Implementing perfSONAR in the South African National Research and Education Network

Kevin DRAAI[1,2], Roderick MOOI[1,3]
[1]*Council for Scientific and Industrial Research, P.O. Box 395, Pretoria 0001, South Africa*
[2]*Tel: +27 841 2342, Email: kdraai@csir.co.za*
[3]*Tel: +27 841 4111, Email: rmooi@csir.co.za*

Abstract
The South African National Research Network (SANReN) serves over 1 million users who expect high performance. To ensure optimal operation, rapid detection and correction of abnormalities is crucial. PerfSONAR is a network measurement toolkit that can be used to test and monitor end-to-end network performance and achieve this goal. The tools provided can verify the limits of a network (primarily in terms of throughput and loss/latency) and reveal faults and issues. The SANReN team use perfSONAR to evaluate the network on a hop-by-hop basis with test results conveyed through a dashboard for rapid visualisation and alerting. There are many cases where it has proven its utility in the SANReN network. This paper provides an overview of the SANReN and perfSONAR prior to covering the deployment scenario and use cases to highlight the benefits provided.

Keywords:
perfSONAR, network performance, network troubleshooting, NREN advanced services.

1. Introduction
Measuring and monitoring the performance of the South African National Research Network (SANReN) has become increasingly important because of the complexity and size of the network. In order to report on the actual delivered bandwidth and quality of the network, throughput, loss and latency were identified as meaningful metrics. Following an investigation of the available tools, the perfSONAR toolkit[1] was identified as the most suitable solution for measuring and monitoring these network performance metrics. The SANReN competency area (CA) of the Council for Scientific and Industrial Research (CSIR) has been deploying perfSONAR nodes on the network since 2012. We now have 15 x 10Gbps dual-homed nodes covering most of the national backbone of the network and both international links. This paper provides an overview of both the SANReN and perfSONAR prior to expanding on the SANReN deployment details. Thereafter we reveal some troubleshooting scenarios where perfSONAR proved its value before concluding the paper.

2. SANReN Overview
The SANReN was conceptualised in 2003 and implemented by the Meraka Institute of the CSIR under contract to the Department of Science and Technology (DST) from 2008 onwards *(TENET, 2013)*. The SANReN has a national backbone capacity of 10Gbps. Typical client handoff is 1Gbps though some main campuses connect at 10Gbps.

---

[1] Available at: www.perfsonar.net

The backbone and metro networks are deployed in ring topologies for performance and redundancy. The SANReN high-level topology is shown in Figure 1.

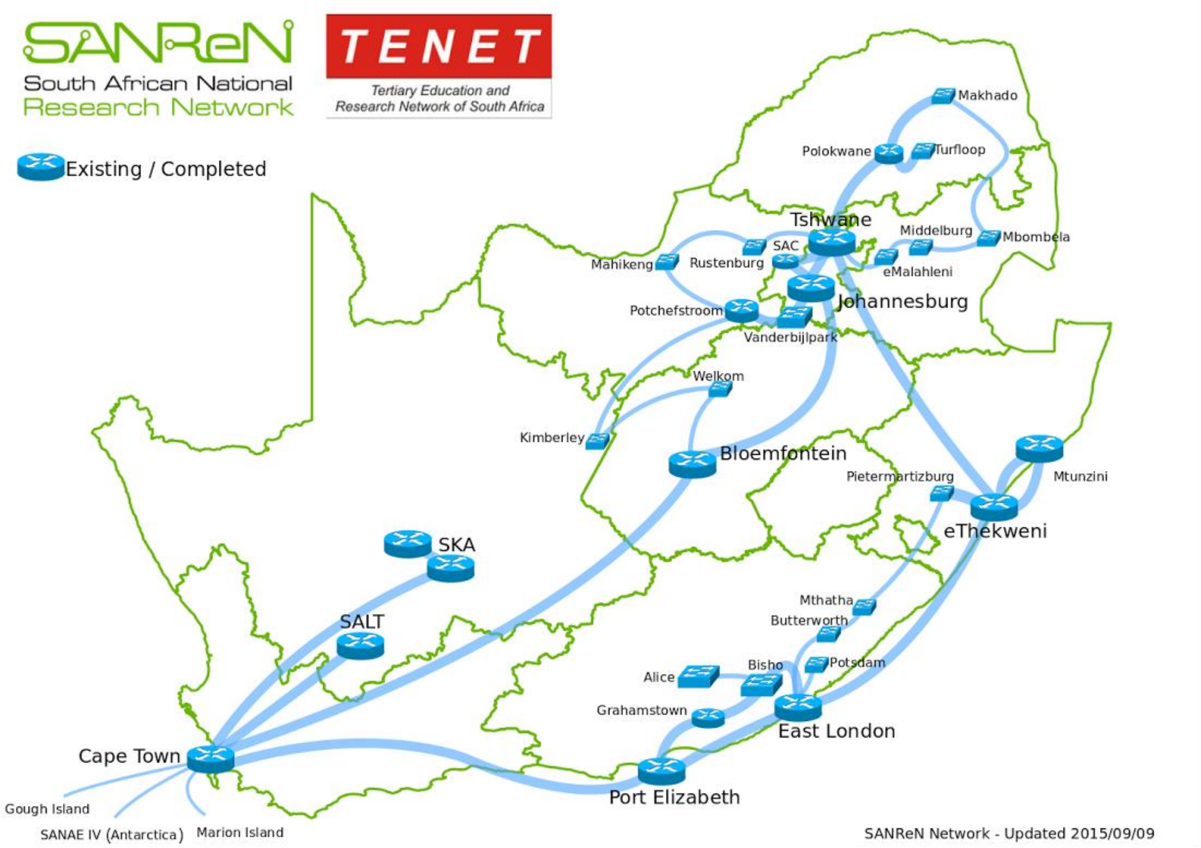

Figure 1: SANReN Topology (by S. Mammen <smammen@csir.co.za>)

The SANReN CA design and build the network while the Tertiary Education and Research Network of South Africa (TENET) operate it. Together these two parties form the South African National Research and Education Network (SA NREN). The SA NREN provides value-added advanced services such as eduroam and video conferencing to its customers and beneficiaries. perfSONAR is currently regarded as an internal service intended for customer-facing expansion.

3. PerfSONAR Overview

PerfSONAR is a joint project that was started by several Research and Education (R&E) networking and other interested parties. The current uses of perfSONAR include the collection and publication of latency, achievable bandwidth, utilization and network topology data as well as the assisting with the diagnosis of performance issues (Tierney, et al., 2009). The aim of the toolkit is to provide a universal representation of network performance thereby aiding network engineers with trouble-shooting and optimisation tasks. PerfSONAR is built on standard Linux distributions enhanced with the web100 kernel. The configuration and tuning is done via the command-line (back end) and tests are scheduled on the web interface (front end). Since our first use of perfSONAR, we have witnessed constant improvements in the toolkit with new features, reliability and support for system administrators in each new version. The mailing list is active, responsive and of high quality[2].

3.1 Version 3.5

The latest version of perfSONAR was released in September 2015. It comes with a new web interface, support for alternative installation methods as well as enhanced auto-configuration and central management features to assist with large deployments. More information on version 3.5 can be found at[3].

3.2 Features and Uses

The perfSONAR toolkit boasts many features and components. These include measurement tools, measurement archive, host management tools, data analysis tools, and a lookup service (Tierney, et al., 2015). The measurement tools include iperf3 (throughput), bwctl (scheduler), owamp (loss/latency), owping (one-way ping) and traceroute. The central measurement archive stores all measurement data in an easily accessible esmond database[4]. Host management tools allow the user to configure the node details, running services, NTP servers and scheduled tests. Test results can be viewed from the web frontend where on-demand ping and traceroute tests can be executed as required.

In addition, the perfSONAR toolkit features a mesh capability. This enables an adWministrator to configure a mesh of tests between any number of perfSONAR nodes on a network. The mesh configuration can be centrally hosted and retrieved by all nodes which then execute the relevant tests and store the results. A central measurement archive can also be used to store the results (which are pushed to the archive by the nodes on test completion). The SANReN perfSONAR team uses both individual as well as central storage for redundancy and quick access to results.

It would make sense to visually represent these tests between perfSONAR nodes in an easy-to-read central dashboard setup. The Monitoring and Debugging Dashboard (MaDDash)[5] is designed for this purpose. With this feature, throughput and loss/latency tests can be viewed between perfSONAR nodes as a grid, issues can be rapidly identified and detailed test results accessed from a central interface. The SANReN perfSONAR dashboard can be seen in section 4.

The Network Diagnostic Tool (NDT) and Network Path & Application Diagnostics (NPAD) tools are included with perfSONAR and assist in identifying bottlenecks[6]. They allow a user to diagnose internal network problems and bottlenecks by testing the path between the local computer and the perfSONAR server. NDT provides network configuration and performance testing to/from and end user's workstation[7]. Lastly, the perfSONAR project provides a

---

[2] see: http://www.perfsonar.net/about/getting-help/

[3] http://www.perfsonar.net/

[4] http://software.es.net/esmond/

[5] http://software.es.net/maddash/

[6] https://fasterdata.es.net/performance-testing/network-troubleshooting-tools/ndt-npad/

[7] http://software.internet2.edu/ndt/

lookup service directory where all registered nodes across the globe can be discovered[8]. The SANReN perfSONAR deployment is discussed in the following section.

4. SANReN PerfSONAR Deployment

The SANReN CA (with TENET's assistance) have deployed perfSONAR nodes across the national backbone as well in London (Telecity) and Amsterdam (NikHef) [connected to the UbuntuNet Alliance routers]. The International nodes are used to test and monitor the throughput and latency of the SEACOM and West Africa Cable System (WACS) links that provide international connectivity to the SANReN network and TENET's customers. The rest of the nodes on the network are deployed at locations on the network that allow tests to cover critical portions of the national backbone. We use DELL PowerEdge R320 rack servers with Intel X520 NICs to enable 10Gbps testing. SANReN also keeps some mobile nodes for ad-hoc troubleshooting up to 1Gbps.

4.1 Topology and Tests

We have implemented a mesh of tests between our perfSONAR nodes. These tests run on a hop-by-hop basis in an effort to troubleshoot each layer 3 network segment on the backbone without duplicating coverage. Sometimes there are a few hops (routers) between the perfSONAR nodes which could affect TCP throughput. To eliminate this and increase the visibility of problematic paths, more nodes can be deployed. Alternatively, temporary routing changes also assist in troubleshooting paths that span multiple routers and switches or use mechanisms like Equal-Cost Multi-Path. To test a routing change, ad-hoc tests can be run via the web interface or the command line of any perfSONAR node (nodes can run tests between other nodes if required). The SANReN perfSONAR deployment logically overlays the national backbone creating a pseudo-ring topology. The deployment is shown in Figure 2.

[8] accessible at: http://stats.es.net/ServicesDirectory/

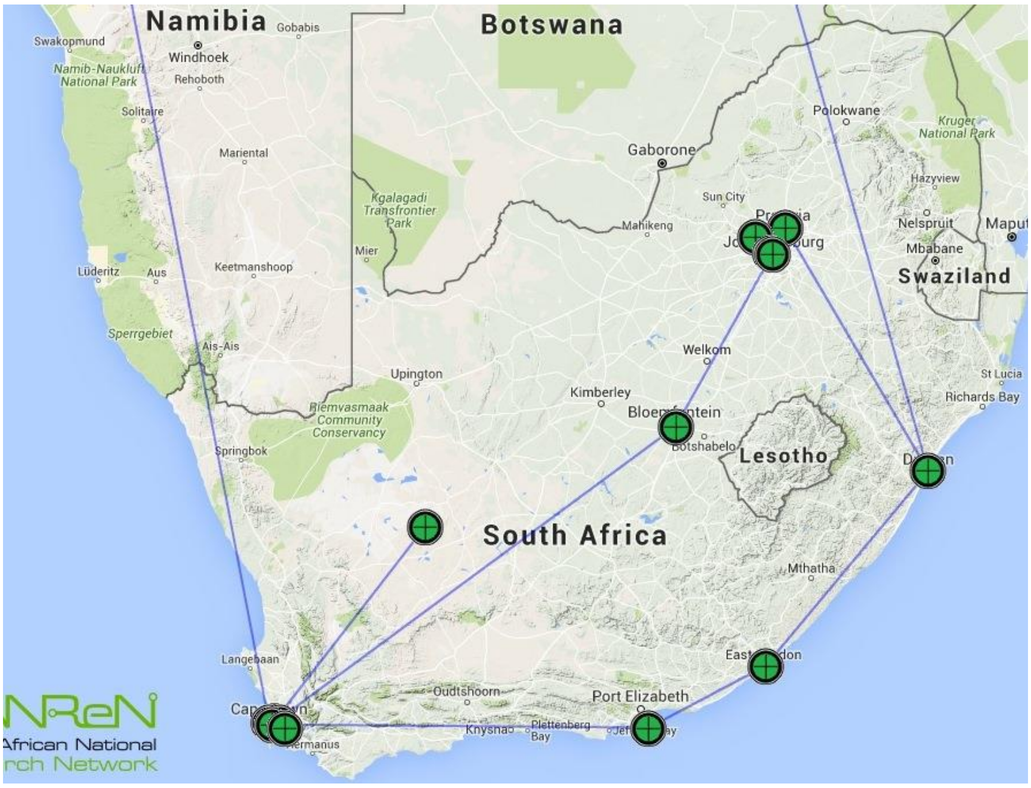

Figure 2: SANReN perfSONAR Topology

All the nodes on the network are dual-homed. This means that throughput and latency tests are separated. It also has an added advantage of connection redundancy. Throughput tests are run on the 10Gbps port and latency tests are run on the 1Gbps port. In this way throughput and latency tests don't interfere with each other.

4.2 MaDDash

The Monitoring and Debugging Dashboard (MaDDash) is deployed on a central node in the SANReN network. This retrieves all tests run between servers and puts it in a grid format as shown in Figure 3.

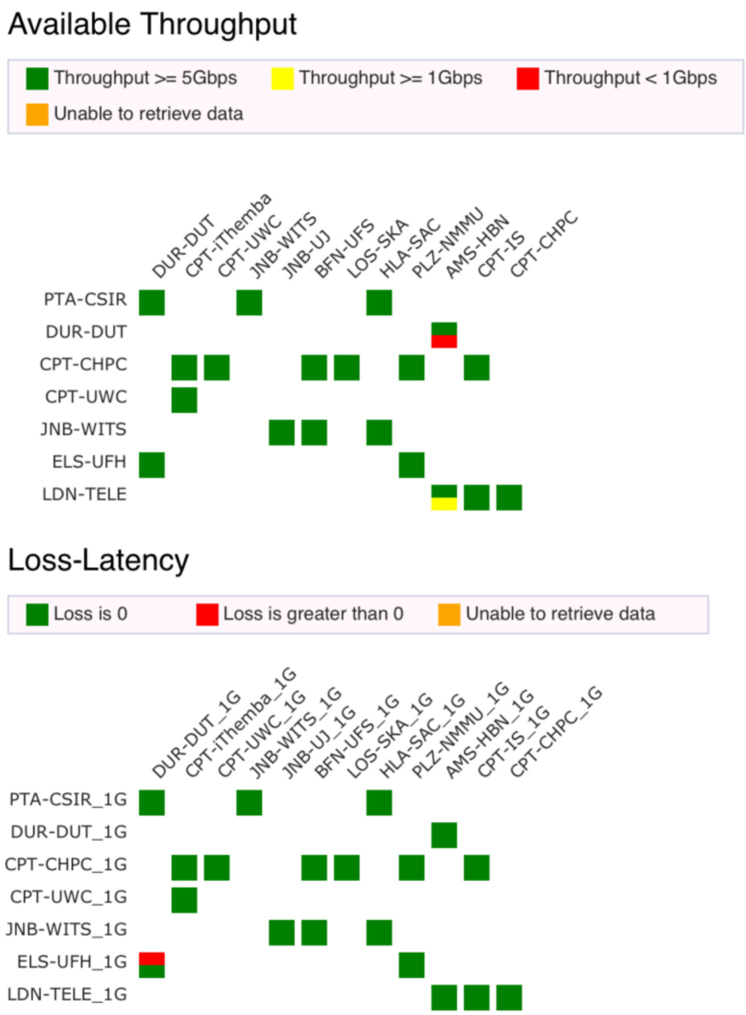

Figure 3: SANReN perfSONAR Dashboard

The dashboard is targeted towards Network Operations Centre (NOC) staff as there are thresholds set for when throughput drops below a certain value or loss greater than zero. When these thresholds are crossed, a state change is triggered causing the blocks to change colour and thereby reflecting a warning (yellow) or critical (red) state. Clicking on a block displays the detailed test results for that link. Utilising the perfSONAR toolkit, and the dashboard in particular, the SANReN CA and TENET staff have been able to detect and isolate problems on the network that were previously unnoticed and difficult to detect otherwise.

4.3 The Usefulness of perfSONAR in the SA NREN

PerfSONAR has proved its usefulness for the SA NREN in numerous scenarios. These range from poor throughput caused by improper configuration of routers on the network to physical problems on the links provided by telecommuncations operators. PerfSONAR picks up on these faults simply by indicating that the throughput, latency or loss measurements between nodes are not according to expectations.

PerfSONAR may further indicate that traffic between nodes is not traversing through an optimal route. An example of an incident where traffic was taking a suboptimal route was when throughput tests were run between Durban and Amsterdam. The throughput was low so a traceroute was done to figure out where the traffic was going and it seemed that all the UbuntuNet Alliance traffic was traversing London before it got to Amsterdam. It was also found that all UbuntuNet Alliance traffic was going to Amsterdam first then back to UbuntuNet Alliance nodes in Africa. Figure 4 highlights the problem and shows the effect of correcting the traffic paths.

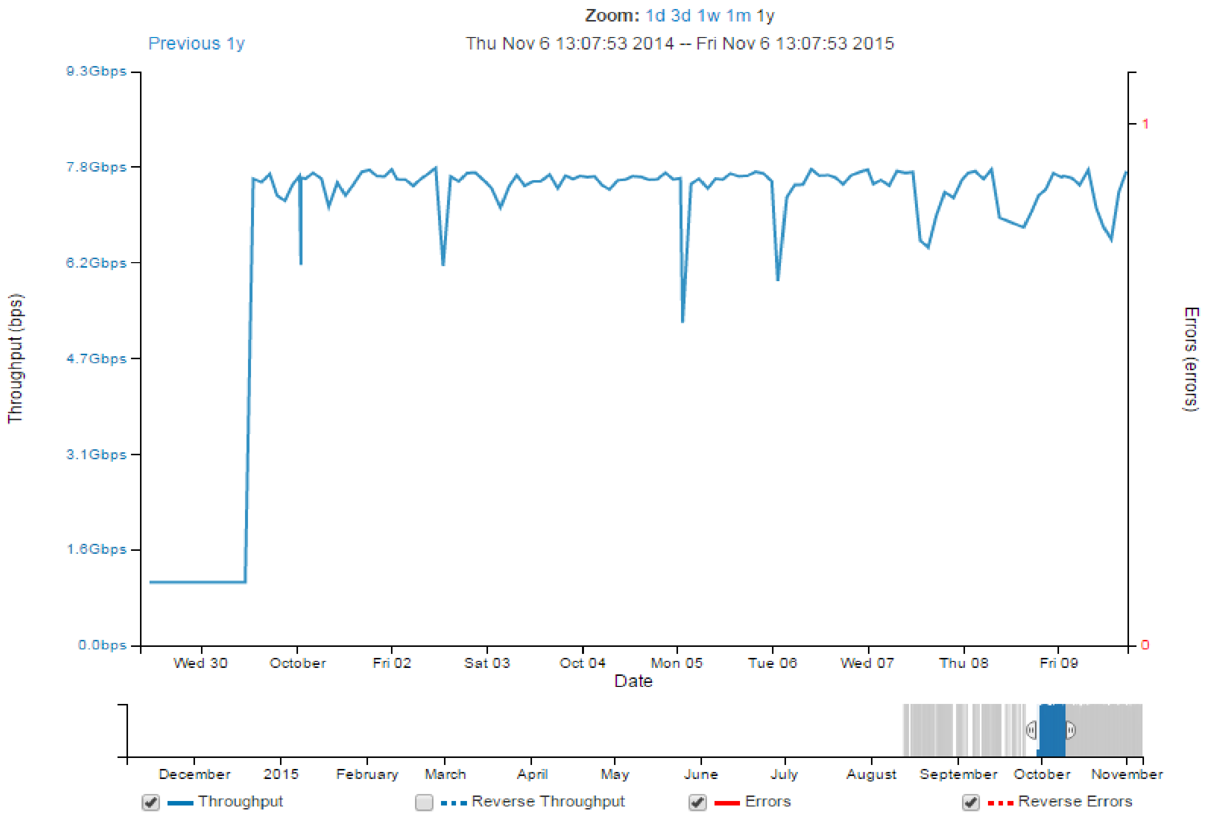

Figure 4: Durban to Amsterdam Routing Problem and Rectification

Jumbo frames need to be enabled for throughput of close to 10Gbps to be achieved. The 9000 byte MTUs need to be configured on the 10Gbps ports of the end nodes as well as routers and switches along the path between perfSONAR nodes. If this is not done, achievable throughput ranges between 1 and 2 Gbps (depending on path length) even for an uncongested 10Gbps link.

To test the TCP throughput of a link that is geographically long the buffer settings on the source and destination perfSONAR nodes need to be tuned to match the Bandwidth-Delay-Product (BDP) of that link. The latency (RTT) of these links are normally greater than 100ms. If the maximum buffer sizes on the source and destination nodes are too small then the throughput will not reflect the available bandwidth at that point in time. An example of an event like this on the SANReN network was between Cape Town and London (over the WACS). After correcting the buffer sizes, the TCP throughput went up from ~150Mbps to 7Gbps. Figure 6 illustrates this event.

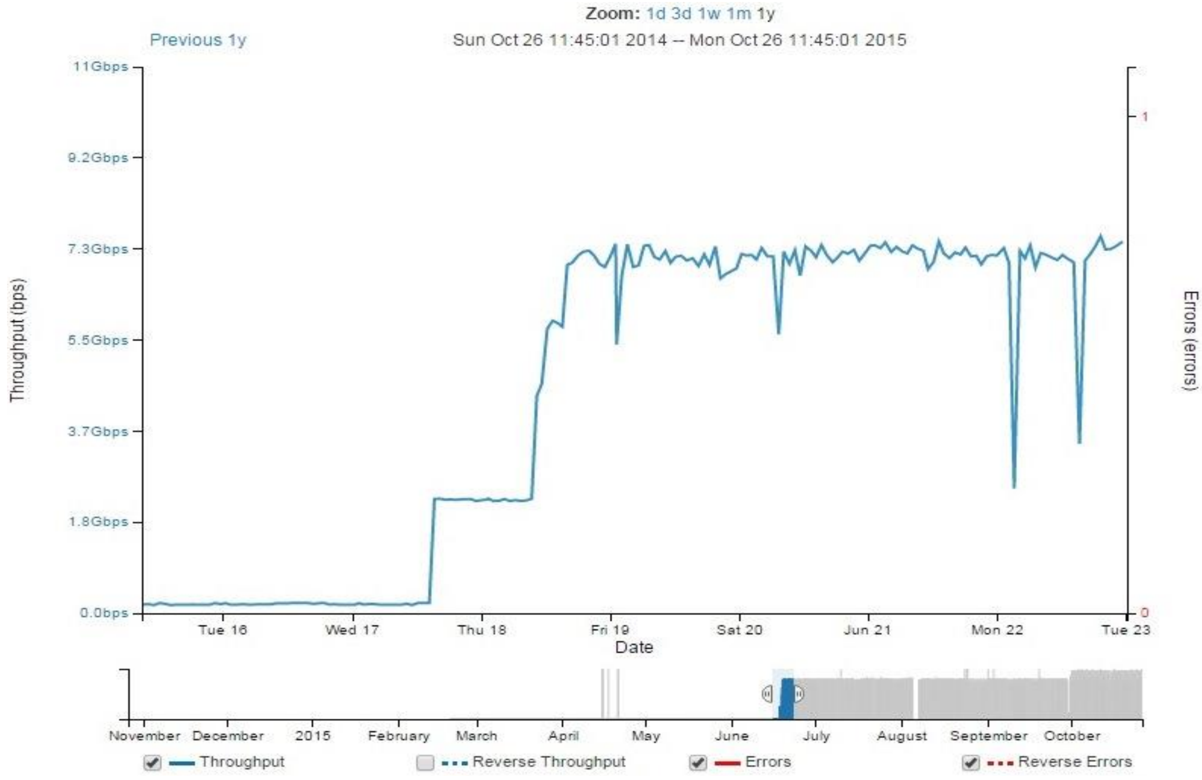

Figure 5: Host Buffer Tuning Fix

On an NREN there might be link breaks. These breaks can be on dark fibre links or managed links and can be difficult to detect as a break or fault may result in degraded performance and not necessarily loss of connectivity. PerfSONAR picks up on these issues as well by indicating loss and low throughput on the link together with latency changes. Figure 7 illustrates an example of a fault between Pretoria and Durban.

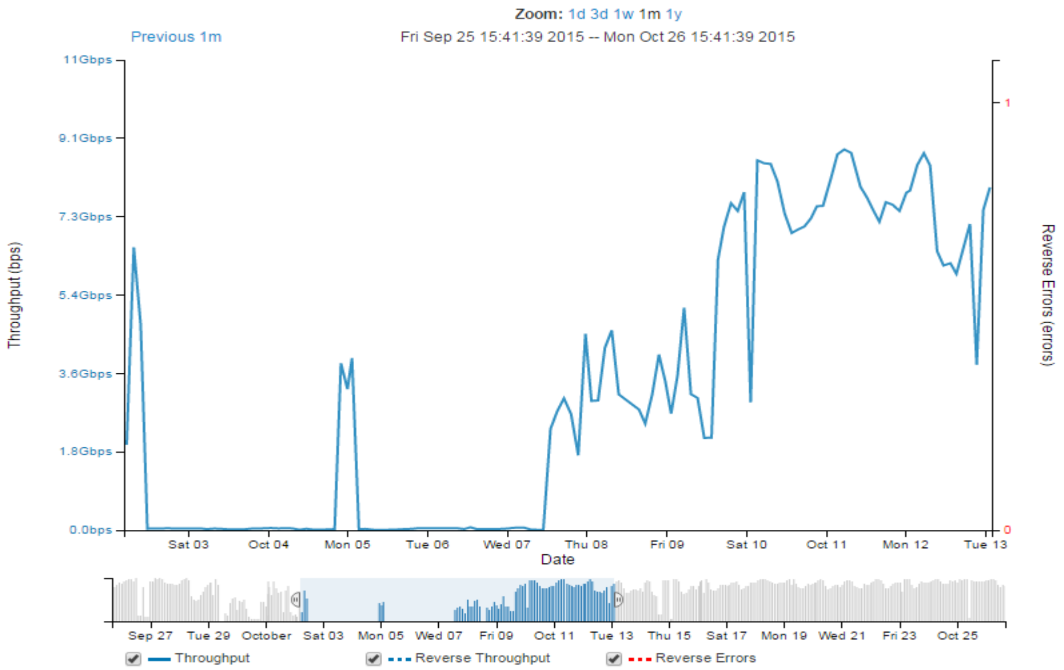

Figure 6: Pretoria to Durban Faulty Link

The link break and restoration events are evident from the graph. Finally, Figure 7 shows the loss and latency events for one week while investigating issues on a link between London and Cape Town.

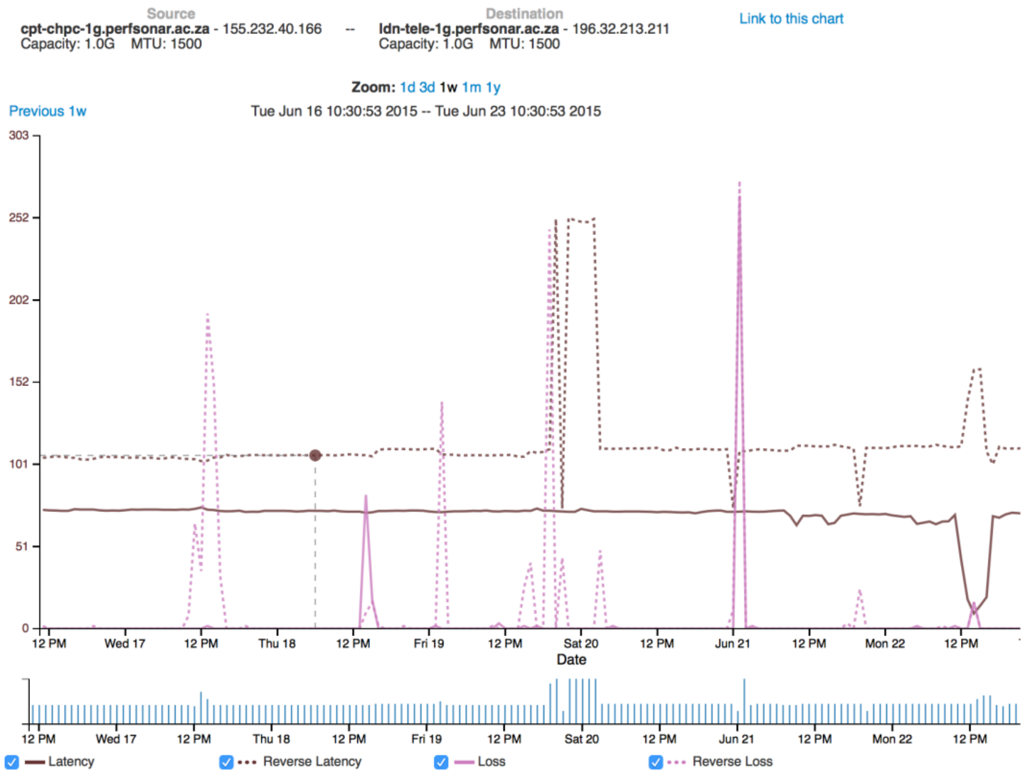

Figure 7: Loss and Latency events between London and Cape Town

5. Usefulness beyond the SA NREN

Two further use cases are worth mentioning. The first is in the investigation of low throughput experienced by the SKA project site connected to the SANReN. The second was a faulty NASA router verified by one of our perfSONAR nodes increasing visibility due to the increased latency between the router and our node.

6. Conclusion

The SANReN team has deployed perfSONAR servers effectively measuring and monitoring the entire national backbone and both international links. The purpose of this service in the SANReN network is to monitor the performance of the network, assist in troubleshooting, tune the network to reach its full potential and assess whether the network can actually reach its potential given the current load on the network. After providing an overview of both the SANReN and perfSONAR, we proceeded to explain SANReN's deployment and highlighted some use cases of the toolkit on the network.

The SANReN CA's next steps for perfSONAR include additional community engagement, MaDDash email/sms alerts, integration of perfSONAR in new link procurements, increased mobile node deployments (for specific trouble-shooting scenarios) and contributing to the African perfSONAR mesh.

Bibliography

Tierney, B., 2001. TCP Tuning Guide for Distributed Applications on Wide Area Networks.

Tierney, B., Fowler, L. & Boyd, E., 2015. perfSONAR: The Road to 100k Nodes. s.l., s.n.

Tierney, B. et al., 2009. perfSONAR: Instantiating a Global Network Measurement Framework. s.l.:s.n.

Biography

Kevin Draai is a SANReN engineer. His primary responsibilities are the implementation of services such as perfSONAR and Network Visualisation as well as Network Roll out and Inventory Management. He is currently studying towards an M-Tech in Communication Networks at the Nelson Mandela Metropolitan University (NMMU).

Roderick Mooi is an experienced software engineer who became part of the SANReN CA in 2012. His projects are focused on advanced (value-added) services development and include perfSONAR and the SA NREN CSIRT. Roderick has an M. Tech. in Information Technology. He also works ad-hoc on tools to improve team productivity as well as participating in network design and tender evaluation activities.