		0.0 kB/s	Geneva	Switzerland	46.234	6.05295	22s	60h46m45s	✓	Private
21	vm55.in.tier2.hep.manchester.ac.uk	195.194.109.199		5375.349 kB/s	Manchester	United Kingdom	53.4671	-2.23418	28s	60h47m1s	✓	Private
22	vm53.in.tier2.hep.manchester.ac.uk	195.194.109.200		5835.086 kB/s	Manchester	United Kingdom	53.4671	-2.23418	28s	60h47m1s	✓	Private
23	ca-mey-frontier-0097bb1c18.cern.ch	137.138.123.72		10243.0 kB/s	Meyrin	Switzerland	46.2296	6.0526	29s	60h46m47s	✓	Private
24	192.168.2.170	129.215.193.72	192.168.2.170	1.0 kB/s	Edinburgh	United Kingdom	55.9466	-3.20103	30s	13h15m21s	✓	Private
25	sampax02.if.usp.br	200.17.30.146	10.10.15.42	0.0 kB/s	São Paulo	Brazil	-23.5614	-46.7308	31s	60h46m43s	✓	Private
26	sampaproxy.if.usp.br	200.17.30.145	10.10.15.41	0.0 kB/s	São Paulo	Brazil	-23.5614	-46.7308	31s	60h46m43s	✓	Private

Outlook for future releases

- presentation
 - web view sorted by column of choice
 - add map with locations
- better load estimation of a squid
- look into switching to mysql for IP->City database
 - we can easily add new entries ourselves instead of waiting until a commercial one supports it
 - maybe in combination with online queries?
- everything already in [GitHub](#) [4]
 - please open an issue there
 - bug reports
 - features you would like to see in the future
 - also Readme-files available for installations
 - *shoal-agent needs to have setup script executed*
 - *shoal-client needs to have config file created*

Summary

- Shoal system **successfully in production for nearly 10 years**
 - first only by us but now many squids are registered from all over the world
- major **rewrite and update of all components supporting Python3** now
 - shoal-server, shoal-client, and shoal-agent
- all code available in **GitHub**
- removed manual configurations as much as possible
- **better testing of squids** in all three parts of shoal
- pre-installed in **CERNVM**
- interesting **for any environment**
 - make new local squids instantly available to a local cluster without reconfiguration
 - usage of different clouds by a single institution
 - contribution of small groups without much local resources to an experiment
 - usage of multiple clouds within an experiment

Summary

- Shoal system **successfully in production for nearly 10 years**
 - first only by us but now many squids are registered from all over the world
- major **rewrite and update of all components supporting Python3** now
 - shoal-server, shoal-client, and shoal-agent
- all code available in **GitHub**
- removed manual configurations as much as possible
- **better testing of squids** in all three parts of shoal
- pre-installed in **CERNVM**
- interesting **for any environment**
 - make new local squids instantly available to a local cluster without reconfiguration
 - usage of different clouds by a single institution
 - contribution of small groups without much local resources to an experiment
 - usage of multiple clouds within an experiment

It would be very helpful if as many squids as possible could have the shoal-agent installed.

- [1] <http://heprc.phys.uvic.ca/>
- [2] <https://link.springer.com/epdf/10.1007/s41781-020-0036-1>
- [3] <http://shoal.heprc.uvic.ca/>
- [4] <https://github.com/hep-gc/shoal>

Thank you!

- [1] <http://heprc.phys.uvic.ca/>
- [2] <https://link.springer.com/epdf/10.1007/s41781-020-0036-1>
- [3] <http://shoal.heprc.uvic.ca/>
- [4] <https://github.com/hep-gc/shoal>