



Improving the OSG-LHC Network Metrics User Interface

Tommy Shearer



IRIS-HEP Project Information & Background

- I started the IRIS-HEP Fellowship under the proposal **Improving the User Interface to OSG-LHC Network Metrics**
 - Initially was only going to work on developing **dashboards**
 - The **dashboards** planned were going to solve some specific use-cases that users of the system may have
- As I worked on the project, there were other opportunities that allowed users to easily navigate the **dashboards**
 - This primarily was the [OSG Toolkit Info Page](#)



Project Plans

- When first introduced to the IRIS-HEP fellowship, the **data** being stored on Elasticsearch at the MidWest Tier-2 really stuck out
 - Have a lot of previous experience with ELK at the ATLAS Great Lakes Tier-2
- All of the **dashboards** already created were only showing very broad data
- Therefore, I wanted to create **dashboards** that would allow users to answer a question on one page instead of grabbing a lot of general information from a few



Project Progress

- Some very rough **dashboards** were made, but mostly introducing myself to various **communities** to see if there were any specific problems they had
 - US ATLAS Facilities
 - OSG Facilities
 - SAND
- The group of people I worked with, including Shawn McKee, had weekly meetings as well
 - Got constant feedback on the work that I was doing initially
- Was introduced to the **[OSG Toolkit Info Page](#)**

Initial OSG Toolkit Info Page

- Was pretty outdated, some basic ideas of adding my **dashboards** and using it as a central location

The perfSONAR Toolkit Information Page

Open Science Grid

WLCG
Worldwide LHC Computing Grid

Select toolkit: Submit

Documentation OSG Network Pipeline OSG Network Services Analytics and Dashboards

Your selected perfSONAR Toolkit is:

Web links for
[This toolkit's web interface](#)
[Testing instructions for this toolkit \(JSON\)](#)
[This toolkit's settings and status](#)

You don't seem to have an x.509 credential.
Some items are not accessible without such a credential
and are removed from this page and its menus.

SAND iris hep NSF
Institute for Research & Innovation
in Software for High Energy Physics



More Progress on Dashboards

- Came across the use-case that I had focused on the most, helping **site-administrators** solve problems that they are having with their site, **even if they don't know much about the technicalities of the data**
- Now with a goal in mind, real development of the **dashboards** could be made



New Perspective on Dashboards

- Starting laying out a foundation that could allow access to all types of **data** based on the users' needs
- When creating the first **dashboards**, the ways I was showing data weren't very useful, and after some feedback, a more clear approach was discovered
 - Make them very **user-friendly**
 - Someone who doesn't know what they are looking for can be **directed** to information
 - Get more specific as the user progresses down the page
- The **dashboards** really started coming along
- <https://atlas-kibana.mwt2.org/s/networking/app/kibana#/dashboard/57e3c0b0-d7f3-11ea-9344-2da4788d78a4>

Examples of a Dashboard

Going to use the AGLT2 as an example site

(src_site:AGLT2 AND dest_site:*)
OR (src_site:* AND dest_site:AGLT2)

Full screen Share Clone Edit

KQL Show dates

To RETURN to the Landing Page, please click [here](#)

Links to Other Informational Dashboards

To access any of the other types of data, please select it from the following list

[General](#) | [Throughput](#) | [Trace](#) | [Packetloss](#) | [Retransmits](#) | [One-Way Delay](#)

To change the site information that is selected, please use the following control tabs to do so. If site information wants to be removed, please select the 'X' on the filter to eliminate the filter.

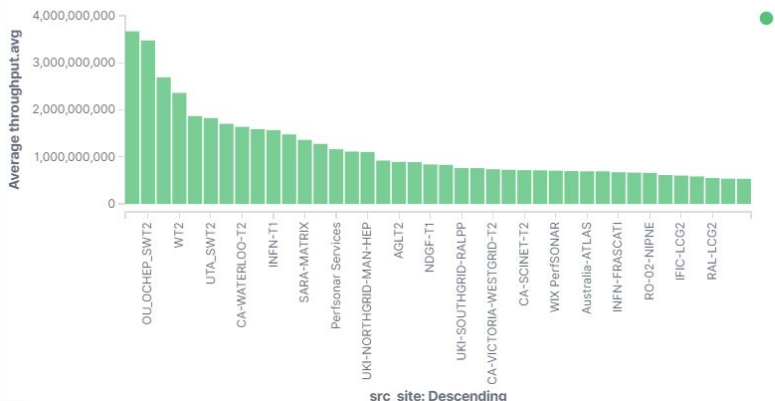
Source and Destination Site Selection

Select Source Site

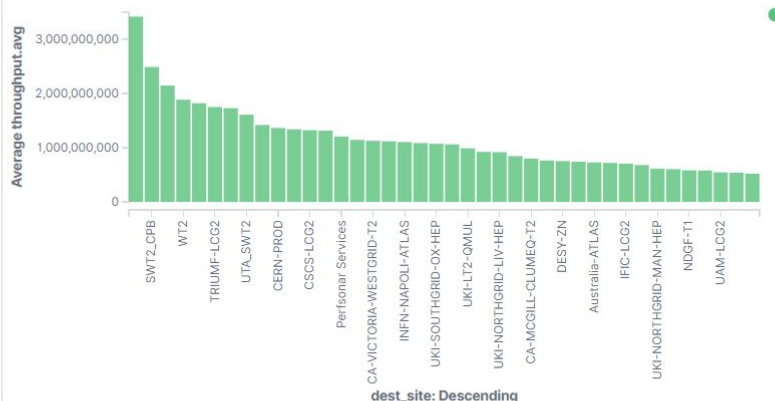
Select Destination Site

General Averages between Sites


Top 40 Largest Throughput Values for Source Sites




Top 40 Largest Throughput Values for Destination Sites



Top 40 Smallest Throughput Value for Source Sites



Top 40 Smallest Throughput Values for Destination Sites



Breakdown of the Site-Site Averages

Source and Destination Site Pair Throughput Average Breakdown ⓘ

src_site: Descending ⓘ	dest_site: Descending ⓘ	Max throughput.avg ⓘ	Min throughput.avg ⓘ	Average throughput.avg ⓘ
AGLT2	AGLT2	5,510,589,696.414	5,158,147,085.567	5,302,737,564.415
AGLT2	MWT2	4,663,621,462.775	1,946,292,893.621	3,418,625,246.391
AGLT2	SWT2_CPB	4,167,967,613	809,365,786.667	2,488,666,699.833
AGLT2	BNL-ATLAS	3,557,654,917.387	932,100,353.415	2,145,888,730.143
AGLT2	CA-VICTORIA-WESTGRID-T2	3,279,111,445.746	214,113,395.617	1,127,075,953.408
AGLT2	WT2	3,260,342,364.085	509,840,756.038	1,885,091,560.061
AGLT2	SARA-MATRIX	3,176,888,310.371	467,176,705.474	1,821,500,120.725
AGLT2	TRIUMF-LCG2	2,991,391,281.06	546,143,915.926	1,751,531,960.172
AGLT2	CERN-PROD	2,919,588,769.634	159,593,926.333	1,364,002,268.283
AGLT2	INFN-T1	2,751,466,813.01	386,453,257.668	1,420,101,566.119

Looking at Host-Host and IP-IP

Source and Destination Host Pair Throughput Average Breakdown ⓘ

src_host: Descending ⌵	dest_host: Descending ⌵	Average throughput.avg ▾
mwt2-ps02.campuscluster.illinois.edu	psmsu02.aglt2.org	5,192,013,872.031
psmsu02.aglt2.org	psum02.aglt2.org	5,168,525,068.973
iut2-net2.iu.edu	psmsu02.aglt2.org	4,705,769,162.423
psmsu02.aglt2.org	mwt2-ps02.campuscluster.illinois.edu	4,427,617,389.344
psnr-farm10.slac.stanford.edu	psmsu02.aglt2.org	4,230,923,090.393
psmsu02.aglt2.org	uct2-net2.mwt2.org	4,170,464,889.17
lhcmn.bnl.gov	psmsu02.aglt2.org	4,126,764,267.769
uct2-net2.mwt2.org	psmsu02.aglt2.org	3,980,045,505.354
psmsu02.aglt2.org	iut2-net2.iu.edu	3,789,952,376.331
ps2.ochep.ou.edu	psmsu02.aglt2.org	3,475,678,404.109

Source and Destination IP Pair Throughput Average Breakdown ⓘ

src: Descending ⌵	dest: Descending ⌵	Average throughput.avg ▾
72.36.96.9	192.41.236.32	5,509,984,216.695
2001:48a8:68f7:8001:192:41:236:32	2001:48a8:68f7:1:192:41:230:60	5,178,903,052.379
192.41.236.32	192.41.230.60	5,158,147,085.567
2620:0:e01:4800::18	2001:48a8:68f7:8001:192:41:236:32	4,874,043,527.367
149.165.225.224	192.41.236.32	4,705,769,162.423
192.170.227.162	192.41.236.32	4,664,024,145.496
192.41.236.32	72.36.96.9	4,663,621,462.775
192.41.236.32	192.170.227.162	4,550,524,656.073
2620:0:210:1::23	2001:48a8:68f7:8001:192:41:236:32	4,392,797,169.938
134.79.118.72	192.41.236.32	4,230,923,090.393

Previews of Related Data

Link to the Packetloss Dashboard

To access more information about packetloss, please click [here](#)

Preview of the Source and Destination Site Pair Packetloss Average Breakdown

src_site: Descending ▾	dest_site: Descending ▾	Max packet_loss.avg ▾	Min packet_loss.avg ▾	Average packet_loss.avg ▾
CA-SCINET-T2	AGLT2	1	0.002	0.501
CA-VICTORIA-WESTGRID-T2	AGLT2	1	0	0.251
CIT_CMS_T2	AGLT2	1	1	1
FZK-LCG2	AGLT2	1	0.001	0.264
IFIC-LCG2	AGLT2	1	0.001	0.27
SARA-MATRIX	AGLT2	1	0	0.208
UKI-NORTHGRID-MAN-HEP	AGLT2	1	0.001	0.5
UKI-SOUTHGRID-BRIS-HEP	AGLT2	1	0.001	0.501
praguelcg2	AGLT2	1	0	0.251
RO-07-NIPNE	AGLT2	0.998	0	0.251

Link to the Retransmits Dashboard

To access more information about retransmits, please click [here](#)

Preview of the Source and Destination Site Pair Number of Retransmits Average Breakdown

src_site: Descending ▾	dest_site: Descending ▾	Max retransmits.avg ▾	Min retransmits.avg ▾	Average retransmits.avg ▾
AGLT2	UKI-SOUTHGRID-OX-HEP	11,222.549	6.478	1,628.811
AGLT2	UKI-SCOTGRID-ECDF	7,393.56	102.174	2,699.71
AGLT2	RO-07-NIPNE	3,357.415	232.209	1,171.031
AGLT2	FZK-LCG2	3,310.874	37.527	878.765
AGLT2	GRIF	2,469.477	15.783	660.721
AGLT2	UKI-NORTHGRID-LANCS-HEP	2,438.464	180.063	868.326
AGLT2	DESY-ZN	2,035.577	93.204	846.828
AGLT2	UKI-SCOTGRID-GLASGOW	1,764.788	480.597	1,027.855
AGLT2	IFIC-LCG2	1,596.595	127.32	577.183
AGLT2	🔍 RAL-LCG2	1,556.152	139.02	563.739



Working on the OSG Toolkit Info Page in Parallel

- You may have noticed, **users** would have to input a string of text in order to look at anything related to their site
- I decided to not only give the **OSG Toolkit Info Page** a redesign, but to add useful features as well

Graphical Redesign

Select toolkit: Submit

Documentation OSG Network Pipeline **OSG Network Services** Analytics and Dashboards

perSONAR Ingest Ra OSG/WLCG Ma WLCG Grafana Network Dashboards

OSG/WLCG pSCor OSG/SAND/WLCG ELK Network Dashboards

OSG perfSONAR Infrastructure details on a Kibana Dashboard

OSG Overview of Packet Loss Kibana Dashboard

The SAND Project

focused on networking analytics for OSG and WLCG

The MEPHI Traceroute Visualizer

using our OSG/WLCG traceroute measurements

Web links for
[This toolkit's web in](#)
[Testing instructions f](#)
[This toolkit's setting](#)

You don't seem to have an account. Some items are not accessible from this account.

This is the OSG/WLCG perSONAR Infrastructure dashboard"

Institute for Research & Innovation
in Software for High Energy Physics

[Contact us about this webpage](#)

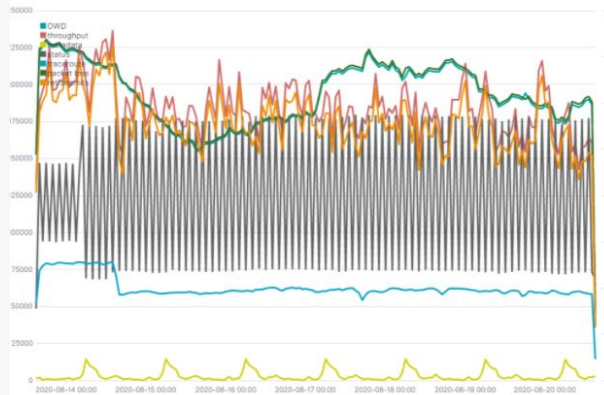
New Graphics

OSG Network Pipelines

Note, Orange Labels Require A SSL Certificate

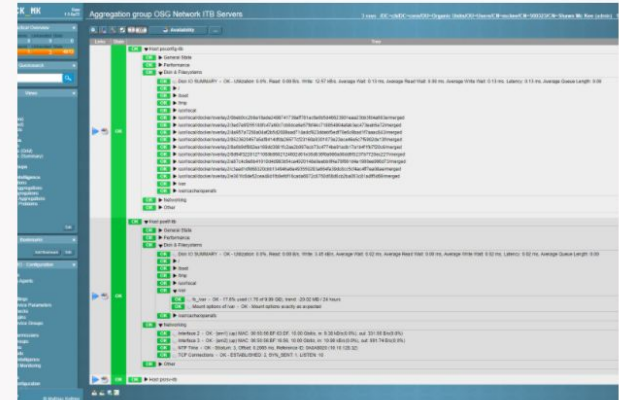
perfSONAR Ingest Rate

This is the Kibana interface to the University of Chicago ATLAS Analytics platform which hosts our network metrics (as well as lots of other analytics data)



Status of ITB Services

This is the OSG/WLCG perfSONAR Infrastructure Test Bed (ITB) service monitoring host, based upon the Experiments Testing Framework (ETF) and Check_MK



New Features



Main Page

Documentation

OSG Network Pipelines

OSG Network Services

Analytics and Dashboards

Toolkit Specific Page

Please Select A Category Based on The Question You Have

[Do You Suspect General Networking Problems with Your Site?](#)

[Are you Interested in your Site's Overall Network Performance?](#)

[Do You Suspect a lot of Packetloss when Transmitting Data to Another Site?](#)

[Are you Interested in the Path that Data Takes from your Site to Another?](#)

[Do You Suspect Slow Connections between Your Site and Another?](#)

[Are you Seeing a lot of Throughput Issues or Retransmitted Data for your Site?](#)



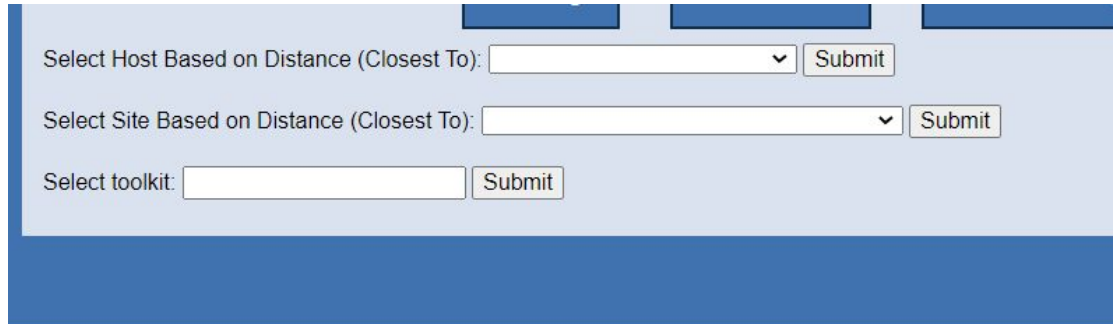
Continuing Work

- I am planning on continuing working with the SAND group to finish some user-friendly features, both for the OSG Toolkit Info and the dashboard networks
- Working on a system that will **query elasticsearch** for hosts, geoiip latitudes and longitudes, and **site** names.
- Will then compare the **user's** latitude and longitude to all of the hosts'
- Will then give the **closest** hosts and sites to the user, creating ease of access for **users who may not know which hosts they should look for.**



With Test Data

This is being developed on the developer OSG Toolkit Info Page that I developed to test new features before pushing them to the production copy.



Select Host Based on Distance (Closest To):

Select Site Based on Distance (Closest To):

Select toolkit:

User IP is: 96.66.30.65

User Latitude is: 42.2709

User Longitude is: -83.7268

Host_name: psonar3.fnal.gov is 234.13111700419 miles away at site Fermilab

Host_name: atlas-owamp.bu.edu is 644.62449360495 miles away at site Boston University

Host_name: heplnx130.pp.rl.ac.uk is 3731.3321786185 miles away at site

Host_name: perfsonar01.datagrid.cea.fr is 3957.1423147232 miles away at site GRIF-IRFU

Host_name: psl-ifae.pic.es is 4209.2889348162 miles away at site pic

Host_name: perfsonar01.lcg.cscs.ch is 4326.0692448523 miles away at site CSCS

Host_name: perfsonar2.mi.infn.it is 4355.9778823942 miles away at site

Host_name: psonar.cis.gov.pl is 4533.303300525 miles away at site National Centre for Nuclear Research

Host_name: ps01.cat.cbpf.br is 5190.6631215573 miles away at site Centro Brasileiro de Pesquisas Físicas

You have not selected a toolkit for customized links.
Please use the above "Select toolkit" drop-down
to pick a specific toolkit.




User's Geoip Script

```
<div>
<?php
    $ip_addr = $_SERVER['REMOTE_ADDR'];
    $geoplugin = unserialize( file_get_contents('http://www.geoplugin.net/php.gp?ip='.$ip_addr) );

    if ( is_numeric($geoplugin['geoplugin_latitude']) && is_numeric($geoplugin['geoplugin_longitude']) ) {

        $lat = $geoplugin['geoplugin_latitude'];
        $long = $geoplugin['geoplugin_longitude'];
    }
    ?>
</div>
```

This piece of code was used to grab the client's IP and grab a geoIP latitude and longitude so that they can be used to calculate the distance between the user that is connected and the hosts.



Distance and Organization

This piece of code is then used to define a distance function, finding the distance between the client and the list of hosts, and then calculate and sort based on the closest hosts

```
<div>
<?php
function distance($lat1, $lon1, $lat2, $lon2, $unit) {
    if (($lat1 == $lat2) && ($lon1 == $lon2)) {
        return 0;
    }
    else {
        $theta = $lon1 - $lon2;
        $dist = sin(deg2rad($lat1)) * sin(deg2rad($lat2)) + cos(deg2rad($lat1)) * cos(deg2rad($lat2)) * cos(deg2rad($theta));
        $dist = acos($dist);
        $dist = rad2deg($dist);
        $miles = $dist * 60 * 1.1515;
        $unit = strtoupper($unit);

        if ($unit == "K") {
            return ($miles * 1.609344);
        } else if ($unit == "N") {
            return ($miles * 0.8684);
        } else {
            return $miles;
        }
    }
}
?>
</div>
<div>
<?php
$csv = array_map('str_getcsv', file('test_query_data.csv'));
foreach($csv as $location => $data)
{
    $calculated_distance = distance($lat,$long,$data[0],$data[1],"M");
    $distance_array[$location][distance] = $calculated_distance;
    $distance_array[$location][host_name] = $data[2];
    $distance_array[$location][site_name] = $data[3];
}
$sorting = array_column($distance_array,'distance');
array_multisort($sorting, SORT_ASC, $distance_array);
?>
</div>
```



Dynamic Dropdown Menus

This piece of code then uses the hosts sorted by distance, and makes a dynamic drop down menu

```
<div class="form">
<form action="/developertoolkitinfo/perfsonartoolkit.php" method="get">
<label for="host" style="color:black">Select Host Based on Distance (Closest To): </label>
<select id="host" name="host">
<option value = ''></option>
<?php
foreach($distance_array as $array) { ?>
<option value= "<?php echo $array['host_name']; ?>"><?php echo $array['host_name']; ?></option>
<?php
} ?>
</select>
<input type="submit" value='Submit'>
</form>
</div>

<div class="form">
<form action="/developertoolkitinfo/perfsonartoolkit.php" method="get">
<label for="site" style="color:black">Select Site Based on Distance (Closest To): </label>
<select id="site" name="site">
<option value = ''></option>
<?php
foreach($distance_array as $array) {
if (!empty($array['site_name'])) { ?>
<option value= "<?php echo $array['site_name']; ?>"><?php echo $array['site_name']; ?></option>
<?php
}
} ?>
</select>
<input type="submit" value='Submit'>
</form>
</div>
```



Moving Forward

- Make some aesthetic and user-appealing changes to the new features for the [OSG Toolkit Info Page](#)
- Finish developing the Python script that will query **Elasticsearch** for the real **data**, and export as a .CSV
- Set up a cronjob to run the **script** every couple of days to **keep updated**
- Make more **dashboards** for more use-cases, and continue to perfect the already created ones
 - Requires user feedback on current designs
 - Any desired features?
 - Any other use-cases they'd like to see?



Thank you! Questions?