STRATEGIC TECH FOR INTERNET

2004-2005 FINAL REPORT

WARD NO. 0231044

AMPATH

UTH

NGF OFFICE DP. CVDERINFRASTRUCTURE (OCI).

> Collaborative Research & Education Operational and Functional Support

Julio Ibarra Principal Investigator Heidi Alvarez Co-Principal Investigator





NSF Award - #0231844

STI: AMPATH Collaborative Research and Education Operational and Functional Support

2002-2005

Final Report

Program Managér Kévin L. Thompson OCI Office of CyberInfrastructure O/D OFFICE OF THE DIRECTOR

Start Daté October 1, 2002 Expirés Séptémber 30, 2005 Awardéd Amount to Daté \$482679 Invéstigator(s) Julio Ibarra julio@fiu.édu(Principal Invéstigator) Héidi Alvaréz (Co-Principal Invéstigator)

> Sponsor Florida International University 11200 SW 8TH ST MIAMI, FL 33199 305/554-2000

NSF Program(s) STRATEGIC TECH FOR INTERNET

This page intentionally left blank

Table of Contents

1. Introduction	1
1.1 Background	1
1.2 Goals and strategies	1
1.3 Value to U.S. science	1
2. AMPATH Organization and Managem	ent2
3. Operations	
3.1 NETWORK INFRASTRUCTURE	
3.2 BUDGET	
4. Milestones	4
5. Progress Reports	7
5.1 Year 3 Work Plan Progress – July, 20 5.1.1 Infrastructure	04 through October, 20057
5.1.2 Outreach	7
5.1.3 Community Building	
6. Merit of funded activity and broader in	1pacts
6.1 INFRASTRUCTURE	9
6.2 E-SCIENCE APPLICATION SUPPO	PRT9
Appendix A: Papers and Presentations, Ir	ternational Meetings and Conferences10
Appendix B: AMPATH Projects and App	lications12
B.1 High Energy Physics	
B.2 Physical and Organic Chemistry	
B.3 Remote Biological, Marine, and Atm	ospheric Sensing12
B.4 Biodiversity	
B.5 Materials Science	
B.6 Unique Facilities and Science Collabo	rations13
B.7 Video/Audio Communications	
B.8 Remote Operations	
B.9 Virtual Laboratories/Observatories	
B.10 Distributed Archives and Libraries	
B.11 Distributed Processing	

Appendix C: AMPATH Performance Statistics	16
C.1 Real World Demonstrations	19
C.2 Observation of AMPATH Annual Traffic	22
Appendix D: AMPATH-STI REU Report: Distributed Network Monitoring 1. Introduction	5 25
2. Anatomy of an international exchange point	
3. NetFlow and Flowtools	27
4. MonALISA and ApMon	
5. Integrating MonALISA and FlowTools via ApMon	
6. Case Study– Monitoring Peer Traffic	
8. Discussions and Future works	

Table of Figures

FIGURE 1 AMPATH NETWORK ARCHITECTURE AND CONNECTIVITY	3
FIGURE 2 AMPATH IN 2004 BEFORE ONS 15454 INSTALLATION	
FIGURE 3 AMPATH NETWORK AFTER CISCO ONS 15454 INSTALLATION	19
FIGURE 4 SC04 BANDWIDTH CHALLENGE - BRAZIL'S PARTICIPATION	
FIGURE 5 SC04 BANDWIDTH CHALLENGE NETWORK DIAGRAM LED BY CALTECH	21
FIGURE 1: AMPATH SCHEMATIC ARCHITECTURE	27

Table of Tables

TABLE 1 AMPATH MANAGEMENT TEAM	2
TABLE 2 AMPATH-STI MILESTONES	6
TABLE 3 NETWORKS AND CONNECTIONS AT THE START OF FINAL YEAR	16
TABLE 4 NETWORKS AND CONNECTIONS AT THE END OF THE FINAL YEAR	17

1. Introduction

1.1 Background

AMPATH (AMéricasPATH) bégan in March, 2000 at Florida Intérnational Univérsity (FIU), in Miami. Its principal mission is to providé high-pérformancé réséarch and éducational nétworking in Latin América and the Caribbéan in the support of scientific collaboration and éducation. By August 2001, AMPATH had alréady dévélopéd coopérativé agrééments with its industrial partnérs. Thésé agrééments includéd a thréé-yéar donation of tén DS3 (45 Mbps) circuits by Global Crossing valuéd at \$25M. Other industrial partnérs by that time includéd Cisco Systems, Lucént Téchnologiés, Junipér Nétworks and Térrémark Worldwidé. With thésé asséts, AM-PATH was ablé to connéct two régional Réséarch & Education Nétworks (RENs): Chilé's RE-UNA and Brazil's RNP2. By April 2002, AMPATH connéctéd two moré RENs: ANSP of Brazil and RETINA of Argéntina, as wéll as thé Arécibo Obsérvatory, thé Univérsity of Puérte Rice, and thé Gémini South, NOAO and CTIO obsérvatoriés in Chip. In August of 2003, thé réséarch and éducation nétwork of Vénézuéla was connéctéd at thé DS-3 lévél. Panama connéctéd at thé DS-3 lévél in 2004. Thésé achiévéments aré thé work of a véry small AMPATH staff. Thé staffing at présent includés just sévén individuals who amount to léss than four FTEs.

1.2 Goals and strategies

AMPATH sééks not only to providé high-spééd connectivity bétwéén réséarchers and students throughout North, Céntral and South América. AMPATH also sééks to foster éfféctivé usé of such connectivity for scientific and scholastic purposés, éspécially thosé of intérest to the U.S. With this latter goal in mind, AMPATH has sought to (1) identify connectivity applications of intérest to U.S. science and enginéering réséarch and éducation, and (2) foster collaborations bétwéén U.S. scientists and théir countérparts in the service area.

1.3 Value to U.S. science

Although AMPATH sééks to stréngthén sciéncé and éducation in its sérvicé aréa per se, it also sééks to énhancé U.S. sciéncé and to fostér closér réséarch and éducational tiés bétwéén thé U.S. and thé sérvicé aréa. In this sénsé, AMPATH providés bénéfits in its sérvicé aréa similar to thosé alréady réalizéd by high-spééd NSF-sponsoréd connéctivity to Europé and thé Asia-Pacific Rim réséarch nétworks. A numbér of U.S. sciéncé initiativés dépénd critically upon facilitiés or énvironments locatéd in Latin América. Oné éxamplé is obsérvational astronomy. Astronomical obsérvatoriés locatéd, or to bé locatéd, in thé Caribbéan and South América includé Arécibe Obsérvatory, Piérré Augér, thé Gémini South, CTIO and NOAO optical téléscopés and thé Atacama Largé Millimétér Array (thé lattér two in Chilé). Anothér éxamplé is thé Intér-Américan Instituté for Global Changé Réséarch (IAI). This intérgovérnméntal organization coordinatés réséarch into énvironméntal and socio-économic changé in thé Américas, and it counts 17 mémbér countriés in thé Latin América aréa as well as thé U.S. and Canada. Also, NASA's Intérnational Spacé Station (ISS) project sééks to providé accéss to thé ISS for sciéntific invéstigators worldwidé, including thosé in thé Latin América. All of thésé U.S.-léd initiativés now dépénd or will dépend crucially upon high-spéed connectivity between the U.S. and the Latin América. Séveral fédéral agéncies currently opératé nétworks in Latin América using point-to-point lowbandwidth circuits. The AMPATH project provides a coordinated and effective approach to these connectivity néeds.

2. AMPATH Organization and Management

The AMPATH management team, noted below, is chaired by Julio Ibarra, of FIU, and includes representation from all collaborators. This committee makes policy and application support plans.

Chile	
	REUNA: Florencio Utreras
	Connected June 2001
Argentina	
U	RETINA: Carlos Frank
	Connected December, 2001
Brazil	,,
	RNP2: Nelson Simoes (Rio)
	Connected July 2001
	FAPESP: Luis Lopez (Sao Paulo)
	Connected March 2002
Panama	Guy Cordoba
	Connected: June 2004
Venezuela	
	CNTI: Jorge Berrizbeitia
	Connected April 2002
	Connecteu April 2003

 Table 1 AMPATH Management Team

3. Operations

3.1 NETWORK INFRASTRUCTURE

The AMPATH network architecture, as shown in Figure 1, provides access to the wide area network provided over an ATM based Internet Protocol (IP) enabled network. Each national REN connects to AMPATH via this architecture. All AMPATH traffic routed to other international networks flows through the NAP of the Americas located in Miami, Florida. From Miami, the traffic is routed to either Abilene or STARTAP/StarLight. Abilene is the US-based research and education network, and STARTAP/StarLight is the ATM/optical exchange point for international traffic, located in Chicago.

AMPATH offers a variety of <u>network services</u>. Among the major services are:

ATM and Optical Ethérnét pééring fabrics Intra-régional pééring ovér Layér 2 sérvicés, including IP VPNs Nativé IPv6 Multicast capablé End-to-end performance measurement and monitoring AMPATH Reflector VRVS server for Video over IP Flow based and QoS based monitoring using NetFlow tools capabilities NOC Services through the Global Research NOC at Indiana University



Figure 1 AMPATH Network Architecture and Connectivity

The figure above presents the basic network architecture employed by AMPATH. This architecture is representative of the research and educational networks worldwide. It also depicts the current AMPATH connectivity to the national RENs. As can be seen, the only interface between AMPATH and the rest of the world is through Miami to STARTAP/StarLight and Abilene. Through the STARTAP and Star Light interfaces in Chicago, virtually any location in the world is accessible to the national RENs. This interface is extremely important to the success of AM-PATH in that it allows truly international collaboration to be conducted.

3.2 BUDGET

The operations of AMPATH are supported through NSF STI Award 6231844, Florida International University, Global Crossing and contributions from the management team.

Annual Costs:	
Abilene Connectivity (622 Mb/s)	\$186,666
Local Loops in Miami	\$177,666
Network Operations Center (Nap Of the Americas)	\$ 24,666
Equipment Mainténance & Réplacement	\$168,888
Senior and Technical Staff	\$185,666
Network Operations Services (Indiana University)	\$ 25,666
Administrative Staff	\$ 61,888
Travel	\$ 126,666
Misc.	\$ 28,666

Tetal: \$968,666

5 DS-3 Circuits	\$4,288,888
	Provided by Global Crossing

Cost Sharing:	
Chile	\$126,666
Argentina	\$ 78,888
Brazil	\$126,666
Venezuela	\$126,666
NSF STI Award	\$146,893
Panama	\$ 65,888
Florida International University	\$329,533

4. Milestones

2000		
April	AMPATH progress report at Brazil's national research networking conference, in	
	Bello Horizonte, Brazil	
August	AMPATH status report at the NLANR Joint Techs meeting in Toronto, Canada	
●ct	AMPATH status report at Internet2 Member Meeting, Atlanta, Georgia	
2001		
June	Chile's Red Universitaria Nacional (REUNA) NREN connects to AMPATH; peer-	
	ing established with Abilene	
June	AMPATH status report at the STARTAP Annual Meeting, INET 2001, Stockholm,	
	Sweden	
July	Brazil's Rédé Nacionalé de Pésquisa (RNP) NREN connects to AMPATH; pééring	
	established with Abiline	

August 15-	First NSF AMPATH Workshop for Identifying Areas of Scientific Collaboration	
17	between the US and the Latin America, ANI-0123388	
August	University of Puerto Rico connects to AMPATH	
August 23-	AMPATH Status Report at the HPIIS Workshop on Developing metrics to quantify	
24	the use of HPIIS network links, San Diego Supercomputing Center	
December	Argentina's Red Teleinformatica Académica (RETINA) connects to AMPATH;	
	peering established with Abilene	
2002		
March	Brazil's Academic Network of Sao Paolo (ANSP) connects to AMPATH; peering	
	established with Abilene	
April	Gemini South connects to AMPATH; peering established with Abilene	
April 12	First NSF International AMPATH Conference in Valdevia, Chilé, ANI-0220176	
May	AMPATH status report at International Task Force, Internet2 Member Meeting, Ar-	
	lington, Virginia	
May	AMPATH status report at the Global Research Networking Summit, Brussels, Bel-	
	gium	
June	AMPATH establishes presence at the NAP of the Americas	
September	AMPATH enables native IPv6 and offers IPv6 peering to members	
October	NSF Strategic Technologies for the Internet Award for AMPATH Collaborative	
	Research and Education Operational and Functional Support, ANI-\$231844	
October	AMPATH status report at International Task Force, Internet2 Member Meeting, Los	
	Angeles, California	
December	University of Puerto upgrades to an OC3c its connection to AMPATH	
2003		
January	AMPATH upgrades to an OC12c its connection to Abilene.	
January	NSF AMPATH Workshop: Fostering Collaborations and Next Generation Infra-	
	structure, ANI-0305876.	
January	AMPATH status report at Internet2 Techs/ESCC Workshop in Miami, Florida,	
	hosted at FIU	
April	Vénézuela's Centre Nacional de Tecnologias de Informacion (CNTI) connects to	
	AMPATH; peering to Abilene established	
April	AMPATH status report at International Task Force, Internet2 Member Meeting, Ar-	
	lington, Virginia	
April	AMPATH participates in REUNA's conference on research networking, in Chile	
May	AMPATH collaborates with US High-Energy Physics community to submit a pro-	
	posal to the NSF Experimental Infrastructure Network program (proposal not	
	funded)	
May	AMPATH participates in RNP's Brazilian Symposium on Computer Networks, in	
	Natal, Brazil	
June	AMPATH status report at the ACURIL conference in Puerto Rico.	
June	Plan developed to connect the University of the Virgin Islands and the NRAO radio	

	antenna on St. Croix to AMPATH	
September	AMPATH member meeting in Buenes Aires, Argentina	
September	NSF awards FIU to dévélop an intér-régional Grid-énabled Center for High-Energy	
	Physics Research Education and Outreach, which includes partial funding for AM-	
	PATH to establish an STM-4 to Rio de Janeiro	
October	AMPATH status report at International Task Force, Internet2 Member Meeting, In-	
	dianapolis, Indiana	
October	Présented a paper on the Digital Divide in Latin América for the Digital Divide	
	Round Table, at the International Center for Theoretical Physics (ICTP), in Trieste Italy	
2004		
February	Established formal partnership with CENIC and NLR to provide North-South con-	
	nectivity in the Western Hemisphere	
February	Presented AMPATH STI project at HeGrid workshop in Rio	
March	Participated in 1st International Grid Networking Workshop in Geneva	
April	Presented at Jet Road Workshop AMPATH measurement activities	
May	Participated in REACCIUN 2 seminar in Venezuela	
June	Panama connected	
July	Presented STI project to NSF at HIPIIS review	
August	Begun working with Global Crossing to facilitate CLARA tender	
September	Established partnership with CENIC and NLR to collaborate on connectivity be-	
	tween the U.S. and Latin America	
October	Held PASI workshop in Argentina	
Nevember	Entered into an MOU with CLARA donating remaining AMPATH circuits	
December	Begun working with FAPESP to establish long-term connectivitiy goals	
2005		
January	RNP's three year circuit expires. Traffic is routed through Europe	
March	Chile's three year circuit expires. Chile negotiates extension	
April	Venezuela launches new high speed network	
May	Sao Paulo transitions from three year DS-3 on Global Crossing to LANautilas	
July	Chilé's connectivity ends, and is routed through Europe on RedClara	
August	RNP upgraded Rio to Sao Paulo connectivity	
September	RédCLARA péers with U.S. in California, all other STI supported countries estab-	
	lished long-term plan past end of STI support.	
October	Project terminates without any countries disconnected. Long-term plans in place or	
	being developed. Options are much better for the entire region than three years ago.	

Table 2 AMPATH-STI Milestones

5. Progress Reports

Thé following séctions détail progréss against the NSF ANIR STI project miléstonés. Thé STI proposal listéd miléstoné activitiés organizéd by Infrastructuré, é-Sciéncé Application Support and Outréach, shown in thé proposal, in Appendix D of this réport.

5.1 Year 3 Work Plan Progress – July, 2004 through October, 2005

5.1.1 Infrastructure

Ongoing: Dévélop rélationships with moré intérnational submarine cable carrièrs that offer service to other countries not serviced by Global Crossing.

ACTION: Collaborated with Global Crossing to respond to the ALICE EC program for regional networking, offering all current and present assets to assist in the establishment of a regional network.

Ongoing: Connect the remaining NRENs or académic networks in Service Area countries where Global Crossing has a presence.

- ACTION Focused on inter-régional connectivity in future activities, and in this period supported the creation of régional, and state sponsored direct connections
- ACTION: Established MOU with CLARA and donated 3 unrestricted three year leases of DS-3s for use as they please throughout the region.

Analyzé and plan for improvéd or new AMPATH connectivity to other international connect points such as StarLight

ACTION: Begun collaboration with SURA to start the AtlanticWave.

5.1.2 Outreach

Préparé for the next AMPATH Mééting in Argéntina for nécessary infrastructure support; vidéo conférences with AMPATH participants and Advisory board mémbér to plan rélévant agénda itéms

ACTION: Conference and working groups meetings held.

Hold a Workshop on leveraging IPV6

ACTION: Participaté in the TIDIA mééting in Sao Paulo addréssing IPV6 and associated téchnologiés.

A Workshop to promote US-ASIAN-LATIN AMERICAN cooperation

ACTION: Held a CANS conférence in Miami

A PASI conference on Grid Computing

ACTION: Held a PASI conference in Argentina

CIARA Application Support Workshops

ACTION: Participating in REN workshops.

5.1.3 Community Building

Organizé external advisory committée and identify its purpose; recruit subject matter experts.

ACTION: None taken.

Schedule presentations to US and International agencies, including the Inter-American Development Bank, on the cultural impact of AMPATH, to increase in-country support for NREN's

ACTION: Active in both the Digital Divide initiative and the European Commission's @LIS, working with U.S. State Department and US AID to ensure policies in sync with U.S. government.

AMPATH présence at global nétworking conférences (such as INET), démonstrating AM-PATH's commitment to global collaborations in scientific réséarch and éducation

ACTION: Participation increased (please reference timeline)

Schédulé démonstrations of intércontinental collaborativé réséarch and applications at national conférencés, such as Intérnét2 Mémbér Méétings

ACTION: INET présentation. AMPATH has facilitated démonstrations by Brazil at Super Computing, IGrid, and Intérnet2 Mémber mééting évents.

Develop and publish the AMPATH Annual Report

ACTION: This document.

6. Merit of funded activity and broader impacts.

The proposal's goal was to contribute to the stability of U.S.-Latin American networking, in order to increase the rate of discovery for U.S. scientist. This goal required a predictable infrastructure, support of existing applications and outreach to bring in more collaborators for the benefit of U.S. science. These three topics are addressed separately below.

6.1 INFRASTRUCTURE

The infrastructure program is a very dynamic space. U.S. investigators need both more broad network access (to regions like the Caribbean) and more robust networking to collaborators and instruments in Brazil and Chile. Before the project begun, there was no advanced (greater than é-1) connectivity from Latin America to the U.S. Now there are direct connections and a regional aggrégation. Many challenges face the regional aggrégations, most significantly the development of domestic networking and last mile challenges. Scientists and broad education use are in Latin America more at odds than elsewhere, as populist agendas often overstep the needs of advancing the state (subsequently out of step with U.S. needs)

6.2 E-SCIENCE APPLICATION SUPPORT

The AMPATH team has reflectively engaged a broad scientific community to determine what is and is not working with application support. For this research came the CIARA project. The CI-ARA project is a model that FIU, the University of Puerto Rico (UPR), the Universities of Sao Paulo and Venezuela are pursuing to increase the rate of discovery for investigators and improve graduate education. The impact of CIARA is systemic. Participants have institutionalized the initiative, and collaborators from Caracas to Sao Paulo are developing local versions. This model effectively provides a rising tide that has lifted all research effectiveness across the region. The overarching benefit for the U.S. is that collaborators and instruments are more accessible and more effectively used with the technology transfer that the CIARA model provides.

6.3 OUTREACH

The network needs of U.S. investigators are growing rapidly in Latin America. The desire to collaborate on issues such as biodiversity and material science is exceeding the existing partnerships, and the scope of AMPATH. Through this STI grant science and education have been demonstratively advanced through new efforts in Physics, Astronomy, BioDiversity, and Computer Science research. The scope of benefit is specifically manifest in the 7 workshops held and indefinitely augmented with the creation of new communities of practice, with new discovery and new partnerships formed for the advancement of science and education in concert with the NSF's mission.

Appendix A: Papers and Presentations, International Meetings and Conferences

2003, June

AMPATH status report at the ACURIL conference in Puerto Rico.

2003, August

Opérational Plan for Upgrading the Intérnét Connectivity of the United States Virgin Islands (USVI): University of the Virgin Islands (UVI), St. Croix VLBA Radio Astronomy Station and Other Réséarch and Education Facilities

2003, September

AMPATH member meeting in Buenos Aires, Argentina.

2003, August 25

An In-Depth Look at AMPATH StarLight Traffic Accounting and Participant Analysis A Réport for the Stratégic Technologies for the Internet (STI): AMPATH Collaborative Réséarch and Education Opérational and Functional Support Réséarch Expérience for Undérgraduatés, Award #331112

2003, October

AMPATH status report at International Task Force, Internet2 Member Meeting, Indianapolis, Indiana.

2003, October 23-24

Expériénces with the Digital Dividé in Latin América Round Tablé on Dévéloping Countries Access to Scientific Knowledge, The Abdus Salam ICTP, Trieste, Italy

2003, December 8-9

Rolé of Sciénce in the Information Society CERN Généva, Switzerland 2004, Fébruary Participated in the Nét@Edu to ensure broad understanding of AMPATH project

2004, February

Demonstrated to HEP community the AMPATH project at teh HeGRID workshop in Rio

2004, March

Participated in 1st International Grid Networking Meeting

2**664**, May

Engaged Fielding Scholars in CIARA model for AMPATH

2004, Juné Présentéd at the workshop for the Virtual Data Grid Laboratory 2004, July Présentéd STI project to the NSF.

2004, September Internet2 Members meeting

2004 December RNP Annual meeting presentation GEMINI IT planning meeting in Miami

2005, March Digital Divide meeting at CERN

2005, July Présented at CLARA mééting in México

2005, August Présented at HEP mééting, DAEGU

Appendix B: AMPATH Projects and Applications

B.1 High Energy Physics

The multi-TeV energy scales needed to advance the study of the nature of matter and its most basic interactions, and to search for new particles and forces, has led to unprecedented challenges in petabyte data access and analysis. This has led in turn to worldwide scientific collaborations where each nation contributes its share to the construction of the experiment, and to the resources needed to analyze the data, while participating in the global process of search and discovery. The Large Hadron Collider (LHC) program at CERN in Geneva, encompasses four major experiments searching for the Higgs particles thought to be responsible for mass, as well as the states of matter and violation of symmetries that existed in the early moments of the universe. The US is a major player in the two largest of these experiments, ATLAS and CMS, with approximately 400 US physicists and engineers (20% of the total) involved in each one.

The CERN/LHC program is a large, but not atypical, example of a collaboration-driven project. The US has expended significant resources in this research area. The principal experimental devices are located in Europe, with major US contributions, and the Computational Data Grids that are being developed will be worldwide. Nevertheless, the program asks questions on such a grand scale that more than 5,000 individual researchers (in four large collaborations) are involved in working with various aspects of the problem. The Parallel and Distributed Processing Group of the Federal University of Rio Grande du Sul, in collaboration with its physics research staff, is an important player in the development of the Grid processing and data handling techniques necessary for the success of the LHC.

This collaboration is possible today in South America because of AMPATH.

B.2 Physical and Organic Chemistry

The University of Utah and the University of Buenos Aires are engaged in collaboration, funded jointly by the NSF and several Argentine agencies, to model chemical shifts in complex organic crystals. The collaboration itself is rooted in special expertise available in Utah and Buenos Aires. One of the objectives of the study is to develop robust techniques to include intermolecular effects in the calculations. Subsequently these techniques will be applied to solve structural problems in biologically active compounds from marine invertebrates from the South Atlantic in collaboration with researchers in the organic chemistry group and Buenos Aires.

B.3 Remote Biological, Marine, and Atmospheric Sensing

Although somewhat influenced by geography, the University of Puerto Rico's remote sensing program represents a collaboration of a different kind one across many disciplines. This program acquires large, generally geographically-based, data sets from remote ocean sensing devices, earth-sensing satellites, and a variety of other sources. These data are then applied to problems in oceanography, marine biology, and a variety of related fields.

B.4 Biodiversity

Another broad collaboration dépendant on AMPATH, based at the University of Buenos Aires and sponsored by IAI, is examining the role of biodiversity and climate in functioning ecosystems. The group is studying the interrelationships between the atmospheric composition, climate, land use, biodiversity, and the ecosystem as a whole and how that impacts on the provision of human services. This is a global program in terms of the systems studied and it relies on a collaboration between several senior scientists in six South and North American countries. Two of the participating researchers are in the United States, one in Mexico, and several more are in South America.

B.5 Materials Science

The Institute Balseire in Argentina is pursuing a program aimed at understanding the thermodynamic, phase stability, and transformation of complex material systems. This effort involves a broad collaboration involving research centers in Florida, France, and Sweden. It includes an important distance-learning component directed at the upper-division and graduate level.

B.6 Unique Facilities and Science Collaborations

Whether sponsored by the NSF, DOE, NASA, or other public or private agencies, today's major US science projects take on economic scales that innately require international participation through partnerships and other forms of collaboration. Thus, for many of these programs the interests of US science and policy are inexorably linked to the international communications infrastructure because the science is being done with internationally shared resources.

Many of the presentations spoke on actual, planned, or desired collaborations to share unique, one-of-a-kind or few-of-a-kind, research tools. There were many examples of in-place programs. The CMS experiment at CERN is a 144-institution partnership across 31 countries, and the ATLAS experiment is a collaboration of similar size and scope. The Gemini observatory is a seven-nation program, managed by the US (NSF), with three South American partners including Chile's CONICYT. The International Space Station

is a 15-nation partnership. The UCD program consists of 17 partner nations in the Americas.

A number of other examples of currently active, planned, or desired programs had the same flavor. ALMA is a major international cooperative project involving the US, Canada, Japan, Chile, and the European Southern Observatory which, itself, has nine member countries. FIU has a center for electron microscopy that can be operated remotely from any part of the world and would like to establish a partnership to fund and operate a remotely accessible high-temperature and high-pressure physics laboratory. NASA/Ames would like to establish a similar facility for the analysis of a wide variety of remote ecological, marine, and atmospheric sensor data, as would the University of Puerto Rice with oceanography data.

In many of these cases, the unique facility includes large expensive pieces of equipment, such as telescopes or particle accelerators. In every case, there were one or more remotely accessed shared databases or data archives. AMPATH is the connection to a number of these instruments today.

B.7 Video/Audio Communications

Vidéoconférencing, both for direct science collaboration and for operations purposes including data gathering, is becoming ubiquitous across AMPATH. Vidéo over IP and Access Grid applications are becoming increasing common. The LHC experiments use the VRVS system for vidéo/IP, including many méetings per day, throughout the year, as an integral part of the collaborations' daily work.

Thèsé téchnologiés aré éspécially uséful adjuncts to programs involving rémoté opérations or rémoté usérs, such as ALMA, Gémini, the University of Utah crystal modéling project with the University of Buénos Airés, and the Intérnational Space Station to méntion a féw. Somé programs with multiplé sités also méntionéd intégrating voice over IP into their téléphoné plants, as a convénience and cost-saving méasuré.

While some the systems use are already well-advanced, these will be developed into integrated environments for multi-site collaborative work; either on their own or in the context of Data Grids.

B.8 Remote Operations

Users through out the US manipulate and control observational or experimental equipment connected to AMPATH. Some described this as remote operation, and others as virtual operation. In any case, there is a lot of cross coupling between remote data taking, virtual laboratories, and archival data analysis.

The Arecibe and Gemini observatories both make use of observing by scientists at remote sites as an alternative to astronomers having to travel to the telescope. It is intrinsic to Gemini operations, and they plan to expand the capability significantly, to include additional remote observing centers in Arizona, Florida, and Chile. Located at 16,000 feet, the ALMA observatory will have to be operated in this mode.

The NASA/Amés Ecosystèm Computer Facility not only is engaged in the analysis of remote sensing data, put has a suite of analysis instrumentation that it plans to enable to allow remote analysis of remotely sensed data. The Utah/Buénes Airés crystal project, the FIU study of the Andéan rivers, and the International Space Station effort all require the remote manipulation of sensor equipment. NASA/Marshall has already developed a remote payload control packed for International Space Station PIs.

B.9 Virtual Laboratories/Observatories

This might be defined as doing remote science from centralized or distributed archives, libraries, and databases, using standardized suites of access and processing tools made for the purpose. One example of such tools is the FIU/NASA Regional Applications Center's Terrafly package that allows the overlay and moving-map visualization of multiple geographically-parsed data sets.

The Caltech/CERN/University of Buenos Aires Large Hadron Collider (LHC) collaboration expects to make extensive use of this approach. ALMA, Gemini, and Arecibo expect to participate in a collaboration called the National Virtual Observatory to permit the mining and analysis of various astronomy data sets with standardized tools.

B.10 Distributed Archives and Libraries

The CERN group will need to transport tens of Petabytes by 2006. The Space Station, ALMA, Gemini, and the Utah/Buenos Aires crystal programs all have multiple international archive sites.

The University of Puerto Rico's Tropical Center for Earth and Space Studies' satellite down-link facility maintains a large digital database of the tropical Western Atlantic Region that is accessed online by NASA and other federal agencies, universities, and the private sector. Ecological programs such as the FIU rivers project and other global change research projects like the Buenos Aires biodiversity and climate study all intimately rely on access to large data sets maintained in many different places.

The IAI global change data and information service project currently has two access nodes (Brazil and Uruguay) and expects to expand to 18 in the near future using AMPATH. This system uses a distributed network of libraries and databases accessed with a Yahoo like search tool that yields relevant metadata and can provide ftp access if required.

B.11 Distributed Processing

The Federal University of the Rio Grande in Brazil has several programs that require distributed processing, including research in that subject itself, biological cell modeling efforts, 3D thermodynamic modeling, and high energy physics problems (with Caltech/CERN). The Buenos Aires/LHC collaboration will take part in the GriPhyN network-based data-grid program, and will be part of the International Virtual Data Grid Laboratory (iVDGL) being supported by NSF. The FIU rivers project expects to use grid techniques in doing soil and fluids models. The Utah/Buenos Aires project expects to use grid processing in its crystals program. Gémini and ALMA expect that grid approaches will become necessary in the future to handle complex image processing with large data sets.

Appendix C: AMPATH Performance Statistics

For the 2884-2885 timeframe, AMPATH serviced the following NRENs and South Florida GigaPOP members with the following connection speeds. Table 3 shows the networks and their connection characteristics at the start of the final year (July 2884); Table 4 shows the networks and their connections at the conclusion of the year (September 2885).

NREN	Location	Connection Speed
REUNA	Santiago, Chile	DS3 (45 Mbps)
AURA/Gemini	La Serena, Chile	DS3 (45 Mbps)
REACCIUN/CNTI	Caracas, Venezuela	DS3 (45 Mbps)
ANSP	Sao Paulo, Brazil	DS3 (45 Mbps)
RNP	Rio de Janeiro, Brazi	1 DS3 (45 Mbps)
RETINA	B. A., Argentina	DS3 (45 Mbps)
UPR	San Juan, Puerte Rice	●C3 (155 Mbps)
FIU	Miami, FL	GigE (1000 Mbps)
NWS	Miami, FL	GigE (1000 Mbps)

Table 3 Networks and connections at the start of final year

NREN	Location	Connection Speed
AURA/Gemini	La Serena, Chile	DS3 (45 Mbps)
REACCIUN/CNTI	Caracas, Venezuela	DS3 (45 Mbps)
ANSP	Sao Paulo, Brazil	STM-4 (622 Mbps)
RETINA	B. A., Argentina	DS3 (45 Mbps)
UPR	San Juan, Puerte Ric	• •C3 (155 Meps)
FIU	Mi z mi, FL	GigE (1888 Maps)
NWS	Miami, FL	GigE (1888 Maps)

Table 4 Networks and connections at the end of the final year

Thé first phasé of thé RédCLARA régional nétwork impléméntation of thé ALICE project, startéd opération in thé last quartér of 2004. Chilé's NREN, REUNA, connéctéd to RédCLARA, shortly aftér. Brazil's NREN, RNP, latér in 2004 connéctéd to RédCLARA. Whén thésé movés wéré madé from AMPATH to RédCLARA, traffic déstinéd to U.S. R&E nétworks startéd being routéd to Europé to transit GEANT, thén Abiléné.

AMPATH during this period also brought into production a new circuit to Abilene's core in Atlanta via FLR's newly established optical network. AMPATH now connects to Abilene via an FLR provided Layer 2 MPLS VPN (compliant with IETF Draft-Martini Layer 2 VPN RFC 3031) which terminates in Atlanta at the Abilene collocation space at SoX. This connection replaced the connection AMPATH had to Quest for connectivity to Abilene. Connected Latin American NRENs and R&E networks are able to continue directly peer with Abilene and the U.S. Federal Networks, with no loss of functionality.

In the period covered by this report AMPATH also established a connection with FLR via a 10 Gigabit Ethermet cross-connect to FLR's collocation space at the NAP of the Americas. This now enables AMPATH's downstream peers to access any State of Florida higher education institution at a significantly faster speed than previously possible; once again enabling collaboration between researchers.

Another significant achievement was the establishment of a 10 Gigabit Ethernet peering with UltraLight; this peering is achieved again through the leveraging of FLR's optical infrastructure; using a dedicated light path through FLR's network which terminates at AMPATH's Ethernet fabric in the form of a 10 GigE port on the Foundry FastIron switch. This newly implemented connectivity will serve as a key step in providing better access to High Energy Physics research projects and again enable collaboration between Physics research scientists which operate in AMPATH's downstream peers.

Connectivity to AMPATH's downstream NRENs and SFGP members was achieved through a mix of Ethernet and ATM circuits supported by AMPATH's core Ethernet fabric in the form of a Foundry FastIron 400 and AMPATH's core ATM fabric in the form of a Lucent CBX 500 ATM switch. The following figure shows the AMPATH network during the 2004 and early 2005 period.



Figure 2 AMPATH in 2004 before ONS 15454 installation

The DS3 link to the Academic Network of Sao Paulo (ANSP) was upgraded to an STM-4 (622Mbps) to support the network and cylerinfrastructure requirements of the NSF Center for High-Energy Physics Research Education and Outreach (CHEPREO) project, award #0312038. During the 2004-2005 period, AMPATH successfully installed and put into production a Cisco ONS 15454 optical switch, collocated at the NAP of the Americas in downtown Miami, FL. The introduction of this new network element allowed for the migration from ATM to SONET/SDH of the ANSP circuit to Sao Paulo, Brazil. This equipment set the stage to leverage resources from other projects to increase bandwidth capacity to Sao Paulo to support U.S. Science and Engineering projects of other NSF activities.

The following figure illustrates the state of AMPATH after the installation of the Cisco ONS 15454 and the ONS's interconnectivity to existing AMPATH network elements:



Figure 3 AMPATH network after Cisco ONS 15454 installation

C.1 Real World Demonstrations

In collaboration with the State of Sao Paulo, the CHEPRE® project upgraded the link to AM-PATH in Miami from a DS3 (45 Mbps) to an STM-4 (622 Mbps). This upgrade was done in time for Brazil and Latin America to participate in the Super Computing 2004 (SC04) event. The upgraded link to Sao Paulo was contracted on STM-16 (2.5 Gbps) ports, giving the CHEPRE® project the ability to increase the bandwidth capacity for an event, such as SC04, to the full port capacity. For the first time, Brazil and Latin America, were able to participate in a global demonstration with bandwidth of such magnitude. The SC04 Bandwidth Challenge provided the opportunity to form a global collaboration for a global demonstration powered by highperformance networks.

The CHEPREO link allowed High-energy physicists at the University of Sao Paulo to participate in the SC04 Bandwidth Challenge. For the first time, physicists were not constrained by bandwidth to run events and do very large file transfers to CERN and FermiLab. After the event, the physicists continued to use the link capacity while it remained available. See Figure 2 for a diagram of the network for the SC04 Bandwidth Challenge.



Figure 4 SC04 Bandwidth Challenge - Brazil's participation

During the bandwidth challenge, a new record was set for a Latin American research network with throughputs up to 2.95 Gbps between the conference Center floor in Pittsburgh and the computational clusters in Brazil. The traffic generated on the link between Sao-Paulo and the SC show floor (see Figure 3 below). This was a first big step in enabling physics groups in Sao Paulo, and also Rio de Janeiro (through the use of the GIGA project) to take part fully in the LHC and FNAL physics programs.



Figure 5 SC04 Bandwidth Challenge network diagram led by Caltech

C.2 Observation of AMPATH Annual Traffic

The following figures of the Multi-Router Traffic Grapher (MRTG) graphs show the average utilization on several of the international links that connect countries participating in AMPATH. The point of reference for the in bound and out bound flows of the MRTG graphs is the AM-PATH Cisco GSR 12012 core router.



Figure 6 - RETINA Utilization

Figuré 6 shows the traffic to and from the NREN of Argentina, RETINA. The Pierre Auger observatory has been using a minimal amount of the bandwidth on the AMPATH connection. Traffic over the AMPATH link has remained about the same for year. Argentina continues to utilize the AMPATH connection to reach the U.S.



Figure 7 - AURA/Gemini Utilization

Gémini South, CTIO, NOAO and other téléscopés in Chilé havé béén incréasing their link utilization to Abiléné (Figuré 7). This is to bé expéctéd as the téléscopés imprové éfficiency of observations using quéué observing, résulting in faster production of science data products that aré transferréd to répositories at partner sités.



Figure 8 - UPR Utilization

The University of Puerto Rico system and the Arecibo observatory are primary users of this link. Usage on the link (Figure 8) has been steady throughout the year with increases towards the latter months.



Figure 9 - CNTI Utilization

Vénézuéla's NREN has béén in formation and working to connect universitiés to the backboné as the goal for the launch of their national Internét2 project. An increase in network usagé is évident towards the latter part of the year (Figure 9), as additional universities were connected. Vénézuéla is also connected to RédCLARA.

Appendix D: AMPATH-STI REU Report: Distributed Network Monitoring

Anatomy of an International Exchange Point: Distributed Network Monitoring Using MonALISA and NetFlow¹

Jose Fernandez, Ernesto Rubi, Xun Su, Iosif Legrand, Heidi Alvarez, Julio Ibarra²

Abstract: In this paper we present a distributed network monitoring system, which exploits MonALISA (Monitoring Agents in A Large Integrated Services Architecture), a distributed web service delivery infrastructure designed to collect and process the network monitoring information. We augment the capability of MonALISA with FlowTools, the popular NetFlow data analysis toolset. We demonstrate how to integrate MonALISA and Flowtools via an UDP-listening agent ApMon, and highlight a case study of AMPATH, an international exchanging point located in Miami and serving a number of South American National Research and Education Networks (NRENS). Our experience showcases the elegant design philosophy of a scalable distributed service deployment platform coupled with the open-source traffic analysis tools and its impact on the daily operation of the production networks.

Keywords: NetFlow, MonALISA, Flowtools, Network monitoring.

1. Introduction

As the Internet expands both in its scope, reach and capacity, it becomes evident that there is a strong need to devélop a distributéd nétwork monitoring infrastructuré that can bé scaléd to support various nétwork topology, traffic granularity and user applications. NetFlow[1] is a widely deployed router-based traffic monitoring mechanism. FlowTools[2] is a open-source NetFlow analysis toolset underlying the data gathering and analysis infrastructure of our project. It is our main motivation to effectively use NetFlow to gain crucial understanding of the traffic charactéristics of the networks we opératé. In particular, we are intérested in understanding how to exploit the kéy advantages and avoid drawbacks of NetFlow-based traffic analysis by augmenting it with a distributed service-deployment platform. Indeed the focus of our work is to integrate MonALISA[3], a distributed monitoring system based on JINI/JAVA and WSDL/SOAP technologies. MonALISA's flexibility as a framework to gather, store and distribute network data collected was crucial to the success of our investigation and it shall become apparent throughout the course of this paper. The MonALISA framework provides a distributed monitoring service that not only is closely intégrated with our monitoring and data distribution philosophy but also acts as a dynamic service system. The goal is to provide the monitoring information from large and distributed systems in a flexible and self-describing way as part of a loosely coupled service architectural model to perform effective resource utilization in large, heterogeneous distributed centers. A salient feature of the MonALISA framework lies in its capability to integrate existing monitoring tools and procedures to collect parameters related to computational nodes, storage devices and network performance. A critical part to the research that we undertook is to integrate the parsed NetFlow data into a MonALISA which would make available the information in a distributed manner through mobile service agents. We will incorporaté réal-time as well as historical information in our system to improve the understanding of the traffic statistics data from the networks being monitored.

¹This work is supported by National Science Foundation's Strategic Technologies for the Internet (STI): Research Experience for Undergraduates Award No. 331112 and a Cisco Systems University Research Program (URP) grant.

² Xun Su (<u>xsu@hép.czltéch.édu</u>) and Iosif Légrand (<u>Iosif.Légrand@cérn.ch</u>) aré with High Energy Physics Départment, California Instituté of Téchnology, Pasadéna, CA. José Hérnandéz (joséf@fiu.édu), Ernésto Rubi (<u>érnésto@cs.fiu.édu</u>), Héidi Alvaréz (<u>héidi@fiu.édu</u>) and Julio Ibarra (<u>Julio@fiu.édu</u>) aré with Céntér for Intérnét Augmentéd Réséarch & Asséssment, Florida Intérnational University, Miami, FL.

A significant part of the paper is dedicated to the use of our system to analyze network traffic behavior across a realworld production network, AMPATH. The AMericasPATH (AMPATH) network is an FIU project sponsored in part by the US National Science Foundation CISE directorate, in collaboration with Global Crossing and other télecommunications product and service providers. Using Global Crossing's terrestrial and submarine optical-fiber networks, AMPATH is interconnecting the research and education networks in South and Central America and the Caribbéan to US and non-US réséarch and éducation nétworks via Intérnét2's Abiléné nétwork. Thé purposé of thé AMPATH project is to allow participating countries to contribute to the research and development of applications for the advancement of Internet technologies. The mission of AMPATH is to serve as the pathway for Research and Education networking in the Americas and to the world and to be the International Exchange Point for Latin America and the Caribbean research and education networks. Additionally AMPATH fosters collaboration for educational outreach to underserved populations both in the US and abroad. The AMPATH pathway serves as the bridge bétwéén Céntral and South Américan National Réséarch Nétworks (NRENs) and the world's réséarch and éducation networks. With the multiplicity of complex networked systems and educational activities served by AMPATH's wide-ranging infrastructure a strong demand for high availability and engineering collaboration arises, which is met through the use of various monitoring agents to provide a strong factual foundation to troubleshooting. Likewise, déciphéring the évéryday activities of our péers is achieved with a distributed approach to data gathéring and dissemination.

2. Anatomy of an international exchange point

Figuré 1 démonstratés AMPATH's current désign, IP addréssés as wéll as Autonomous Systém Numbérs (ASNs) aré omittéd. Wé will use simple NREN namés to identify our intérnational péérs.³ It is this nétwork that sérvéd as the backdrop for our study. Two coré routers éxist: a Cisco GSR 12612 as wéll as a Junipér M10, both routers havé NétFlow accounting énabléd and aré désignéd to export this data to a colléction workstation.

³ For more detailed information please visit <u>http://mrtg.ampath.net</u>.



Figure 6: AMPATH schematic architecture.

3. NetFlow and Flowtools

NétFlow was originally dévélopéd as a switching path and today it is primarily uséd for nétwork accounting. Flow récords aré génératéd and exportéd by a routér. Each flow récord contains information about all packéts that aré catégorizéd to havé the samé combination of the following IP héadér fiélds:

Source IP address Destination IP address Source port Destination port Layer 3 protocal type TOS byte Input logical interface

NétFlow récords only unidiréctional traffic in-bound to any intérfacé on the router. Noté howévér, éven though this traffic is unidiréctional NétFlow accounts for all traffic going in and out of the router by récording both transit traffic and traffic déstinéd for the router. By storing only the router's flow-lévél information and néglécting payload it bécomés féasible to summarizé and "skétch" largé amount of data traffic.

A kéy téchnical noté is thé sampling modé at which thé routér is running. In thé GSR's casé a 1-166 sampling raté is spécifiéd. This méans that for évéry 106 packéts processed by thé forwarding énginé or routé processor théré will bé oné packét éxtractéd and réportéd to thé NétFlow process running at thé routér. This is thé lowést allowéd sampling raté on our GSR running NétFlow v5 and cléarly wé can séé that this causés limitations on thé analysis of nétwork data. It is not uncommon to havé short host-to-host séssions whéré thé ovérall transmission doés not éxcééd 106 packéts. It is béyond thé scopé of this papér to discuss NétFlow sampling algorithms but baséd on our éxpériéncé it is likély that if thé transmission is 106 packéts or less NétFlow will not account for it. This introducés a

margin of error to any analysis of data flows but especially to those UDP flows transmitted over our core. With their inherent error correction mechanisms TCP flows are less prone to being ignored by the collection process.

Flow Tools is a collection of programs and libraries used to collect and process NétFlow data. These tools allow users to process stored flow data from a series of command line interfaces. Commands like flow-filter and flow-sort allow the user to filter and sort NétFlow data by IP address, port, AS number and any other parameter present in that data collected. The data is presented on the command line in a table format. However these tools do not provide a dynamic way of dynamically monitoring flow data. Through the use of MonALISA we have used NétFlow data collected and processed by Flow Tools to create a graphical interface by which to view certain characteristics of the Ampath network in a near-real time fashion. For our analysis we implemented flow-tools version 0.66 on a dual Xeon 2.66 GHz system with a copper Gigabit Ethernet network connection to FIU's campus network. The operating system of choice was Federa Core 2.

Having discussed the data gathering techniques in the following we focus on the data dissemination mechanism. For the rationales that we detail in the next section MonALISA is our platform of choice for this purpose.

4. MonALISA and ApMon

MonALISA (Monitoring Agénts in A Largé Intégratéd Sérvicés Architécturé) is a distributéd sérvicés architécturé te colléct, process and act upon réal-timé monitoring information. Whilé its initial targét fiéld of application is nétworks and Grid systéms uséd by thé global high énérgy and nucléar physics collaborations, MonALISA is broadly applicable to many fiélds of data inténsivé sciéncé, and to the monitoring and managément of major réséarch and éducation nétworks. MonALISA is baséd on a scalablé dynamic distributéd sérvicés Architécturé, and is impléméntéd in Java using JINI[21]and WSDL[22] téchnologiés. Thé scalability of thé systém dérivés from thé usé of a multi-thréadéd énginé to host a variéty of loosély coupléd sélf-déscribing dynamic sérvicés, and thé ability of éach sérvicé to régistér itsélf and thén to bé discovéréd and uséd by othér sérvicés or cliénts that réquiré such information. Thé framéwork intégratés many éxisting monitoring tools and procédurés to colléct paramétérs déscribing computational nodés, applications and nétwork pérformancé. Spécializéd mobilé agénts aré uséd in thé MonALISA framéwork to pérform global optimization tasks or hélp imprové thé opération of largé distributéd systém by pérforming supérvising tasks for différént applications. MonALISA is currently running around thé clock monitoring sévéral Grids and distributéd applications on approximatély 150 sités.

The core of the MonALISA monitoring service is based on a modular system design used to perform the data collection tasks in parallel, independently. The modules used for collecting different sets of information, or interfacing with other monitoring tools, are dynamically loaded and executed in independent threads. In order to reduce the load on systems running MonALISA, a dynamic pool of threads is created once, and the threads are then reused when a task assigned to a thread is completed. This allows one to run concurrently and independently a large number of monitoring modules, and to dynamically adapt to the load and the response time of the components in the system. If a monitoring task fails or hangs due to I/O errors, the other tasks are not delayed or disrupted, since they are executing in other, independent threads. A dedicated control thread is used to properly stop the threads in case of I/O errors, and to reschedule those tasks that have not been successfully completed. A priority queue is used for the tasks that need to be performed periodically.

A schématic viéw of this méchanism of collécting data is shown in Figuré 2. This approach makés it rélativély éasy to monitor a largé numbér of hétérogénéous nodés with différent résponsé timés, and at the samé timé to handlé monitoréd units that are not résponding without affécting other méasuréments. The clients, other sérvices or agents can gét any réal-timé or historical data by using a prédicaté méchanism for réquésting or subscribing to séléctéd méasuréd valués. Thésé prédicatés aré baséd on régular expréssions to match the attribute déscription of the méasuréd valués a cliént is intéréstéd in. Théy may also bé uséd to imposé additional conditions or constraints for séléct-

ing the values. In case of requests for historical data, the predicates are used to generate SQL queries to the local database. The subscription requests create a dedicated thread, to serve each client. This thread performs a matching test for all the predicates submitted by a client with the measured values in the data flow. The same thread is responsible to send the selected results back to the client as compressed serialized objects. Having an independent thread per client allows sending the information they need, in a fast and reliable way, and it is not affected by communication errors which may be occurring at other clients. In case of communication problems these threads will try to reestablish the connection or to clean up the subscriptions for a client or a service that is no longer active.



Figure 2: A schematic view of the MonALISA data collection mechanism based on a multi-threaded engine.



Figuré 3: The MonALISA monitoring service for Abilene, shown at a time we injected more than 8 Gigabits/sec.

Figure 3 is a snapshot of the MonALISA monitoring network for Abilene network of the Internet2. It shows all the active nodes running MonALISA services for this particular "farm", discovered automatically by a graphical Mon-ALISA client. The client can display the real time global views and connectivity, as well as the usage and load of the farms. In this particular instance we captured a highly intensive data transfer event on June 19th, 2004 where a group of 12 disk servers in CERN concurrently sent TCP traffic via LHCNet and Abilene to their destinations in Caltech. Note that in this case MonALISA reported a throughput reaching 8.4 Gbps on the Abilene links from Chicago Kansas City Denver Sunnyvale Los Angeles.

A salient feature of the MonALISA design is its extensibility. It allows user-defined monitoring modules, specific to user-specified system and network information, to be easily implemented and integrated in the MonALISA framework. This facilitates the work reported in this paper, as well as our ongoing project to integrate NLANR PMA[5] real-time packet trace analysis and MonALISA. From a systems point of view, MonALISA provides scalable architectural support for collecting, visualizing and responding to the operating conditions of large-scale distributed systems, and it is especially suited for monitoring and controlling large computing systems and networks used in Grid applications.

ApMon[4] is an Application Programming Intérfacé that facilitatés usér-spécific applications to intéract with the MonALISA sérvicés. ApMon allows any application to send paramétérizéd monitoring information to MonALISA. The data can be sent as UDP packets to multiple hosts running the MonALISA sérvicé.

Through the use of ApMon MonALISA services can receive parameterized data in (name, type, value) tuples. When transmitting a data point to MonALISA the application specifies the name of the parameter about to be sent, the type of the parameter (string, object, integer, double) and the actual value of the parameter. For our implementation we decided to use a 64 bit real number as the type for all of our data transmissions with ApMon. This type is represented in ApMon by the constant ApMonConst::XDR_REAL64(). The ApMon module on the MonALISA service will then receive this data and create any needed fields on its database for new parameters or populate existing fields if a particular parameter name already exists.

The resulting data stored in MonALISA is a set of parameters and the values of those parameters over time. Mon-ALISA then provides an interface by which to view one or multiple parameters in a real-time or historical graph.

5. Integrating MonALISA and FlowTools via ApMon

NétFlow allows for the monitoring of a large number of paraméters. For this project we décided to limit the paraméters monitor to the following:

UDP/TCP port déstination/sourcé traffic IP déstination/sourcé traffic IP protocol traffic IP Next Hop traffic AS déstination/sourcé traffic Préfix déstination/sourcé traffic

For most of these parameters we will be monitoring that total traffic in octets over a period of time. The initial period of time was 5 minutes. In the flow-capture startup script above the parameter n288 indicates that we want flow-capture to generate 288 files per day which results in a new file generated very 5 minutes.

Our application will use the most recent file to retrieve the desired parameters. This will result in parameterized NetFlow data being retrieved by our application reflecting 5-minute averages for each of the monitored parameters. Hence the value of each parameter will represent the total number of octets associated with that particular parameter over a 5-minute period. We will use flow-stat to collect all of the parameters specified above. Below is a sample output using flow-stat. As we can see for this particular 5-minute interval the HTTP port (80) was most heavily used at 11274635 octets roughly 10.75 MB.

```
$ flow-cat ft-v05.2005-01-19.135207-0500 | flow-stat -f 5 -S 2
# --- ---- Report Information --- ---
# Fields:
           Total
# Symbols: Disabled
# Sorting: Descending Field 2
# Name:
          UDP/TCP destination port
# Args:
          flow-stat -f 5 -S 2
#
#
# port
        flows
                               octets
                                                    packets
```

#			
80	42816	11274635	82997
25	2877	3770113	7623
2010	28	2543195	1747
4662	2095	2290946	3770
2009	28	2091493	1413

Thé application writtén will parsé thé flow-stat réport and génératé a namé/valué pair which will bé sent to MonaLISA via ApMON. This namé will bé oné of thé following: (1) a port numbér or namé; (2) an IP addréss; (3) a protocol numbér or namé; (4) an AS numbér; (5) a Préfix.

Thé valué of thé paramétér will always bé thé total numbér of réportéd octéts for that paramétér. This will génératé a largé numbér of paramétérs which is difficult to visualizé. ApMon providés thé ability to classify thé paramétérs submittéd to MonALISA by spécifying a clustér and a nodénamé along with thé paramétér namé/valué pairs. A clustér is a group containing a list of nodénamés. Each nodénamé is a group containing a list of paramétérs. Whén submitting data to MonALISA wé will définé éach routér intérfacé as a clustér and thé déscription of thé paramétérs béing méasuréd as a nodénamé within that clustér. For éxamplé thé "NWS Intérnét2" intérfacé in thé Junipér routér will bé considér a clustér (naméd Juniptér – NWS Intérnét2) and it will contain thé following nodénamés: "Egréss Déstination IP", "Egréss Déstination AS", "Egréss Sourcé Ports UDP/TCP", étc. Each nodénamé will contain thé actual paramétér namé/valué pairs. In ordér to givé thé usér thé fléxibility to spécify what valués to monitor and how to classify thém in MonALISA a configuration filés was créaté for our application. Bélow is a déscription of thé allowéd séttings in thé configuration filé.

Setting Name	Description	Example	
cluster_name:	The description of the interface	GSR-StarLight VLAN	
nødé_namé:	The description of the parameter being moni- tored	Ingréss Déstination Ports	
flew_file_directory:	Path to the directory containing the NetFlow files	/heme/netflew/flews/gsr/	
flow_stat_options:	Options passed directly to flow-stat	-f5 S2 n	
flow_filter_options:	Options passed directly to flow-filter	-i 60	
valuė_param:	Thé column numbér from thé flow-stat réport which will bé uséd as thé valué for thé pa- ramétér. Whéré thé first column numbér is 6 .	2	
max_params:	The maximum number of name/value pairs which will be processes and transmitted to MonALISA for this nodename.	15	
Next	End of the nodename's configuration.		

Our application would parse this information and generate the following report:

```
$flow-cat ft-v05.2005-01-19.224516-0500 | flow-filter -i 60 | flow-stat
-f 5 -S 2 -n
# --- ---- Report Information --- ---
#
# Fields:
           Total
# Symbols:
           Enabled
# Sorting: Descending Field 2
           UDP/TCP destination port
# Name:
#
# Args:
           flow-stat -f 5 -S 2 -n
#
#
```

# port #	flows	octets	packets
" 57893	1	913876	1106
32808	1	730500	487
6881	52	210301	283
32920	1	151132	101
32918	1	85132	57
1144	3	84854	61
http	39	68171	106
1279	1	67500	45
eDonkey-20	67	64742	107
3113	2	63044	49
33498	3	58552	40

The application will parse the generated report and retrieve the parameters name and values as specified in the configuration. For this example the parameter names will be the port numbers and the values will be the octets as specified by the configuration file (value_param:2). The following data would then be sent to MonALISA:

Cluster	Nodename	Parameter	Parameter
		Name	Value
GSR-StartLight VLAN	Ingress Destination Ports	57893	913876
GSR-StartLight VLAN	Ingress Destination Ports	32808	151132
•••			
GSR-StartLight VLAN	Ingress Destination Ports	http	68171
• • •	•••	•••	•••

6. Case Study- Monitoring Peer Traffic

In this section we demonstrate the use of our monitoring system in practice. We choose to examine a peer, the NWS Internet2 Gigabit Ethernet link that carries all HPC traffic to and from New World Symphony (NWS), a postdoctorate institution at Miami Beach that has AMPATH be its upstream provider of Internet2 traffic as well as commodity internet.

Oncé monALISA starts⁴ thé scréén préséntéd in Figuré 4 is a global viéw of all running sités / farms. With a propérly éstablishéd working group under MonALISA wé choosé a moré détailéd viéw from thé léft ménu shown in Figuré 5. Thé "TabPan" viéw givés détailéd information about thé monALISA farm but also servés as the starting point to access paramétérs that aré uniqué to thé sité.

⁴ This requires Java to be installed.



Figure 4 - MonALISA startup screen with 'test' group defined.

00	00				N	IonALISA				
<u>F</u> ile	Discovery	Groups Security	<u>H</u> elp						18 farms,	2668 nodes, 69684 paran
÷	and the	Regional Center [select to	Local	MonaLisa	Group	Free Nodes Load [0 -> 0.25]	ML UpTime	CPU_usr mean	Load mean	RateIN [KB/s] mean/total
- 7	P 42 - 25	Unimelb	05:04 (EST)	1.2.20	test	N/A	02 day(s) 13:55:22	Unknown	Unknown	Unknown
~	WMap	IU_ac	13:04 (EST)	1.2.15	test	1 (100%)	N/A	7.1	0.07	0.02 / 0.02
		hpc	13:04 (EST)	1.2.22	test	N/A	02 day(s) 22:30:29	Unknown	Unknown	Unknown
<u> </u>	- 0	MLRepos	19:04 (CET)	1.2.21	test	N/A	29 day(s) 01:55:52	Unknown	Unknown	Unknown
-	GMap	aidong_homepc	11:04 (MST)	1.2.22	test	N/A	13:01:46	Unknown	Unknown	Unknown
1. 40. pet		I2Monitor	13:04 (EST)	1.2.22	test	N/A	13:19:49	Unknown	Unknown	Unknown
	<u></u>	CIT_HEP_XSU	10:04 (PST)	1.2.22	test	N/A	08 day(s) 15:24:42	Unknown	Unknown	Unknown
		RoGrid	20:04 (EET)	1.2.20	test	3 (100%)	04 day(s) 08:28:36	0.23	0.01	0.0 / 0.01
-	TabPan	upb	20:04 (EET)	1.2.22	ml,test	1 (20%)	01 day(s) 07:23:11	2.42	0.01	0.56 / 0.56
	19 M	Alberta_THOR-G	11:04 (MST)	1.2.20	test	44 (77%)	06 day(s) 08:28:00	2.94	0.38	3.48 / 198.16
31		madreselva	19:04 (CET)	1.2.22	test	4 (100%)	03 day(s) 08:04:17	3.53	0.08	0.05 / 0.22
	Load	CIT-HEP	12:00 (PST)	1.2.20	test	N/A	48 day(s) 16:40:24	Unknown	Unknown	Unknown
1		HIP_CERN	19:04 (CET)	1.2.22	test	N/A	01 day(s) 04:44:23	Unknown	Unknown	Unknown
		fesb	19:04 (CET)	1.2.15	test	N/A	N/A	Unknown	Unknown	Unknown
_	WANI .	MonaLemon	19:04 (CET)	1.2.21	test	N/A	17 day(s) 02:27:35	Unknown	Unknown	Unknown
1		qcdoc-bnl	13:04 (EST)	1.2.22	test	3 (50%)	20:04:41	5.82	0.36	0.09 / 0.35
A		ARDA	19:04 (CET)	1.2.20	test	N/A	49 day(s) 09:37:14	Unknown	Unknown	Unknown
		spud.fnal.gov	12:04 (CST)	1.2.22	test	N/A	00:08:04	Unknown	Unknown	Unknown
	o jobs									
間目										

Figure 5 - MonALISA farm overview: note the MonALISA service I2Monitor that we started

In Figuré 5 wé aré présentéd with a snapshot of thé current staté of thé MonALISA group "tést". Thé MonALISA sérvicé that pértains to our particular study is titléd I2Monitor. Choosing thé I2Monitor farm, wé aré thén présentéd with thé AMPATH spécific data which wé havé chosén to intégraté into our analysis, as shown in Figuré 6.

I2Monitor	I2Monitor@lisa.fiu.edu:9002				
Local Time : 13:08 (EST) MonALISA	(EST) MonALISA Version: 1.2.22				
 I2Monitor Master MonaLisa Juniper Aggregate Juniper - Commodity to FIU Juniper - NWS Internet2 Juniper - NUS Internet2 Juniper - OC12 to CBX500 Juniper - Rack LAN for NOTA GSR Aggregate GSR - SENACYT GSR - SENACYT GSR - StarLight VLAN GSR - UPR Juniper - AURA / Gemini Juniper - Global Crossings Juniper - UFG / El Salvador 	Parameters History Plot Realtime Plot Nodes Summary Cluster Summary Modules				
Site info					

Figure 6 AMPATH Farm specific data

It is évident which paraméters we have specified to be monitored. Two core routers (a Cisco GSR and a Juniper M10) are monitored. As an example we will analyze the Juniper NWS Internet2, as shown in Figure 7.



Figure 7 - Traffic parameters studied in this paper

In Figuré 8 by choosing Egréss Sourcé AS wé vièw thé currênt sét of storéd AS numbérs bélonging to flows léaving our Junipér routér and déstinéd to the Néw World Symphony nétwork. With thé "Paramétérs" corrésponding to thé AS numbérs travérsing our routér déstinéd to NWS wé havé a cléar top-lévél vièw of flows of thé nétwork, and wé can délvé déépér into this data by showing a réaltimé plot or history plot of the AS data gathéréd. This is givén in Figuré 9.

I2Monitor@lisa.fiu.ed	u:9002
Local Time : 13:24 (EST) MonALISA Version: 1.2.22	
🔻 🤤 I2Monitor	Parameters
Master	RIPE-ASNBL
MonaLisa	BATI-MIA
Juniper Aggregate	INKTOMI-LA
Juniper - Commodity to FIU	HOTMAIL-AS
🔻 🤤 Juniper - NWS Internet2	MICROSOFT-
Egress Destination Ports UDP/TCP	GBLX
Egress Destination IP	CWUSA
Egress Destination AS	WAN
Egress Destination Prefix	TASF
Egress Source Ports UDP/TCP	VERIO
Egress Source IP	ASN-QWEST
Egress Source AS	YAHOO-US2
Egress Source Prefix	CYMITAR
Egress IP Next Hop	CITI13
Egress IP Protocols	BATI-ATL
Ingress Destination Ports UDP/TCP	INTERNAP-B
Ingress Destination IP	AOL-PRIMEH
Ingress Destination AS	History Plot Realtime Plot
Ingress Destination Prefix	(History Fist) (Heddianie Fist)
Ingress Source Ports UDP/TCP	Nodes Summary 💎 Cluster Summary 💎
Ingress Source IP	Modules
Ingress Source AS	monXDBUDP
Ingress Source Prefix	indiscertoer
Ingress IP Next Hop	
Ingress IP Protocols	
Juniper - OC12 to CBX500	
GOD A server state	
GSK Aggregate	
Juniner - AllPA / Gemini	
Juniper - Backup Commodity to FIU	
Juniper - Global Crossings	
Juniper - UFG / El Salvador	
Site info	

Figure 8 - ASNs traversing AMPATH towards NWS



Figure 9 - ASN traffic destined to NWS

In Figure 10 we show the protocols of the traffic traversing the GSR router, the main provider of Internet2 to AM-PATH IXP peers. We now drill down through the GSR Aggregate data to show the IP Protocols currently being recorded by the NetFlow process on the router. A historical look can show the recent or long-term utilization of IPv6; and we can visualize it by means of a history plot in Figure 11. We can quickly note that IPv6 data has not been a significant load on AMPATH during the time period specified.



Figure 18 - Router Aggregate NetFlow data



Final Report 2002 - 2005

Figure 11 Two-hour snapshot of IPv6 data traversing AMPATH

8. Discussions and Future works

NétFlow data contains a rich amount of nétwork information with a multitudé of applications. Through thé usé of thé distributéd monitoring énvironment providéd by MonALISA and thé réporting fléxibility émbéddéd in thé FlowTools API it is possiblé to éncapsulaté and summary this data in visually friéndly mannér. In doing saw wéré ablé to créaté dynamic and réal timé vièws of thé AMPATH Intérnét 2 nétwork traffic and its béhaviors. Diréct application of this téchnology could bé uséd to furthér undérstand nétwork traffic béhavior and trénds in compléx réséarch nétworks.

The traffic views generated by MonALISA and the NetFlow application focused on the "top talking" flows. That is MonALISA only received the top x talking flows for a particular period of time. We saw that this technique produced a seemingly random setup of results. Further study should be made to understand this trend and determine its inherent properties.

Oné furthér téchnology that wé inténdéd to éxploré was that of thé National Laboratory for Appliéd Nétwork Réséarch (NLANR) PMA (Passivé Monitoring Agént) [5]. Théré aré kéy différéncés bétwéén PMA data and NétFlow data worthy of sérious réséarch éffort. With our résults wé hopé to providé a stablé platform from which nétworks of varying dégréés can bé closély monitoréd, théir traffic pattérns cléarly idéntifiéd and thé appropriaté décisions takén to réctify issués which négativély impact pérformancé or augment thosé which havé a positivé impact on thé délivéry of sérvicé to an énd usér.

References:

- 1. NétFlow Sérvicés Solution Guidé, Cisco Systèms Inc.
- 2. FlowTools: http://www.splintered.net/sw/flow-tools/.
- 3. MonALISA: http://monalisa.caltech.edu.
- 4. ApMon: http://monalisa.caltech.edu.
- 5. PMA: http://moat.nlanr.net