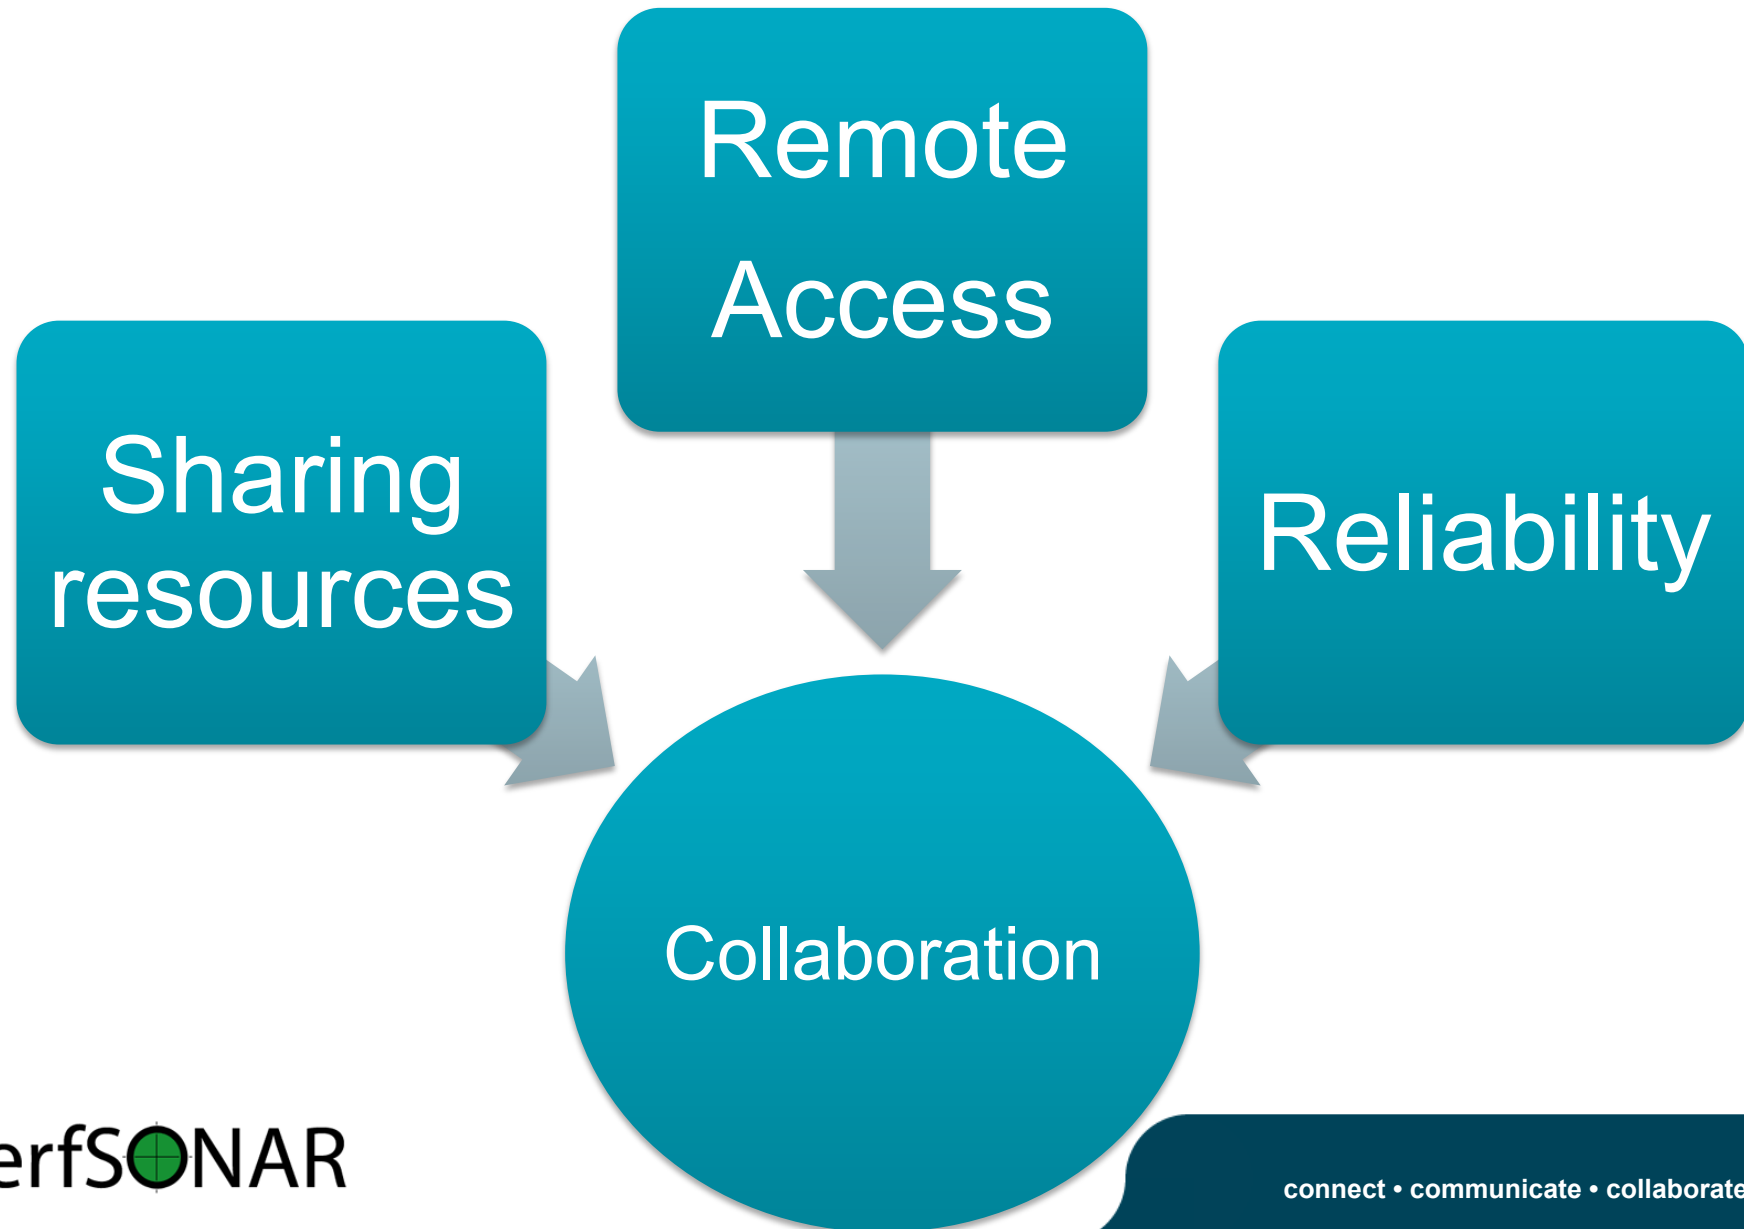


Network Monitoring and Troubleshooting with perfSONAR MDM

Domenico Vicinanza
DANTE, Cambridge UK


Performance U! – First Winter School – Zurich, 6-8 March 2013



New field: Arts and Humanities



- Advanced networking is becoming increasingly popular in the field of the arts and humanities
 - Enabling collaboration
 - Sharing experiences
 - Creating virtual stages



In a way which is unimaginable only a few years ago!



GÉANT launch event performance:
Musicians in Stockholm
Dancers in Kuala Lumpur

<http://www.geant.net/Events/LaunchEvent/Pages/EventHighlights-Day1.aspx>

Why research and education networks?

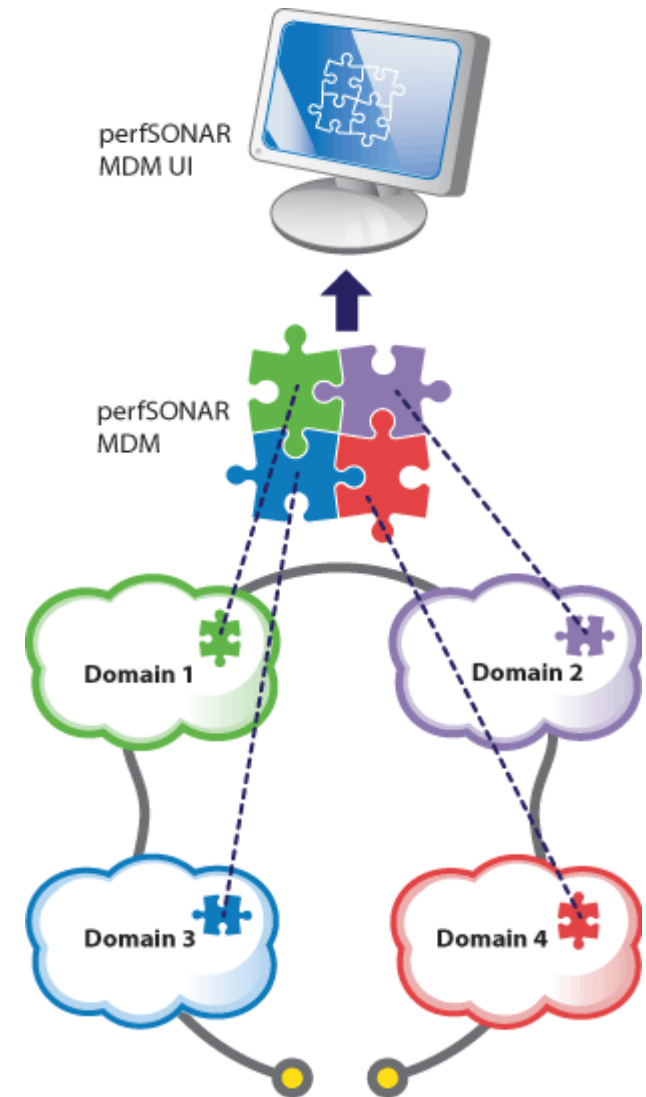


- Scientific and artistic collaboration are using R&E networking
- Ideal platform:
 - Accessible, high-quality (over-provisioned) IP network
 - Availability of point-to-point circuits
 - Support and expertise available within GÉANT/NRENs
 - *to facilitate the use of advanced networking for the arts*

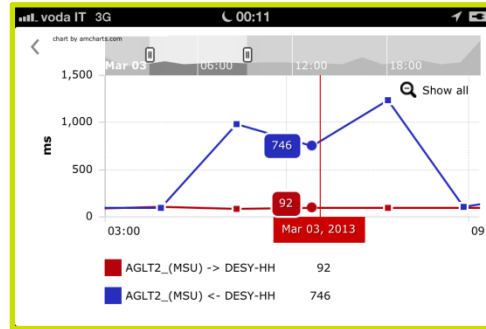
Brilliant performances Crucial monitoring → perfSONAR



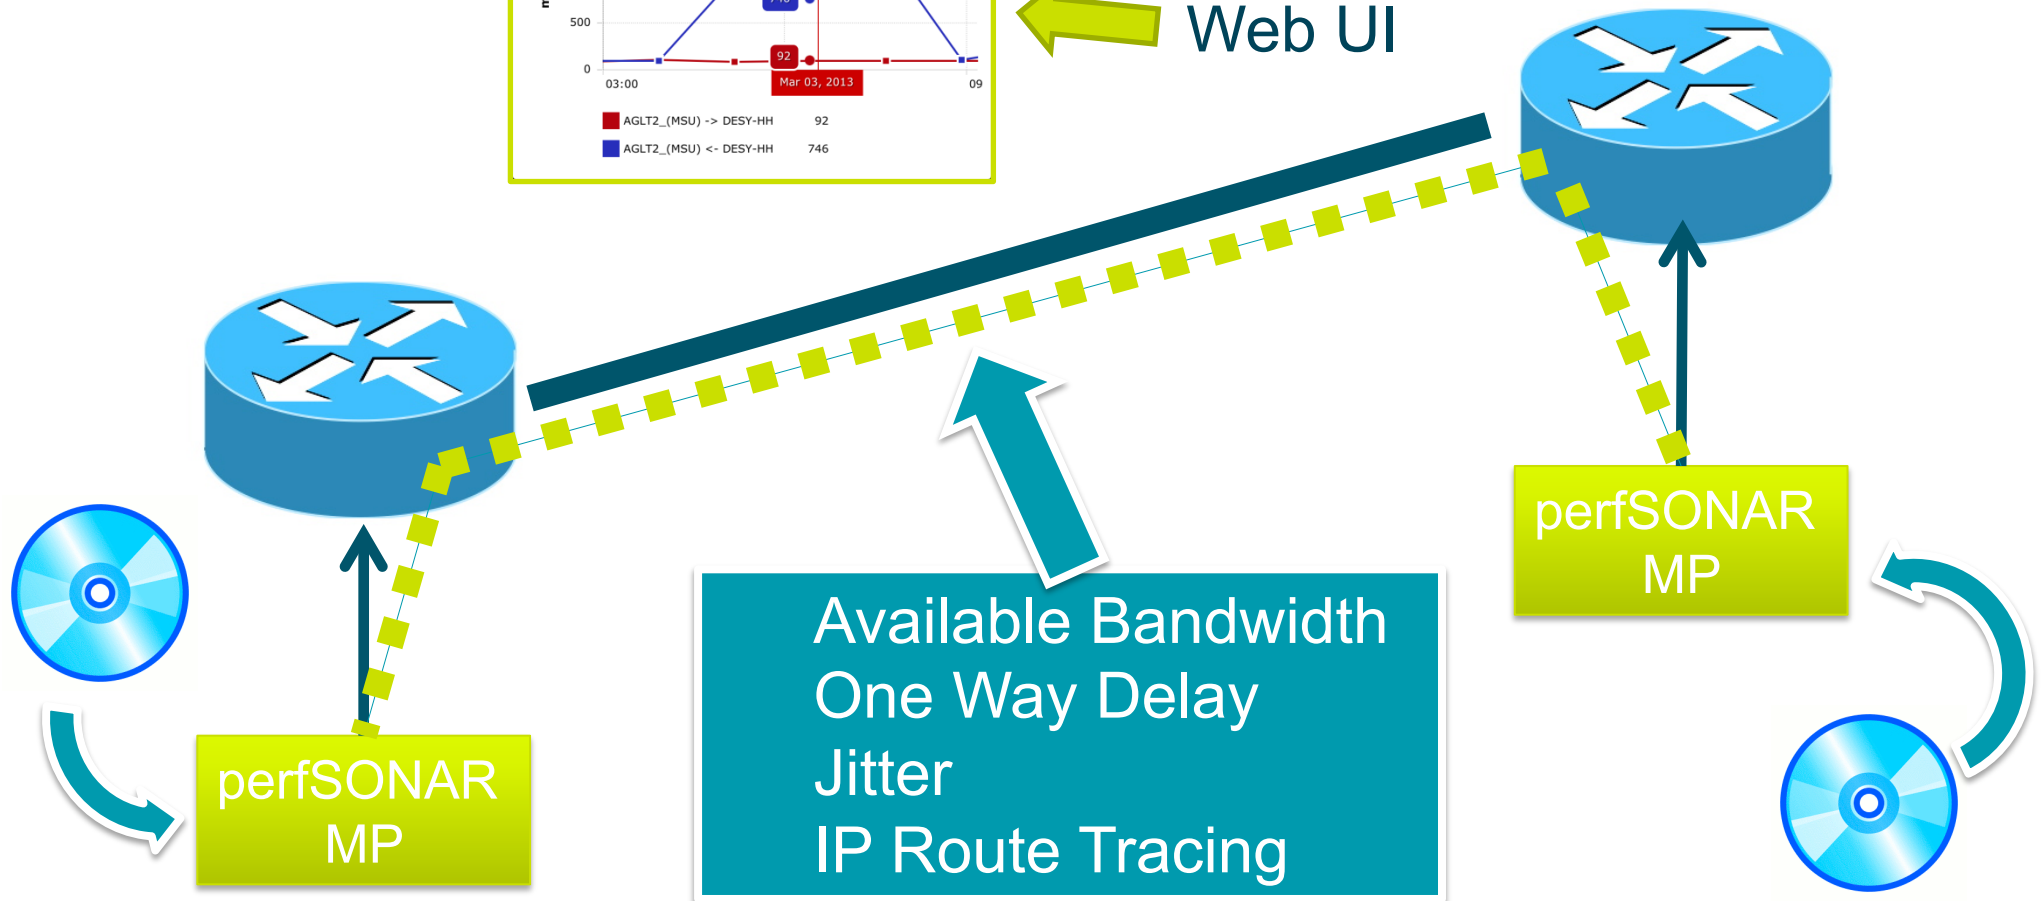
- Advanced networking requires
 - Advanced monitoring
 - Advanced troubleshooting
- International scale:
 - Different network domains
- The solution: perfSONAR
 - Distributed multi-domain monitoring
 - Interoperable on a global scale
 - Easy to use



How does perfSONAR work?



Web UI



perfSONAR

What does perfSONAR measure and how?

Link utilisation, input errors, output drops (RRD-MA)



- **Purpose:**
 - Monitor link utilisation, input errors, packet drops
 - Provide access to historical measurements
- **Strategy:**
 - Query router interfaces statistics using SNMP
 - Store data into RRD files
 - *made accessible through web-service*

Link Utilisation User Interface




perfSONAR WEB USER INTERFACE

Link (L) to DK-GN

	Address	Domain	Capacity
1	reth104.0	DANTE Server LAN 10.100.4.0/24	0 bps
2	xe-2/1/0.320	Link to PIONIER Backup (Prx Ids)	10 Gbps
3	xe-4/0/0.0	Link (L) to LITnet	10 Gbps
4	ae0.0	Link (L) to GRnet	20 Gbps
5	as3.0	Link (L) to DK-GN	19.9 Gbps

Interface details for as3.0

Name	Status
as3.0	OK



OWD, jitter, packet loss, traceroute (HADES/OWAMP)



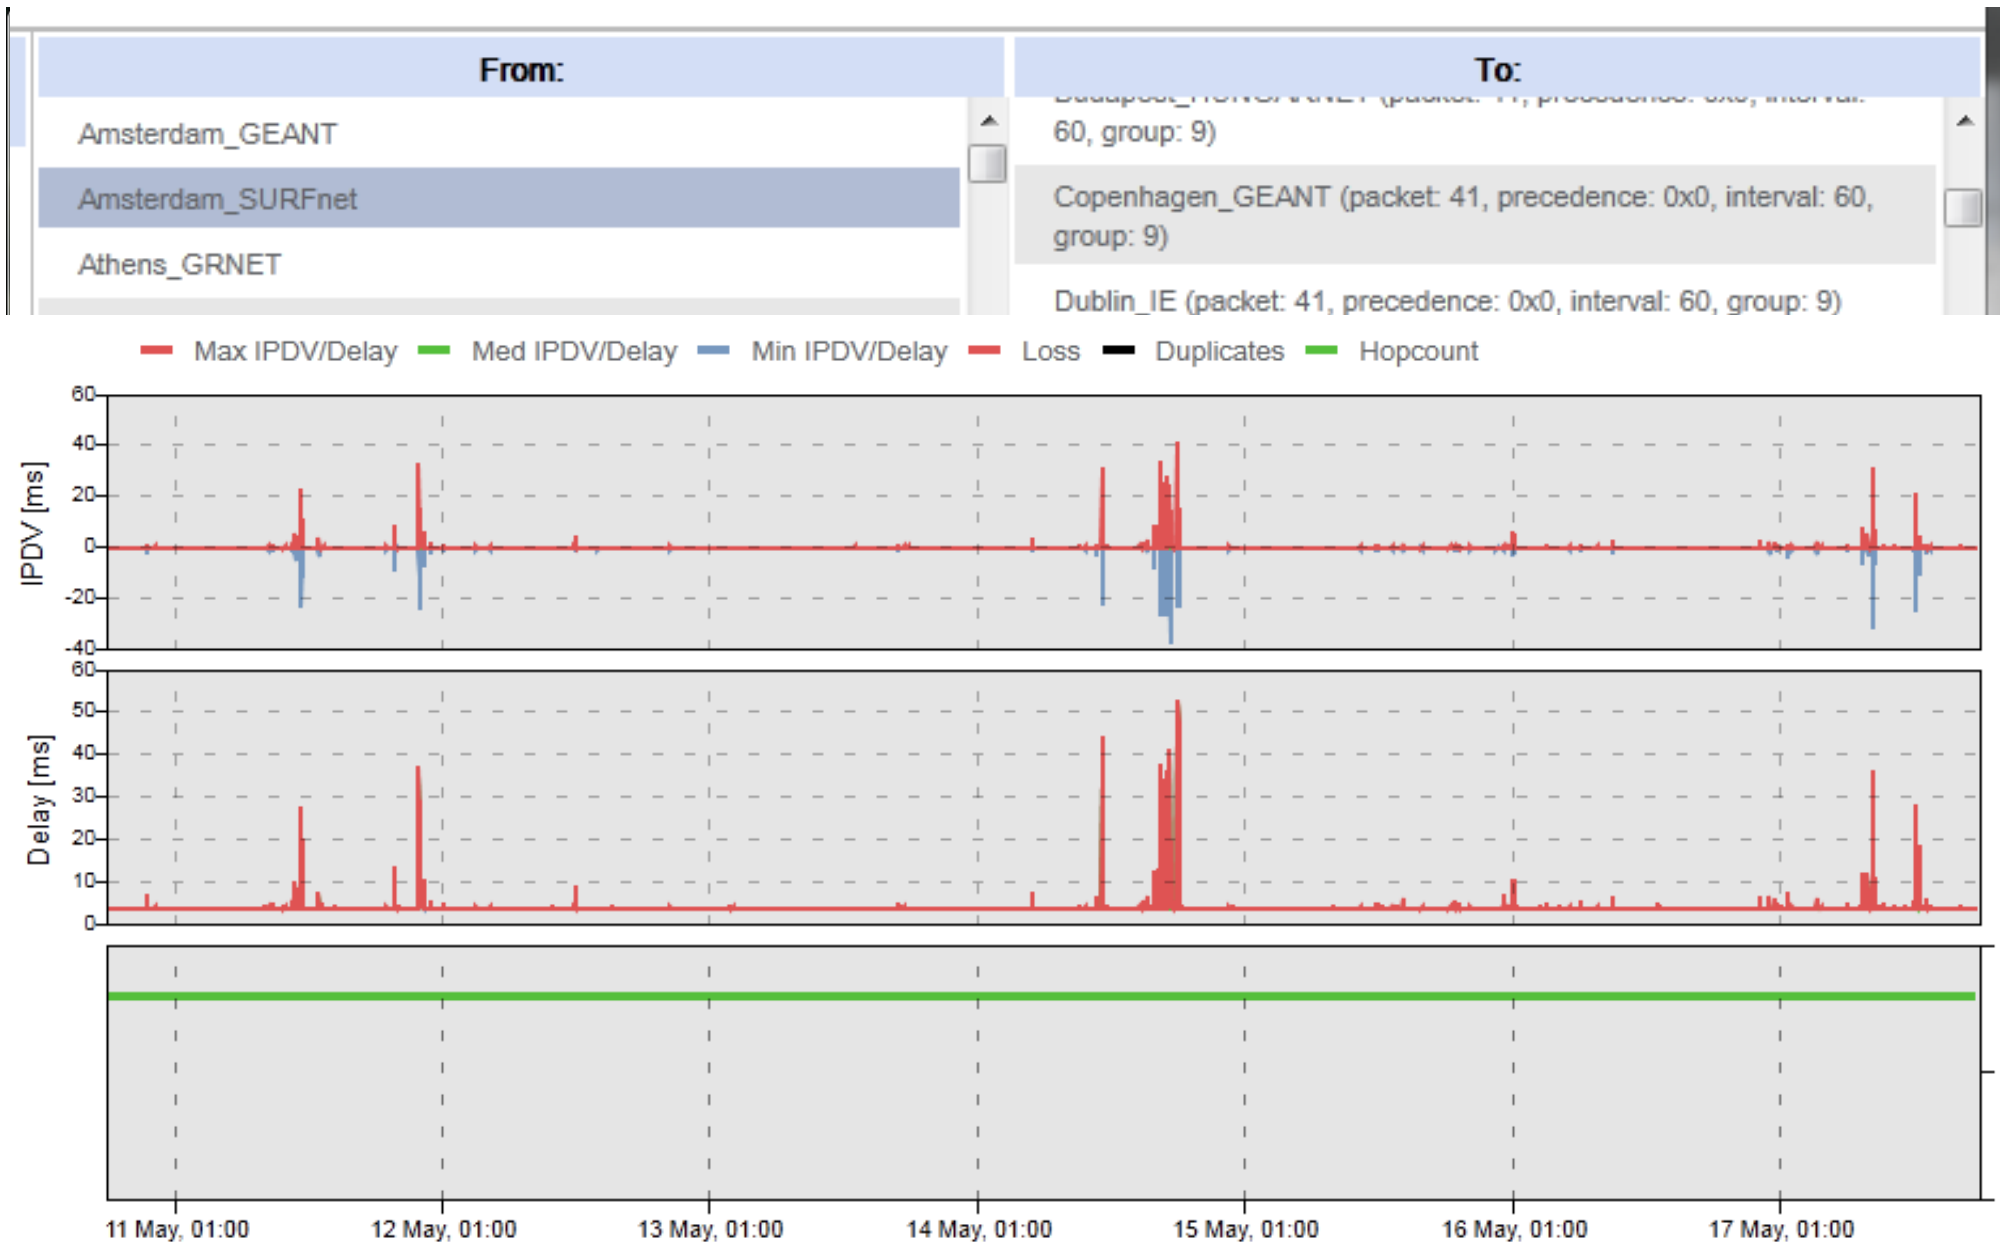
- **Purpose:**

- Monitor OWD, jitter, packet loss, traceroute variations
 - *Regularly scheduled*
 - *On demand (to be implemented)*
- Provide access to historical measurements

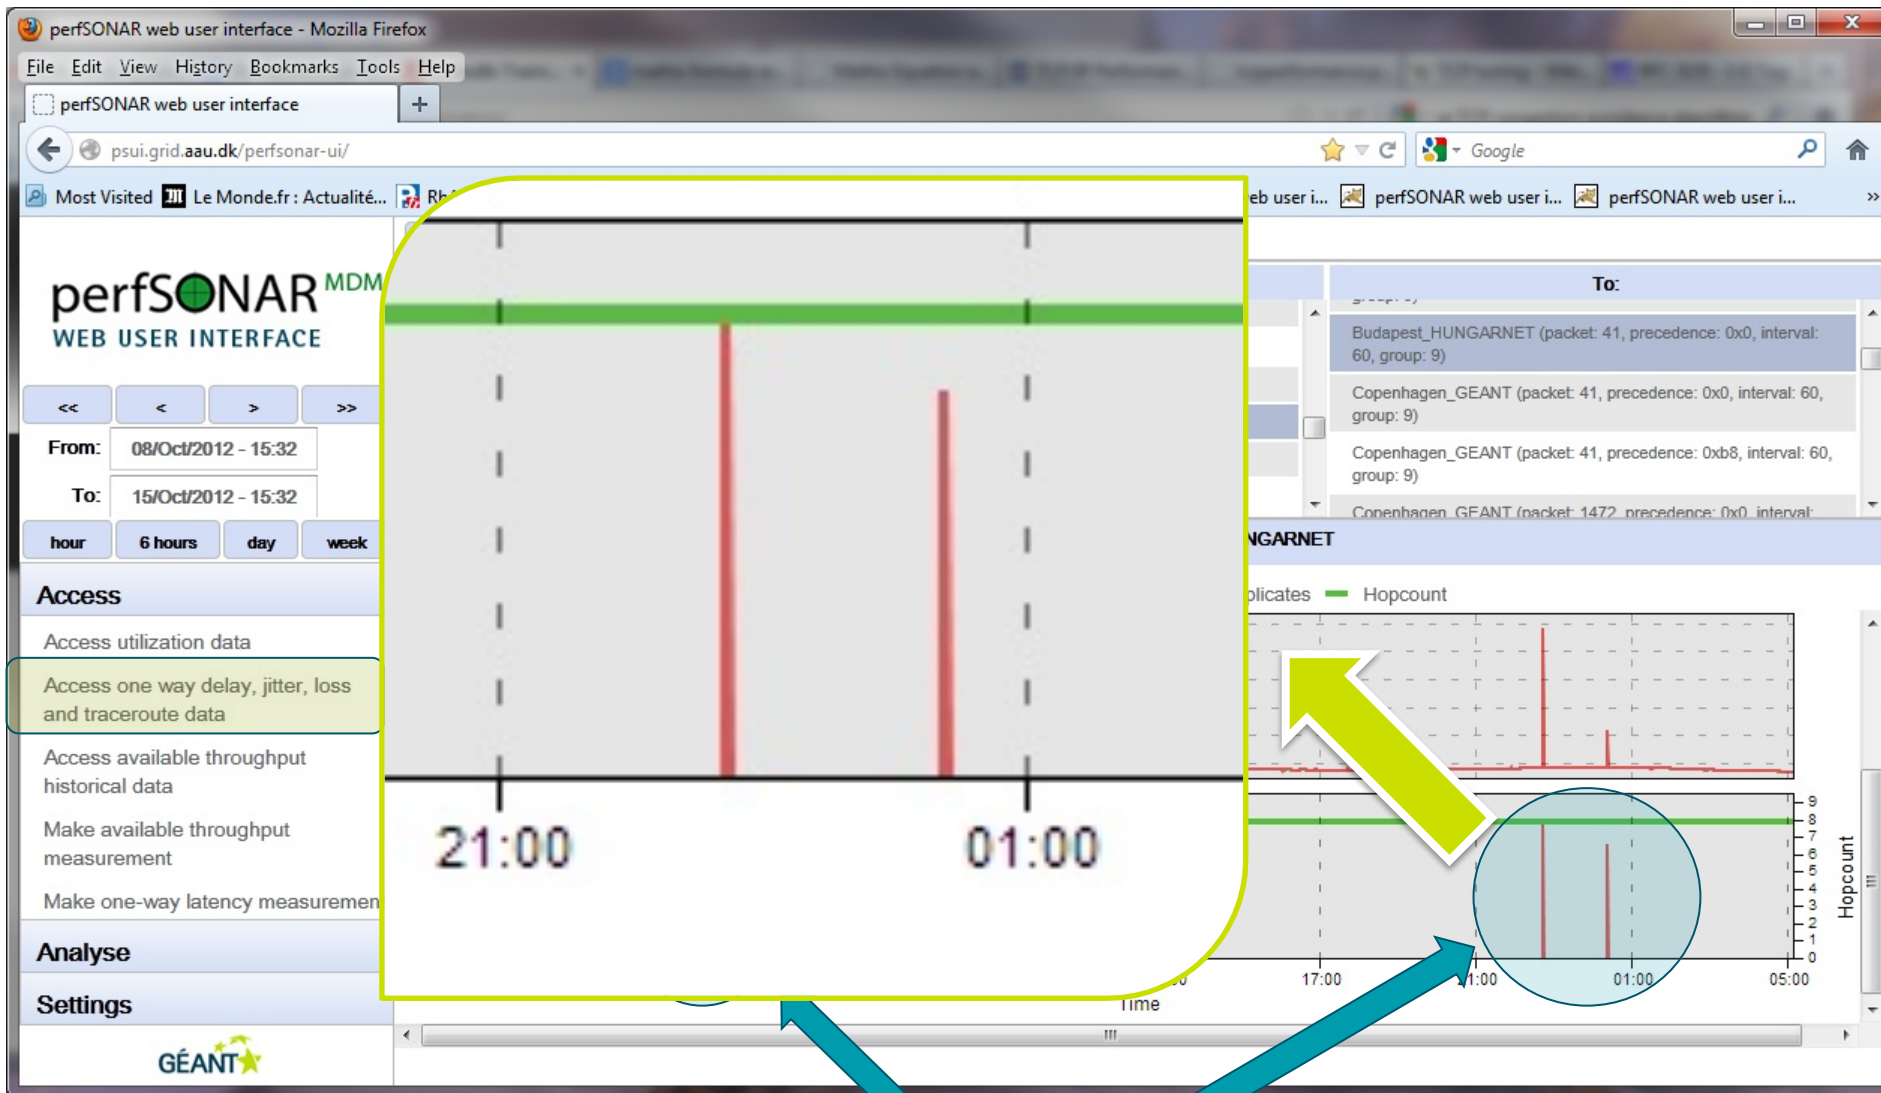
- **Strategy:**

- Sending 9 packets/minute between MPs (HADES)
- Sending 10 packets/second between MPs (OWAMP)
 - *Measure OWD, jitter, packet loss and tracking IP route*
- Store data into archives
 - *made accessible through web-service*

Web User Interface – OWD, jitter packet loss



Example of packet loss: from regularly scheduled measurements



perfSONAR

Look at here for packet loss

connect • communicate • collaborate

1-way delay on-demand (OWAMP)



perfSONAR MDM WEB USER INTERFACE

Pick source: OWAMP-Testing
Pick destination: Internet2-Seattle
Perform test
Swap endpoints

Packet count: 200
more options

From: 15/Oct/2012 - 03:47
To: 15/Oct/2012 - 15:47
hour 6 hours day week

Measurement summary	
Min packet delay	83.44 ms (0.08343899995088577 s)
Max packet delay	85.42 ms (0.08542399993166327 s)
Min TTL	243
Max TTL	243
Max est. error (±)	0.10 ms
Sent packets	200
Lost packets	0
Duplicate packets	0

Sequence number: 0 1 2 4 5 6 7 8 9 10 11 12 13

83.54 ms 83.49 ms 83.47 ms 83.47 ms 83.49 ms 83.52 ms 83.48 ms 83.51 ms 83.50 ms 83.51 ms 83.51 ms 83.50 ms 83.45 ms

Measurement summary	
Min packet delay	83.44 ms (0.08343899995088577 s)
Max packet delay	85.42 ms (0.08542399993166327 s)

Estimated error: 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms 0.10 ms

Web User Interface – route comparison



PerfSONAR UI (Web mock-up)

Service to query: Pick service Fetch data

Pick an endpoint pair to get measurements

From:	To:
FCCN_Aveiro	FCCN_Aveiro
FCCN_Coimbra	FCCN_Coimbra
FCCN_Lisbon_JRA1	FCCN_Porto
FCCN_Porto	GEANT_Budapest
GARR_Bari	GEANT_Geneva

Measurement ID : 9559
Packet size : 41
Precedence : 0x0
Interval : 60
Group size : 9

From: 25/May/2011 - 08:13
To: 25/May/2011 - 20:13

hour 6 hours day week month

Analyze path

Explore

- Explore RRD MA service
- Explore HADES MA service
- Explore BWCTL MA service
- Perform BWCTL test

Test

FCCN_Lisbon_JRA1 to GEANT_Budapest

Route 7

1	130.59.35.145	swiEZ2-G4-12.switch.ch
2	130.59.36.205	swiLS2-10GE-1-1.switch.ch
3	130.59.37.1	swiCE2-10GE-1-3.switch.ch
4	62.40.124.21	switch.rt1.gen.ch.geant2.net
5	62.40.122.3	fe1-vlan71.srv3.gen.ch.geant2.net

Route 2

1	212.191.227.33	UNKNOWN
2	62.40.124.181	pioner.rt1.poz.pl.geant2.net
3	62.40.122.61	so-7-1-0.rt1.fra.de.geant2.net
4	62.40.122.162	so-5-0-0.rt1.gen.ch.geant2.net
5	62.40.122.106	switch-lb2-gw.rt1.gen.ch.geant2.net
6	130.59.36.210	swiCE3-10GE-1-4.switch.ch
7	130.59.36.2	swiZH2-10GE-1-1.switch.ch
8	130.59.36.129	swiIX1-10GE-1-3.switch.ch
9	130.59.36.170	gn2-bwctl3-eth1.switch.ch

IPDV (ms)

Delay(ms)

Loss/Dups

Simple route comparison

Achievable bandwidth (BWCTL)



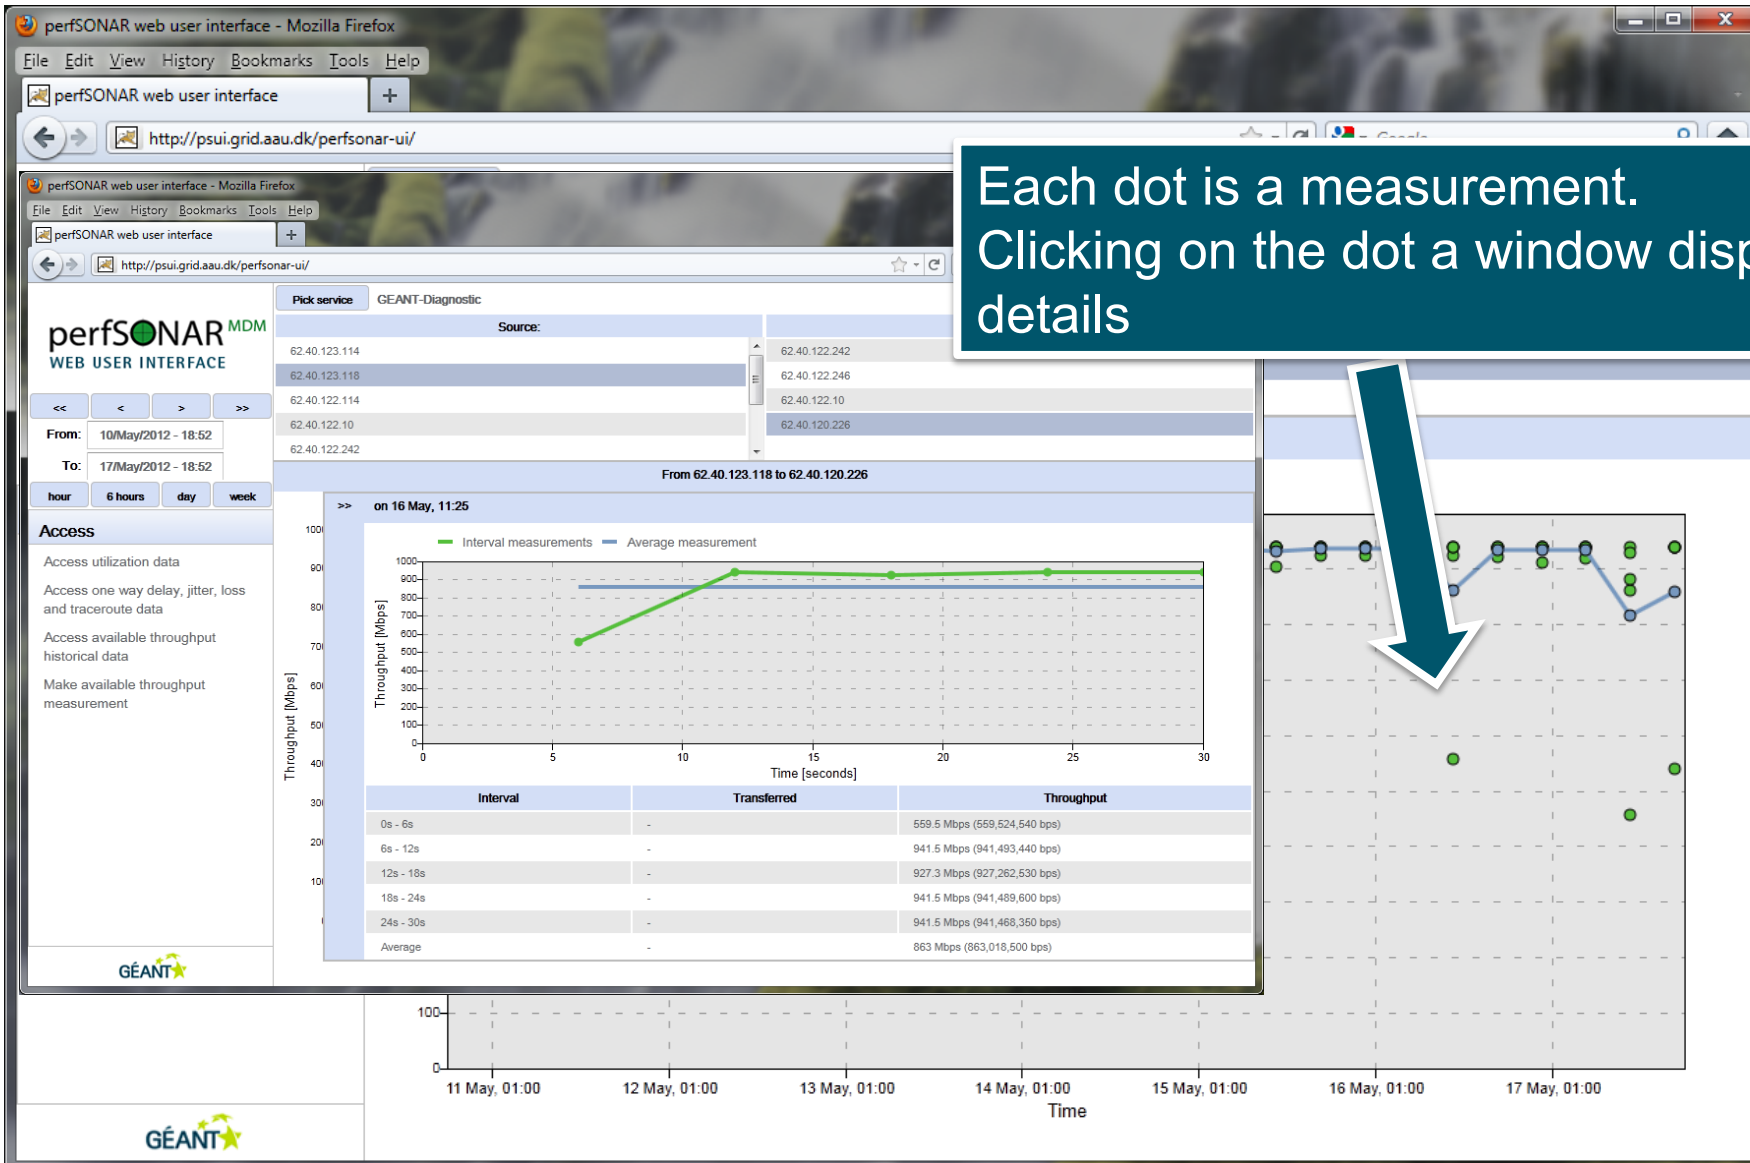
- **Purpose:**

- Measure the achievable bandwidth between two MPs
 - *Regularly scheduled and*
 - *on demand (only for NREN NOC/PERT engineers)*
- Provide access to historical measurements

- **Strategy:**

- Run bandwidth test between MPs using a web-service interface to BWCTL
- Display data with graph and store into perfSONAR SQL-MA archive
 - *made accessible through web-service*

Accessing Historical Bandwidth Measurements



Each dot is a measurement. Clicking on the dot a window displays the details

...and getting the results in two clicks from the web interface



Pick source GEANT_Geneva
Pick destination GEANT_Milano

Reporting interval (s): 6
Type of Service bits:

Graph

Interval measurements Average measurement

Interval	Transferred	Throughput
0s - 6s	612 MB (641,757,548 B)	855.7 Mbps (855,676,731 bps)
6s - 12s	611.8 MB (641,524,072 B)	855.4 Mbps (855,365,429 bps)
12s - 18s	559 MB (586,134,288 B)	781.5 Mbps (781,512,384 bps)
18s - 24s	610.9 MB (640,586,692 B)	854.1 Mbps (854,115,589 bps)
24s - 30s	605.6 MB (635,026,528 B)	846.7 Mbps (846,702,037 bps)
Average	2.9 GB (3,145,064,448 B)	838.6 Mbps (838,575,034 bps)

Textual output

Support for mobile devices

iPhone App (currently being developed)



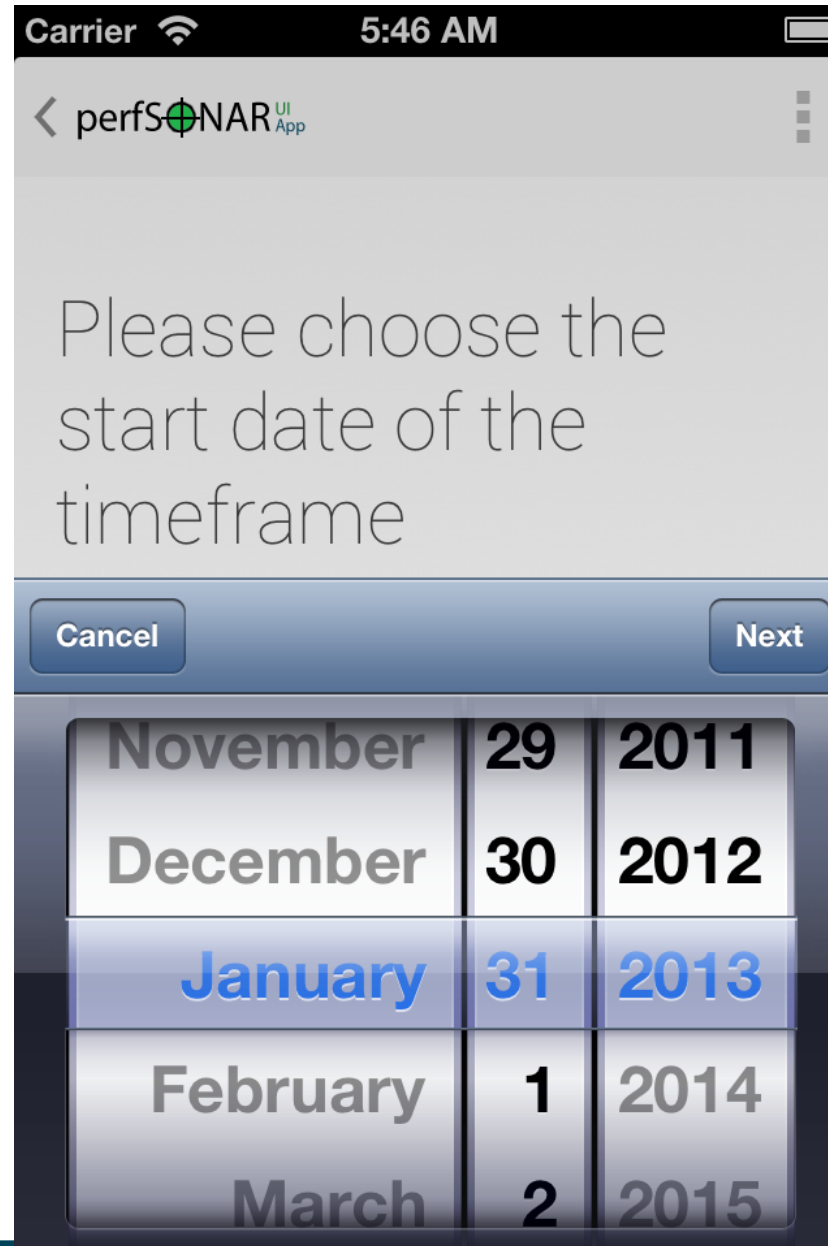
WELCOME




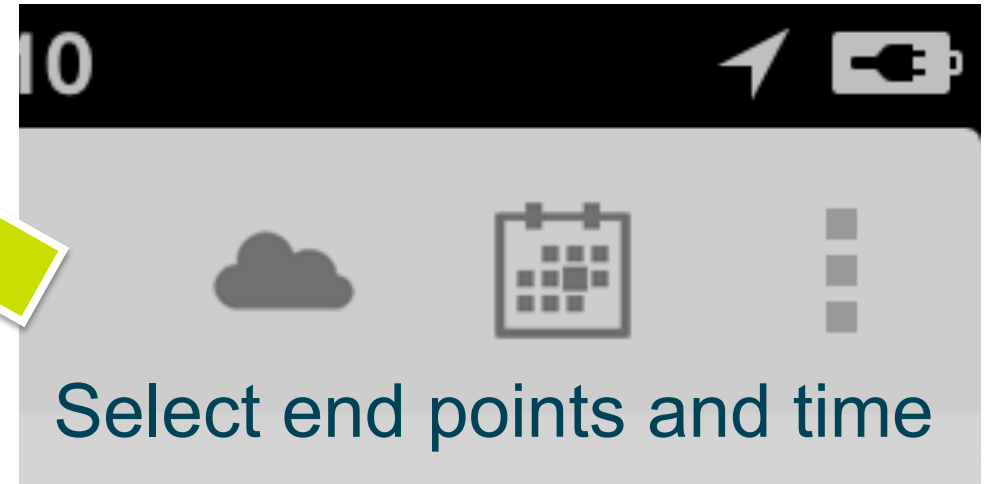
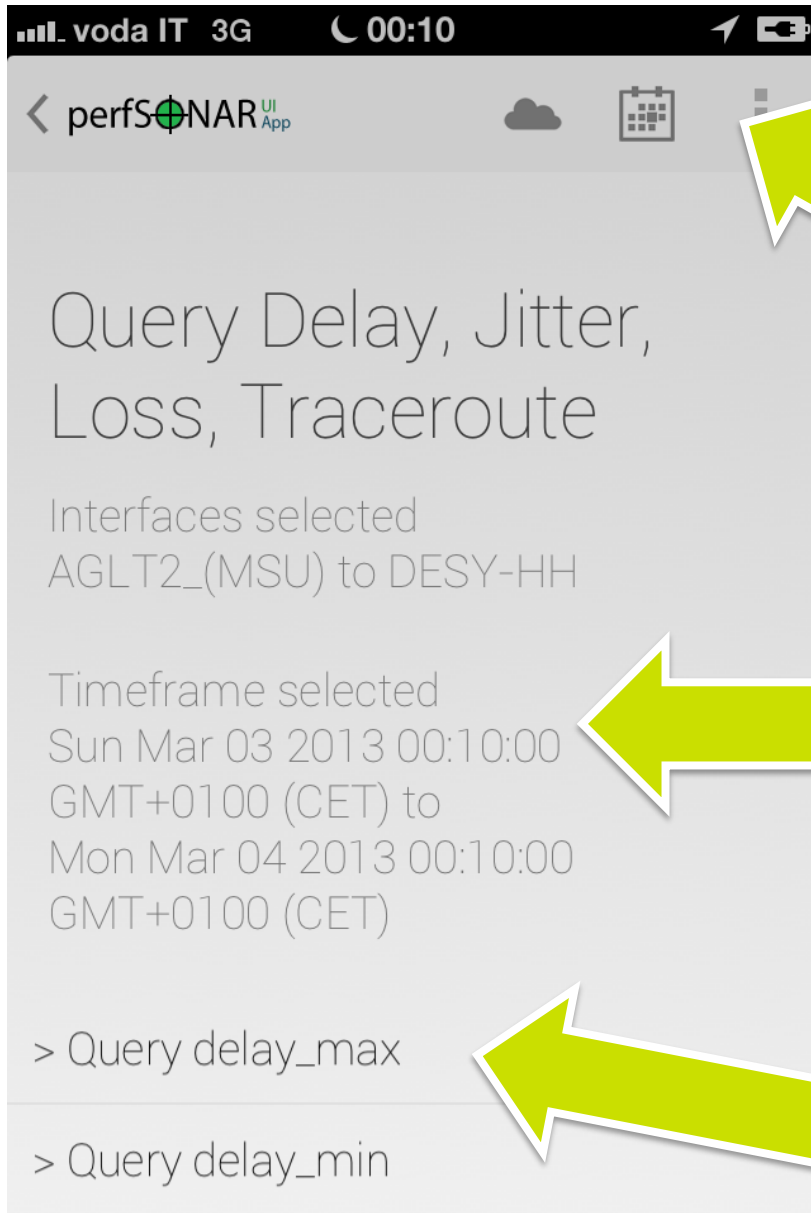
perfSONAR
powered



Native smartphone integration




Example: Query One-Way Delay



The result



voda IT 3G 00:11

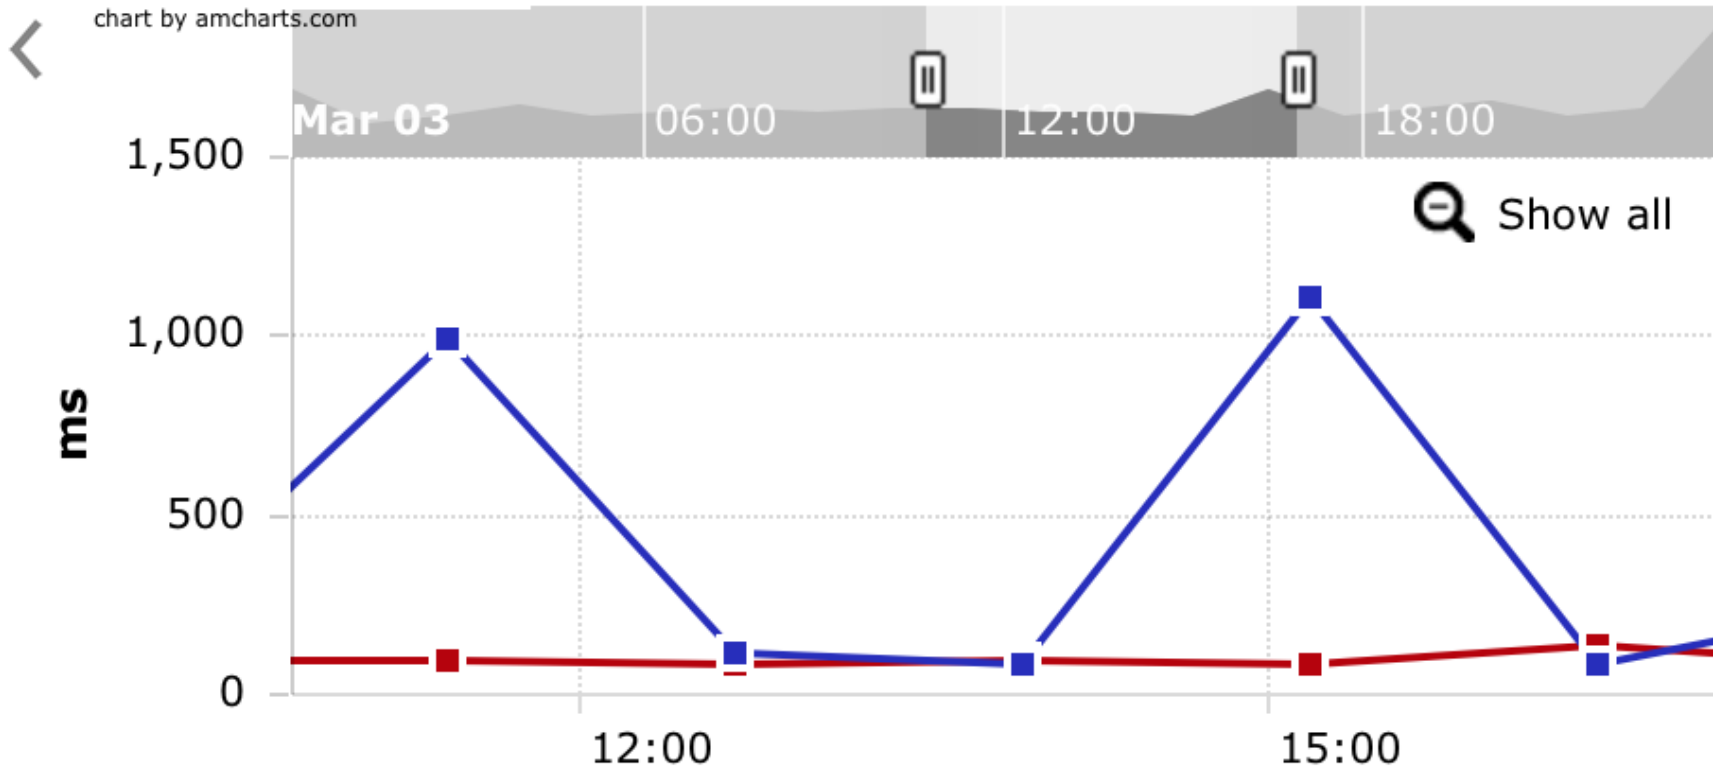


- AGLT2_(MSU) -> DESY-HH
- AGLT2_(MSU) <- DESY-HH

Zoom in/out



voda IT 3G 00:11




- AGLT2_(MSU) -> DESY-HH 90
- AGLT2_(MSU) <- DESY-HH 86

Detailed inspection



voda IT 3G 00:11



- AGLT2_(MSU) -> DESY-HH 92
- AGLT2_(MSU) <- DESY-HH 746

Weathermaps with perfSONAR

Several metrics to choose



WebCNM (cnmlx1.srv.lrz.de) - Mozilla Firefox (Private Browsing)

File Edit View History Bookmarks Tools Help

WebCNM (cnmlx1.srv.lrz.de)

cnmlx1.srv.lrz.de:4080/webtest/topo-dynamic/cgi-bin/cnm-devel.cgi?rm=start_webcnm_gwt&initial_map_id=2584&initial_metric_util&map_image_ty

Most Visited News Le Monde.fr : Actualité... Rhône - Le Progrès Live Traffic Conjugaison de tous l... Google Traduction

Map View: Topology of lhcopn abstract

To Parent/Prev Map map mode: OSM link metric: Hades one_way_delay max Open LinkList Cbx Help Map-related

- none
- BWCTL
- Hades duplicate packets
- Hades jitter
- Hades one_way_delay
- Hades packet loss
- IP hop
- IP Interface bandwidth
- IP Interface errors
- IP Interface utilization
- Status

link metric: Interface utilization

- none
- BWCTL
- Hades duplicate packets
- Hades jitter
- Hades one_way_delay
- Hades packet loss
- IP hop
- IP Interface bandwidth
- IP Interface errors
- IP Interface utilization
- Status

Chicago

An application: troubleshooting for the arts

e-Pormundos Afeto: A multi-stage dance performance



Remote dancers from Brazil

Stage: Barcelona



e-Pormundos Afeto: pictures from the performance

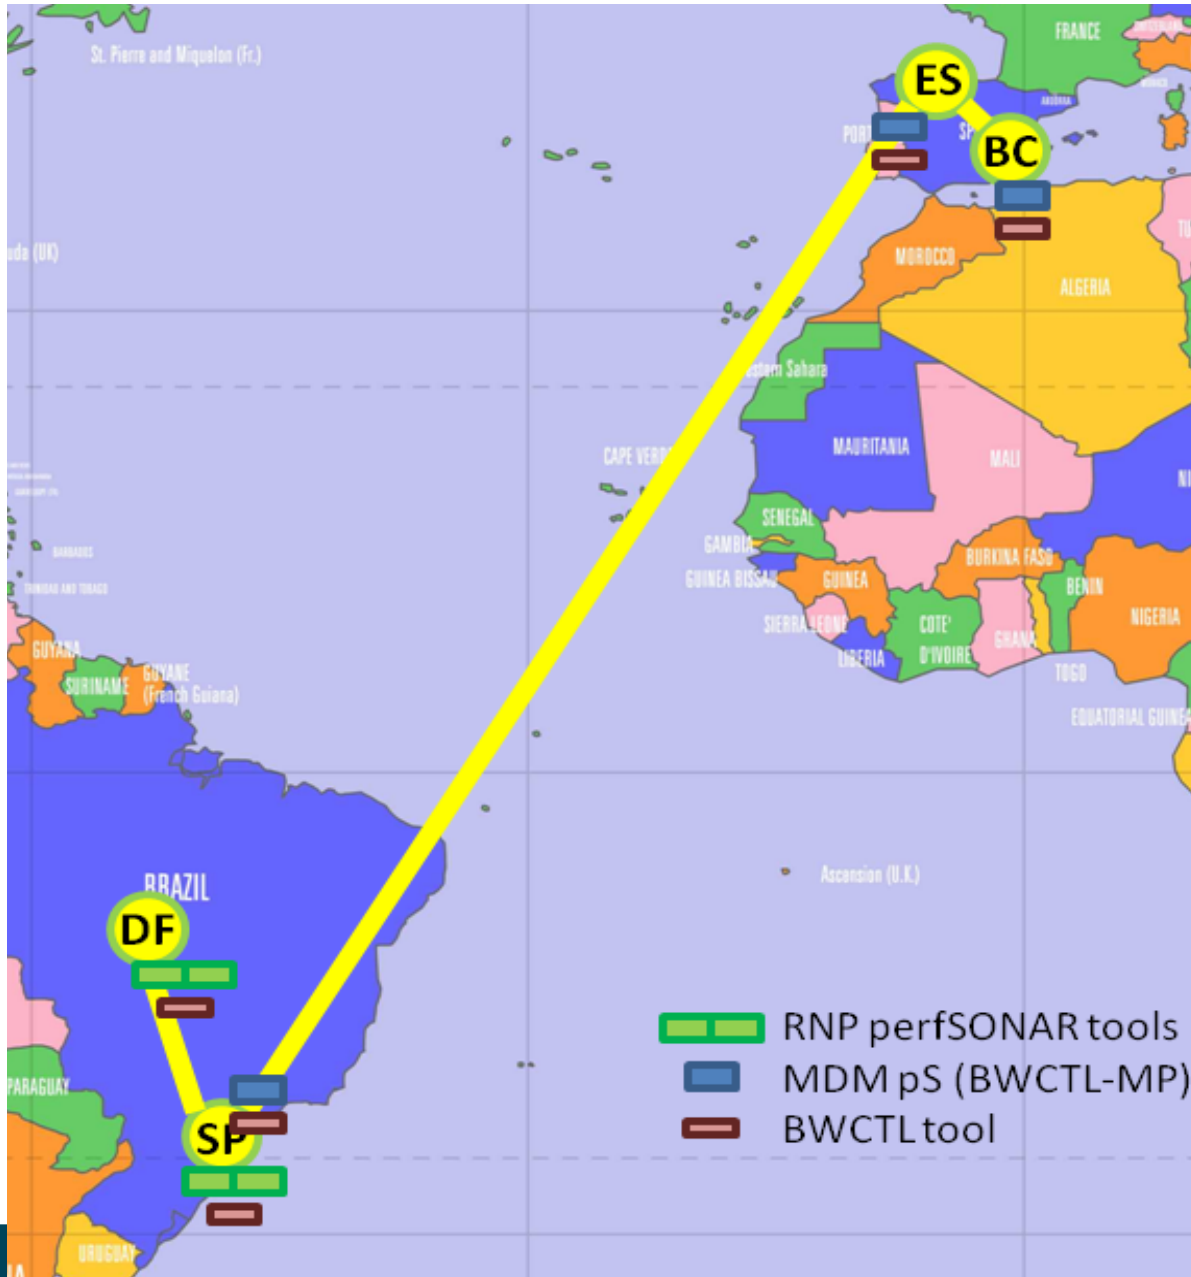


e-Pormundos Afeto performance: The network console



"e_pormundos afeto" (2011)
Concepção /Conception Ivani Santana
Co-criação/ co-creation Konic Thtr
Dançarinos/Dancers Mab Cardoso, Jean Therese (Brasil)/Sachiko Fullito
(Espanha)
Tecnologia: LAVID, RNP, i2Cat

Requirements and objectives

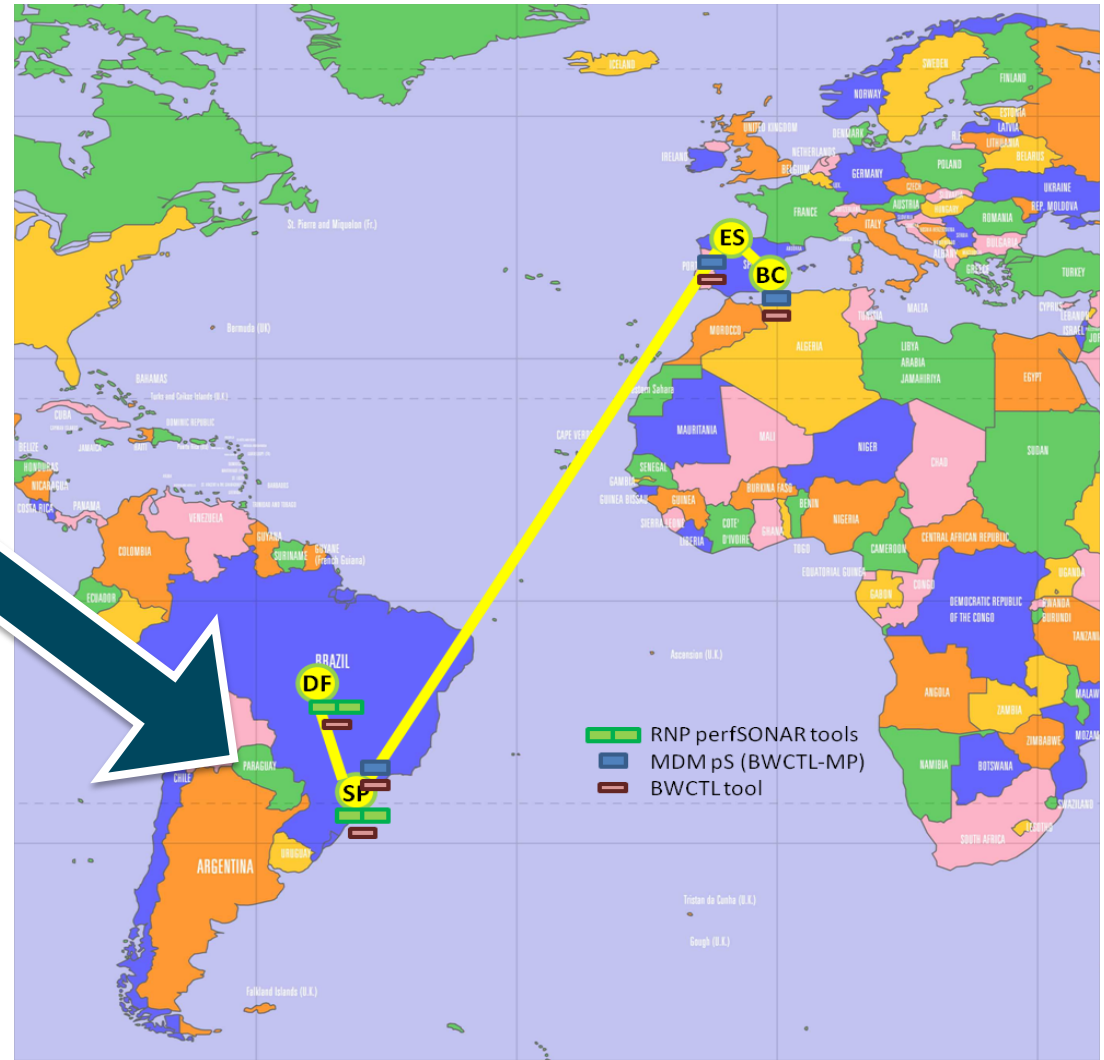



Traffic requirements:
25-30 Mbps
UDP traffic
Low latency
Low Jitter

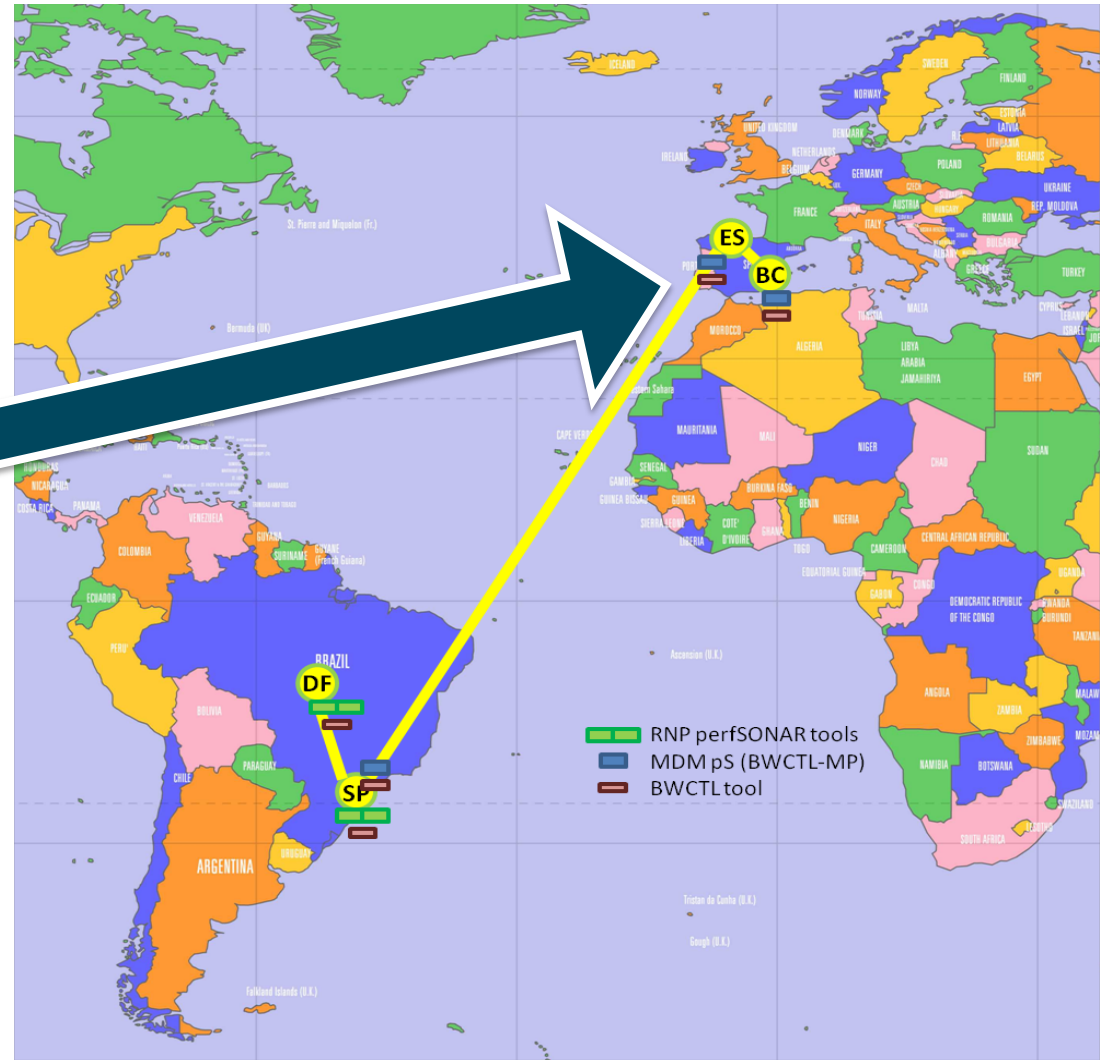

Objectives:
Validate connectivity
Find bottlenecks
Troubleshooting
Support

- Solution based on perfSONAR MDM for troubleshooting
- perfSONAR MDM provided:
 - Reliable multi-domain monitoring
 - *Common interface both in Europe and Brazil*
- The measurement and troubleshooting strategy:
 - Regular and on-demand bandwidth tests:

Latin American side



European side



perfSONAR Available MPs



- GEANT_Madrid
- RNP_Sao_Paulo_VM
- RNP_Brasilia_VM
- RNP_Sao_Paulo
- RNP_Brasilia
- REDIRIS_Madrid
- REDIRIS_Madrid_Delay
- ANELLA_Barcelona
- ANELLA_perfSONAR
- CIEMAT
- REDIRIS-BWCTL-MP



Select Measurement point to be queried

The list is loaded from C:\Fausto\Projects\Madrid Test\barcelona-event

Select MA(s)

Use	Measurement archive	URL	Schema	Status	Status explan...
<input type="checkbox"/>	GEANT_Madrid	http://mp1.mad.es.geant2.net:8090/services/MP/B...	bwctl	echo	
<input type="checkbox"/>	RNP_Sao_Paulo_VM	http://200.133.192.43:8090/services/MP/BWCTL	bwctl	echo	
<input type="checkbox"/>	RNP_Brasilia_VM	http://200.19.119.90:8090/services/MP/BWCTL	bwctl	echo	
<input type="checkbox"/>	RNP_Sao_Paulo	http://bwctl.pop-sp.rnp.br:8090/services/MP/BWCTL	bwctl_destination	org.apache.axis...	java.net.Unkno...
<input type="checkbox"/>	RNP_Brasilia	http://bwctl.pop-df.rnp.br:8090/services/MP/BWCTL	bwctl_destination	org.apache.axis...	java.net.Socket...
<input type="checkbox"/>	REDIRIS_Madrid	http://perfsonar.rediris.es:8090/services/MP/BWCTL	bwctl	echo	
<input type="checkbox"/>	REDIRIS_Madrid_Delay	http://perfsonar-delay.rediris.es:8090/services/MP/...	bwctl	echo	
<input type="checkbox"/>	ANELLA_Barcelona	http://84.88.16.114:8090/services/MP/BWCTL	bwctl	echo	
<input type="checkbox"/>	ANELLA_perfSONAR	http://perfsonar.anella.cesca.cat:8090/services/MP/...	bwctl	echo	
<input type="checkbox"/>	CIEMAT	http://perfbw.ciemat.es:8090/services/MP/BWCTL	bwctl	echo	
<input type="checkbox"/>	REDIRIS-BWCTL-MP	http://hades-bwctl.rediris.es:8090/services/MP/BWCTL	bwctl	echo	

URL: http://perfsonar.anella.cesca.cat:8090/services/MP/BWCTL

Name: ANELLA_perfSONAR Schema: bwctl

Event type: bwctl Supported event types: N/A


Service status: Load Save New Remove

.conf file Lookup Service

perfSONAR Results for UDP and TCP throughput




```
Interval 36.0-42.0; Transferred 35.053139 MB (36755880 Bytes); Throughput 49.00784 Mbps (49007840 bits/sec)
Interval 42.0-48.0; Transferred 35.793343 MB (37532040 Bytes); Throughput 50.04272 Mbps (50042720 bits/sec)
Interval 48.0-54.0; Transferred 35.088186 MB (36792630 Bytes); Throughput 49.05684 Mbps (49056840 bits/sec)
Interval 54.0-60.0; Transferred 35.794744 MB (37533510 Bytes); Throughput 50.04468 Mbps (50044680 bits/sec)
Average throughput 49.653116 Mbps (49653116 bits/sec)
```



GEANT_Madrid > RNP_Sao_Paulo
 GEANT_Madrid > RNP_Sao_Paulo

RNP_Sao_Paulo > REDIRIS_Madrid

perfSONAR results for packet loss: no lost packets



Lost packets

0

Duplicate packets

0

Conclusions




- Network data accessible
 - *from multiple domains*
 - *In a standardised format.*
- Effective collaboration
- Configuration problems solved
 - Quickly
 - Effectively
- Smooth dance performance
 - Happy dancers
 - Happy engineers 😊



perfSONAR

“

perfSONAR MDM enabled the engineering teams across the domains between Barcelona and Brazil end points to quickly pin-point possible network bottlenecks and ensure the network configuration was adequate to support the dance performance transmission for the TERENA Network Performing Arts Production Workshop in Barcelona.



Fausto Vetter, RNP Research & Development Coordinator

perfSONAR MDM website: <http://perfsonar.geant.net>



Goals:

- Single point of access for perfSONAR
- Contact points, FAQs, resources & downloads, and support
- **Host news and success stories from Users**

communicate • collaborate

perfSONAR Twitter: @perfSONARMDM



The screenshot shows the Twitter profile for @perfSONARMDM. The browser window title is "perfSONAR MDM (perfSONARMDM) on Twitter - Windows Internet Explorer". The address bar shows "http://twitter.com/#!/perfSONARMDM". The profile header includes the name "perfSONAR MDM", the handle "@perfSONARMDM", and the location "Cambridge". A bio states: "perfSONAR MDM (Multi-Domain-Monitoring) is the multi-domain monitoring service for the GÉANT Service Area (GSA) http://perfsonar.geant.net". There are 7 tweets, 1 following, 18 followers, and 0 listed. A yellow button says "Sign up & follow perfSONAR MDM". Below the bio, there are four tweets:

- Tweet 1: "#perfSONARMDM Correlation of existing measurements and new web-based user interface discussed with user representatives in Prague. 39 minutes ago"
- Tweet 2: "#perfSONARMDM First User Panel concluded: 18 participants from 9 NRENS discussed the future direction of perfSONAR. 41 minutes ago"
- Tweet 3: "#perfSONARMDM First perfSONAR User Panel next week in Prague. Visualisation tools and circuit mon. among the topics http://bit.ly/knukeQ 11 May"
- Tweet 4: "New perfSONAR website was announced in GÉANT newsletter 'Connect': http://bit.ly/eks0VL #perfSONARMDM 7 Apr"

The right sidebar contains a list of links: About, Help, Blog, Mobile, Status, Jobs, Terms, Privacy, Advertisers, Businesses, Media, Developers, Resources, and © 2011 Twitter.

Weekly tweets

Messages re-tweeted by other sister networks and organisations

Growing community of followers around the world

**@perfSONAR
MDM**

act • communicate • collaborate

perfSONAR MDM. Be part of it.



Follow perfSONAR at:

<http://twitter.com/#!/perfSONARMDM>

- Website: <http://perfsonar.geant.net>
- Twitter: @perfSONARMDM
- Info: domenico.vicinanza@dante.net

perfSONAR

connect • communicate • collaborate

Thanks for the collaboration and support from...



- GÉANT:

- DANTE:



- Domenico Vicinanza
 - Fausto Vetter
 - Thomas Fryer

- DFN:



- Roland Karch

- redIRIS:



- Alberto Escolano Sanchez

- Anella/CESCA:



- Jordi Gaya

- RNP (Rede Ipê):



- Daniela Brauner
 - Leonardo Carneiro
 - Murilo Vetter
 - Rafael Costa
 - José Augusto Suruagy Monteiro



**Extra slides for questions about
impact of packet loss on bandwidth**

- TCP performances depends heavily on packet loss
- Because of the way TCP works, waiting for acks
 - Losing packets rapidly leads to performance degradation
- Mathis formula:

$$\text{Max rate} < MSS / RTT * \sqrt{p}$$

- where:
 - Max Rate: is the TCP transfer rate in bps
 - MSS: is the maximum segment size (e.g.1460) in bytes
 - RTT: is the round trip time in seconds
 - p: is the packet loss rate or loss probability (fraction)

Formula from Mathis, et al. *The macroscopic behavior of the TCP congestion avoidance algorithm*. CCR, 27 (3), July 1997

- The following examples assume $MSS=1460$ Bytes
- Example 1:
 - $RTT= 88$ ms
 - Loss probability: 1%
→ Max BW=**1.33 MB/s**
- Example 2:
 - $RTT= 20$ ms
 - Loss probability: 0.1% → Max BW=**18.5 MB/s**
- Example 3:
 - $RTT= 20$ ms
 - Loss probability: $1E-6$ → Max BW=**5.8 GB/s**

Measuring packet loss with perfSONAR MDM



- The packet loss rate is then a crucial measurement to assess the health of a TCP/IP connection
- perfSONAR MDM supports packet loss measurements from the perfSONAR web user interface
 - Regularly scheduled
 - *every minute*
 - On demand
 - *New feature introduced in the last version*
 - *Available at <http://psui.geant.net>*