

Over-provisioned network performance analysis

OvNet

Objectives

- ... to understand the behavior of production-like networks

Description of the activity (1)

- No exact definition of over-provisioning in the bibliography
- ... to provide enough bandwidth to always match committed network service guarantees
- ... 10% occupancy

Overprovision

- Bandwidth ~ 10 % Capacity (at least)
- Delay (~ Transmission Delay)
- Jitter (processing Delay)

Instrumentation

- Synchronized packet generation at least at 1 Gbps
- NTP or GPS
- NTP ~ 10 – 20 ms accuracy margin
- GPS instrumentation
- i.e. smartbits, ixia

Program (1)

- Test scenario A: Non-multi-hop path
- Test scenario B: Multi-hop path
- Test case 1: Empty network
 - One (1) flow
 - Measurements: packet loss as function of network load (in OvNet packet loss is expected 0), delay & jitter
- Test case 2: Background noise
 - One (1) flow
 - + production-like traffic
 - Measurements: delay & jitter

Program (2)

- Test case 3: Aggregation - Background noise
 - N flows
 - + production-like traffic
 - Measurements: N, delay & jitter
- Test case 4: PQ- Background noise
 - One (1) flow (with PQ)
 - + production-like traffic
 - Measurements: delay & jitter
- Test case 5: PQ - Aggregation - Background noise
 - N flows (with PQ)
 - + production-like traffic
 - Measurements: N, delay & jitter

Program (3)

- Test case X: MPEG-1/MPEG-2 streams over RTP/UDP/IP (measure jitter on application level with/without background traffic)
- Test case Y: Packet re-ordering

Requirements (1)

- Framework (IP Performance Metric charter)
- Test boxes that are able to generate high data rates
- Traffic generators producing realistic background traffic
- High bandwidth network
 - Local test labs
 - International tests
 - Fast connection over large distance for increased RTT's ?
- QoS measurement software (monitor congestion on short term)

Requirements (2)

- Ping, traceroute
- Pathchar, pchar, etc.
- Ttcp/netperf/tcpblast
- Chariot
- RIPE TTM/Surveyor
- Netflow

Requirements (3)

Metric	Surveyor	RIPE	PingER	AMP	Skitter
Method	1 way delay & loss	1 way delay & loss	2 way ping	2 way ping	traceroute like
Hosts	Dedicated	Dedicated	"selected"	Dedicated	Dedicated
Time synchronization	GPS	GPS	NTP	NTP	NTP
Frequency (load average)	~2*2/s (~2kbps)	~3/min (0.330kbps)	~0.01/s (~0.1kbps)	~ 1/minute	Hourly
Scheduling	Poisson <2/s>	Poisson <1/min>	bursty (30 min)	Linear random about 1st 15 seconds of min.	~30 min.
Packet size	~ 40Bytes	100Bytes	100Bytes & 1000Bytes	64Bytes	52Bytes
Locations	US, CA, CH, NL & NZ	EU, IL, US	10 monitoring site countries, 22 remote site countries	US, NZ, NO	Monitors in Asia, CA, UK, US
Monitors	~51 (Jul-99)	~32 (Jul-99)	18 (Jul-99)	~70 (Jul-99)	20
Pairs	~1000	1024	~1200	~4600	35000
Data start	1997	1998	1995	1999	1998
Data availability	Upon request	Upon request	Public access via Web	Public access via Web	?
Data storage	~38MB/pair/mo	2Mbytes/pair/mo	~0.6MB/pair/mo	~1.3MB/pair/mo (0.5MB zipped)	?
Sponsors/Community	CSG / Advanced	RIPE / European R&E sites	DOE / ESnet / HENP / XIWT	NSF / NLANR / Internet 2	DARPA / NSF / CAIDA

- [AMP](#): the [National Laboratory for Applied Network Research \(NLANR\)](#) Active Measurement Program (AMP) for High Performance Computing (HPC) awardees is intended to improve the understanding of how high performance networks perform as seen by participating sites and users, and to help in problem diagnosis for both the network's users and its providers. The community of interest is the [National Science Foundation \(NSF\)](#) HPC program awardees.
- [PingER](#): the [DOE / MICS Internet End-to-end Performance Measurement \(IEPM\)](#) project to provide active monitoring of end-to-end performance of Internet links. The community of interest is [ESnet](#), [High Energy & Nuclear Physics \(HENP\)](#) and the [Cross Industry Working Team \(XIWT\)](#).
- [RIPE's](#): project goal is to do independent measurements of connectivity parameters, such as delays and routing-vectors, in the Internet. It's community of interest is European Internet Service Providers (ISPs) and their users.
- [Skitter](#): is primarily intended to be used to measure forward IP paths (each "hop") from a source to many destinations. It is supported out of [Cooperative Association for Internet Data Analysis \(CAIDA\)](#). The community of interest and much of the funding comes from the [Defense Advanced Research Project Agency \(DARPA\)](#), and the NSF.
- [Surveyor](#): uses active tests of one-way delay and packet loss along paths between measurement machines at [CSG](#) sites, and some associated sites.


Surveyor ~ RIPE

- Dedicated PC running Unix to be placed at each monitoring site.
- Global Positioning System (GPS) device to obtain accurate time and to synchronize time.
- Poisson randomized time intervals
- One way end-to-end delay and loss measurements.
- Concurrent traceroutes.
- The community for Surveyor is Internet 2, though there are monitors at non Internet 2 sites, and in particular at 3 Higher Energy Physics (HEP) sites CERN, FNAL and SLAC that are also PingER monitor sites. The community for RIPE is European Internet Service Providers (ISPs), though again there are RIPE machines at CERN and SLAC.


Requirements (4)

- ITU
- TEQUILA
- IETF (IPPM)
- Internet2 E2E performance
- Eurescom?


Network (1)



Network (2)



Network (3)



Time scale

- Provide experimental results at the end of ...



Participants

- GRNET
- INFN-GARR
- INFN-CNAF
- HEAnet / SURFnet
- University of Utrecht