Adding Network Measuring to Data Federation Caches

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

CENTER FOR
HIGH THROUGHPUT
COMPUTING

iris
hep

Ziyang Ye (UW–Madison)
Mentor: Brian Lin (UW–Madison)

Project Goal

- set up two perfSONAR instances on two data caches.
- schedule network measuring tests between them.

Background: Data Federation

- it is provided by OSG.
- LHC experiments store data on origin servers.
- data is streamed from data origins to distributed data caches all over the world.

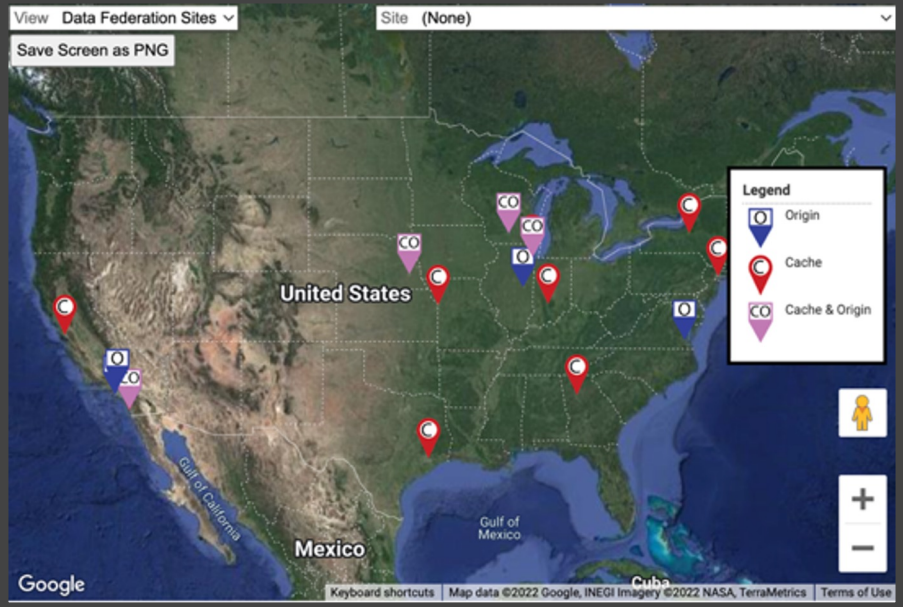

Background: perfSONAR

- It is a performance Service-Oriented Network monitoring ARchitecture
- perfSONAR tests all nodes on the network path and looks for areas of low performance
- user can set different test types based on test purposes. (Latency tests, Throughput tests, etc)

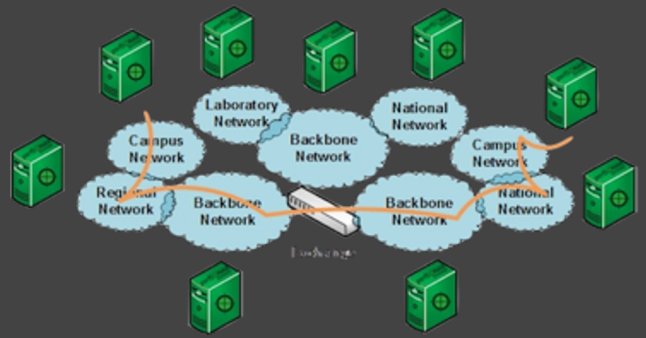

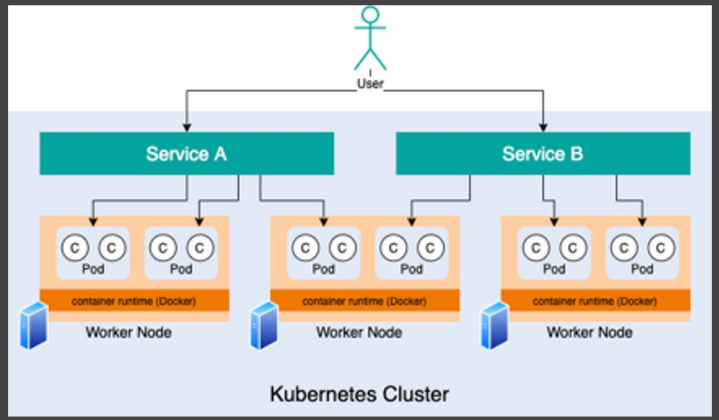

Set up perfSONAR

- build a perfSONAR image that pulls test configs from psconfig web admin
- deploy perfSONAR on docker container using the image

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS	NAMES
9de9a70b186d	perfsonar/testpoint	"/bin/sh -c '/usr/bi…"	About a minute ago	Up About a minute		test_perfsonar_testpoint

```
=== pScheduler Agent ===
Added remote configuration http://psconfig.opensciencegrid.org/pub/config/osg-xache
=== pScheduler Agent ===
[
   {
      "url" : "http://psconfig.opensciencegrid.org/pub/config/osg-xache"
   }
]
```

Set up perfSONAR on Tiger cluster

- deploy the container alongside the data cache deployment on Tiger cluster.
- set up public ip
- set up ports needed for perfSONAR to communicate with other nodes

```
Name:                   stash-cache-perfsonar
Namespace:              osgdev
CreationTimestamp:      Tue, 30 Aug 2022 12:22:03 -0500
Labels:                 app=stash-cache
                        fluxcd.io/sync-gc-mark=sha256.aS6WV9pdAEu6MRRZldNoBI3amV1XAdoAtnMkb1zbwDE
Annotations:            deployment.kubernetes.io/revision: 1
                        fluxcd.io/sync-checksum: 4375b56732986344710472fc574ecd59770ab572
Selector:               app=stash-cache
Replicas:               1 desired | 1 updated | 1 total | 1 available | 0 unavailable
StrategyType:           Recreate
MinReadySeconds:        0
Pod Template:
  Labels:  app=stash-cache
  Init Containers:
   chown-xrootd-log-dir:
    Image:      hub.opensciencegrid.org/library/alpine:3
    Port:       <none>
    Host Port:  <none>
    Command:
      chown
      -R
      10940:10940
      /var/log/xrootd
    Environment:  <none>
    Mounts:
      /var/log/xrootd from stash-cache-xrootd-logs (rw)
  Containers:
   perfsonar-container:
    Image:        zye298/perfsonar-container:v5.0
    Ports:        443/TCP, 5001/TCP, 5201/TCP, 5000/TCP, 5101/TCP, 8090/TCP, 80/TCP
    Host Ports:   443/TCP, 5001/TCP, 5201/TCP, 5000/TCP, 5101/TCP, 8090/TCP, 80/TCP
    Environment:  <none>
    Mounts:       <none>
   stash-cache:
    Image:       hub.opensciencegrid.org/opensciencegrid/stash-cache:3.6-development
    Ports:       1094/TCP, 8000/TCP, 8443/TCP
    Host Ports:  1094/TCP, 8000/TCP, 8443/TCP
    Limits:
      cpu:     4
      memory:  16Gi
    Requests:
      cpu:     1
      memory:  8Gi
    Environment:
      XC_RESOURCENAME:  RESOURCE_NAME_IN_TOPOLOGY
      CACHE_FQDN:       FQDN.IN.TOPOLOGY
    Mounts:
      /etc/grid-security/hostcert.pem from stash-cache-cert (rw,path="hostcert.pem")
      /etc/grid-security/hostkey.pem from stash-cache-cert (rw,path="hostkey.pem")
      /var/log/xrootd from stash-cache-xrootd-logs (rw)
      /xcache from stash-cache-cache (rw)
  Volumes:
   stash-cache-cache:
    Type:       PersistentVolumeClaim (a reference to a PersistentVolumeClaim in the same namespace)
    ClaimName:  stash-cache-cache-perfsonar
    ReadOnly:   false
   stash-cache-xrootd-logs:
    Type:       PersistentVolumeClaim (a reference to a PersistentVolumeClaim in the same namespace)
    ClaimName:  stash-cache-xrootd-logs-perfsonar
    ReadOnly:   false
   stash-cache-cert:
    Type:        Secret (a volume populated by a Secret)
    SecretName:  stash-cache-cert
    Optional:    false
```

Set up perfSONAR on River cluster

- deploy the container alongside the data cache deployment on River cluster.
- set up public ip
- set up ports needed for perfSONAR to communicate with other nodes

```
Name:               stash-cache-perfsonar
Namespace:          osg-dev
CreationTimestamp:  Tue, 19 Jul 2022 16:29:59 -0500
Labels:             app=stash-cache
                    kustomize.toolkit.fluxcd.io/name=osg-dev
                    kustomize.toolkit.fluxcd.io/namespace=osg-dev
Annotations:        deployment.kubernetes.io/revision: 5
Selector:           app=stash-cache
Replicas:           1 desired | 1 updated | 1 total | 1 available | 0 unavailable
StrategyType:       Recreate
MinReadySeconds:    0
Pod Template:
  Labels:  app=stash-cache
  Init Containers:
   chown-xrootd-log-dir:
    Image:      hub.opensciencegrid.org/library/alpine:3
    Port:       <none>
    Host Port:  <none>
    Command:
      chown
      -R
      10940:10940
      /var/log/xrootd
    Environment:  <none>
    Mounts:
      /var/log/xrootd from stash-cache-xrootd-logs (rw)
  Containers:
   perfsonar-container:
    Image:        zye298/perfsonar-container:v5.0
    Ports:        443/TCP, 5001/TCP, 5201/TCP, 5000/TCP, 5101/TCP, 8090/TCP, 80/TCP
    Host Ports:   443/TCP, 5001/TCP, 5201/TCP, 5000/TCP, 5101/TCP, 8090/TCP, 80/TCP
    Environment:  <none>
    Mounts:       <none>
   stash-cache:
    Image:       hub.opensciencegrid.org/opensciencegrid/stash-cache:3.6-development
    Ports:       1094/TCP, 8000/TCP, 8443/TCP
    Host Ports:  1094/TCP, 8000/TCP, 8443/TCP
    Limits:
      cpu:     4
      memory:  16Gi
    Requests:
      cpu:     1
      memory:  8Gi
    Environment:
      XC_RESOURCENAME:  RESOURCE_NAME_IN_TOPOLOGY
      CACHE_FQDN:       FQDN.IN.TOPOLOGY
    Mounts:
      /etc/grid-security/hostcert.pem from stash-cache-cert (rw,path="hostcert.pem")
      /etc/grid-security/hostkey.pem from stash-cache-cert (rw,path="hostkey.pem")
      /var/log/xrootd from stash-cache-xrootd-logs (rw)
      /xcache from stash-cache-cache (rw)
  Volumes:
   stash-cache-cache:
    Type:       PersistentVolumeClaim (a reference to a PersistentVolumeClaim in the same namespace)
    ClaimName:  stash-cache-cache-perfsonar
    ReadOnly:   false
   stash-cache-xrootd-logs:
    Type:       PersistentVolumeClaim (a reference to a PersistentVolumeClaim in the same namespace)
    ClaimName:  stash-cache-xrootd-logs-perfsonar
    ReadOnly:   false
   stash-cache-cert:
    Type:        Secret (a volume populated by a Secret)
    SecretName:  stash-cache-cert
    Optional:    false
```

Schedule network measuring tests

- set up test configs on psconfig web admin at https://psconfig.opensciencegrid.org/

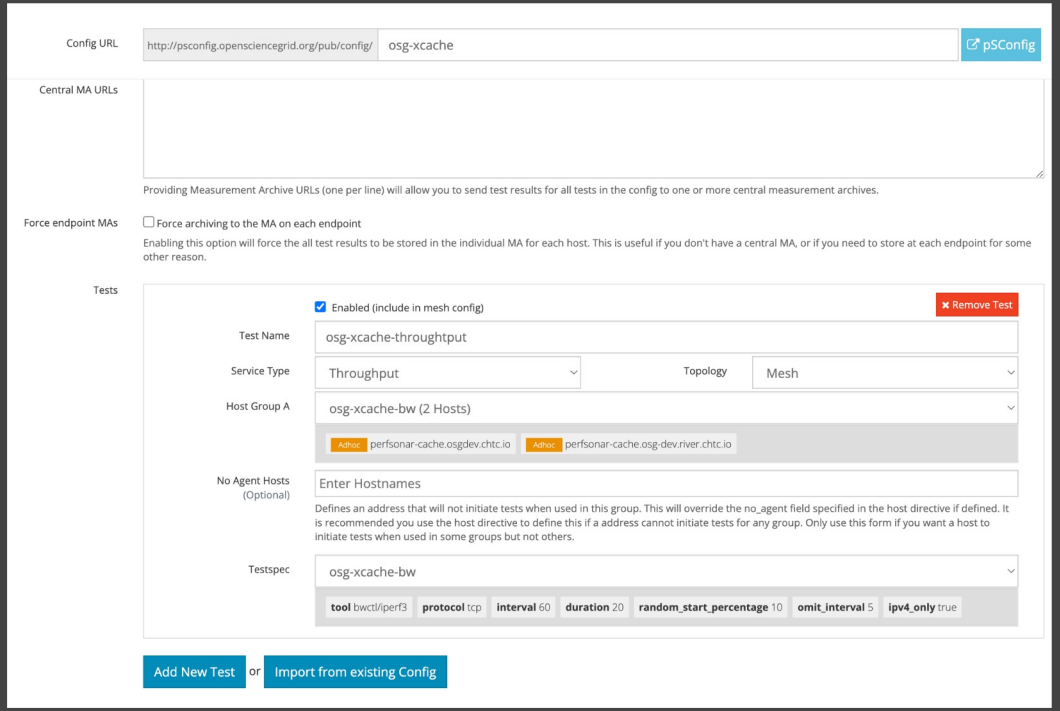

Major issues countered

- container does not run properly in kubernetes cluster (stuck in some failing states)
- container cannot receive test configs from web admin (container network fails, wrong ports opened)
- upstream perfsonar image does not provide full services for network tests

- perfSONAR on Tiger: has been stuck on pod initializing state.

```
Containers:
  perfsonar-container:
    Container ID:
    Image:          zye298/perfsonar-container:v4.0
    Image ID:
    Ports:          443/TCP, 5001/TCP, 5201/TCP, 5000/TCP, 5101/TCP, 8090/TCP, 80/TCP
    Host Ports:     443/TCP, 5001/TCP, 5201/TCP, 5000/TCP, 5101/TCP, 8090/TCP, 80/TCP
    State:          Waiting
      Reason:       PodInitializing
```

- perfSONAR on River: multiple services are not running.

```
[root@stash-cache-perfsonar-fdbfc597b-pqhmm perfsonar]# psconfig pscheduler-tasks
Unable to open /var/log/perfsonar/psconfig-pscheduler-agent.log: No such file or directory
[root@stash-cache-perfsonar-fdbfc597b-pqhmm perfsonar]# pscheduler debug on scheduler
Warning: scheduler does not appear to be running.
[root@stash-cache-perfsonar-fdbfc597b-pqhmm perfsonar]#
```

Schedule network measuring tests

- perfSONAR instance pulls test config from web admin and schedule the tests based on the config
- possible future work: a Central Measurement Archive to save test output is needed.

```
[root@stash-cache-perfsonar-69899966c7-r4ndf perfsonar]# psconfig pscheduler-tasks
{
   "tasks" : [
      {
         "archives" : [],
         "test" : {
            "spec" : {
               "source" : "perfsonar-cache.osgdev.chtc.io",
               "dest" : "perfsonar-cache.osg-dev.river.chtc.io",
               "duration" : "PT20S",
               "schema" : 1,
               "omit" : "PT5S",
               "ip-version" : 4
            },
            "type" : "throughput"
         },
         "reference" : {
            "psconfig" : {
               "created-by" : {
                  "user-agent" : "psconfig-pscheduler-agent",
                  "uuid" : "DE16479A-28A3-11ED-AD90-B3BD0D20D889"
               }
            }
         },
         "tools" : [
            "bwctliperf3",
            "iperf3"
         ],
         "schedule" : {
            "until" : "2022-08-31T23:18:56Z",
            "sliprand" : true,
            "repeat" : "PT1M",
            "slip" : "PT30S"
         }
      }
   ]
}
```

Thanks!

This material is based upon work supported by the National Science Foundation under Grant Nos. 1836650. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.