


# Network and Transfer WG Metrics Area Meeting

Shawn McKee, Marian Babik

Network and Transfer Metrics WG Meeting  
18<sup>h</sup> February 2015



# Outline

- Next Meetings
  - 18 March, 8 Apr, 6 May, 3 June (all at 4pm CEST)

## Today

- perfSONAR Status and Plans
- Datastore demo
- Integration
  - Experiment's interface to perfSONAR
  - FTS Performance
- AOB



# perfSONAR Status

# perfSONAR Status

- Update campaign finalized – deadline for sites was Monday – almost all instances now run 3.4.1
- Announced production instance of the infrastructure monitoring
  - <http://pfomd.grid.iu.edu>
- Pilot instances of datastore and dashboard
  - <http://psds.grid.iu.edu> (quick demo today)
  - <http://psmad.grid.iu.edu> (connects to psds)
- Many sites updated, but didn't re-configure, pfomd helps to identify issues, but follow up is needed
- Re-configuration campaign was run in LHCOPN and LHCONE bringing all instances to correct configuration
  - All LHCOPN instances OK
  - All LHCONE instances OK (except GEANT - working on update)

# perfSONAR Status

- For LHCOPN/LHCONE started to investigate if sonars are consistently delivering metrics
  - Added new checks to OMD to query freshness of local MA (esmond) and identify missing links
- Analysis still on-going, but we have already found a bug in perfSONAR's regular testing
  - Several cases found when some of the regular tests stop and never recover – it was be linked to improper process clean-up by the regular testing daemon
- Other issues were reported based on our experience with 3.4.1 - fixes provided in 3.4.2 that should be out soon and deployed via auto-update
- Very good experience with mesh configuration interface
  - IPv6 testing mesh was setup by Duncan
  - Moved traceroute tests from latency to bandwidth nodes

# Mesh Configurations

- Current
- All inter-cloud tests are disabled
- Latency
  - Full mesh within each cloud – btw LT nodes, continuous (10Hz) [1]
- Bandwidth
  - Full mesh within each cloud – btw BW nodes, every 6hrs (duration: 30s)
  - All WLCG – full mesh, btw BW nodes, once a week
- Traceroutes
  - All WLCG – btw BW nodes, hourly
  - LHCONE, USCMS, USATLAS, GOC – btw BW nodes, every 20mins

# Mesh Configurations

- Proposed
- Keep all inter-cloud tests disabled
- Latency
  - Keep Full mesh within each cloud – btw LT nodes, continuous (10Hz) [1]
  - Add Top-k WLCG – full mesh btw. top 50 sites – ramp up to all sites once this is working (starting with LHCOPN and LHCONE)
  - In case specific links need to be investigated we can do it in parallel
- Bandwidth
  - Keep Full mesh within each cloud – btw BW nodes, every 6hrs (duration: 30s)
  - Disable All WLCG – full mesh, once a week
- Traceroutes
  - Keep All WLCG – btw BW nodes, hourly
  - Keep LHCONE, USCMS, USATLAS, GOC – btw BW nodes, every 20mins



# Datastore Demo



# Datastore

- esmond – Postgress + Cassandra
  - populated by RSV probes
- REST API available – python/perl libs
  - curl  
"http://archive.example.net/esmond/perfsonar/archive/fce0483e51de49aaa7fcf8884d053134/histogram-owdelay/base?time-range=86400"
- Data organized in events
  - packet-trace
  - histogram-owdelay – one way delays over time period
  - ntp-delay – round trip delay time to NTP server
  - packet-loss-rate – number of packets lost/packets sent
  - packet-count-sent – packets sent
  - packet-count-lost – packets lost
  - packet-retransmits – packets retransmitted for a transfer using TCP
  - throughput – observer amount of data sent over period of time
  - failures – record of test failures

# Structure

```
[
  {
    "source": "10.1.1.1",
    "destination": "10.1.1.2",
    "event-types": [
      {
        "base-uri": "/esmond/perfsonar/archive/f6b732e9f351487a96126f0c25e5e546/packet-retransmits/base",
        "event-type": "packet-retransmits",
        "summaries": [
          ],
        "time-updated": 1397482734
      },
      {
        "base-uri": "/esmond/perfsonar/archive/f6b732e9f351487a96126f0c25e5e546/throughput/base",
        "event-type": "throughput",
        "summaries": [
          {
            "summary-type": "average",
            "summary-window": "86400",
            "time-updated": 1397482735,
            "uri": "/esmond/perfsonar/archive/f6b732e9f351487a96126f0c25e5e546/throughput/averages/86400"
          }
        ],
        "time-updated": 1397482735
      },
    ],
    "input-source": "host1.example.net",
    "input-destination": "host2.example.net",
    "ip-transport-protocol": "tcp",
    "measurement-agent": "10.1.1.1",
    "metadata-key": "f6b732e9f351487a96126f0c25e5e546",
    "subject-type": "point-to-point",
    "time-duration": "20",
    "time-duration": "14400",
    "tool-name": "bwctl/iperf3",
    "uri": "/esmond/perfsonar/archive/f6b732e9f351487a96126f0c25e5e546/"
  }
]
```

# Throughput vs OWdelay vs Trace

```
[
  {
    "ts":1397421672,
    "val":7016320000.0
  },
  {
    "ts":1397442692,
    "val":7225480000.0
  },
  {
    "ts":1397466492,
    "val":7095460000.0
  },
  {
    "ts":1397482700,
    "val":7042540000.0
  }
]
```

```
[
  {
    "ts":1397504013,
    "val":{
      "34.4":506,
      "34.5":85,
      "34.6":5,
      "34.7":4
    }
  },
  {
    "ts":1397504052,
    "val":{
      "34.4":510,
      "34.5":80,
      "34.6":7,
      "34.7":3
    }
  },
  .....
]
```

```
{
  "ts":1397566094,
  "val":[
    {
      "error_message":null,
      "ip":"198.124.238.65",
      "mtu":"9000",
      "query":"1",
      "rtt":"0.246",
      "success":1,
      "ttl":"1"
    },
    {
      "error_message":null,
      "ip":"198.124.238.65",
      "mtu":"9000",
      "query":"2",
      "rtt":"0.195",
      "success":1,
      "ttl":"1"
    }
  ],
}
```



# Use Cases and Pilot Projects on Integration

# Use cases

- Feedback received from all experiments, but still missing some specific input – please add it ASAP
- Common use cases:
  - Provide latency and trace routes that can be integrated with throughput from transfer systems
  - Provide mapping between sites/storages and sonars
  - Uniform access to the network monitoring
  - Follow up on the WAN issues and help commission new links
- More details in [1]

# Pilot projects

- Experiment's Interface to perfSONAR/esmond
  - Backed by LHCb
  - Provide experiment agnostic prototype implementation to access datastore
  - Understand and propose solutions for the underlying issues – mappings, translations and post-processing needed
  - Initially focused on DIRAC, but intent is to come up with an implementation that can be used by other experiments

# Pilot projects

- FTS performance

- Currently done in ATLAS to identify bad links – possibility to extend this to other experiments

- Global report (cronned): <http://egg.bu.edu/atlas/adc/fts/plots/>
    - Single channel report: [http://egg.bu.edu/atlas/adc/fts/plots/singles/T1-T2D\\_internal/BNL-NET2/plots/index.html](http://egg.bu.edu/atlas/adc/fts/plots/singles/T1-T2D_internal/BNL-NET2/plots/index.html)
    - IP level summary: [http://egg.bu.edu/net2/studies%7btype:egg.Hatch%7d/internet-NET2-May-13/US\\_internet\\_from\\_NET2.html](http://egg.bu.edu/net2/studies%7btype:egg.Hatch%7d/internet-NET2-May-13/US_internet_from_NET2.html)

- Similar activity is also starting in the FTS dashboard

- <http://dashb-fts-transfers.cern.ch/ui/#>

- Integration of FTS and perfSONAR beneficial to follow up on bad links, but also useful to tune up FTS optimizer

# AOB

