

AFRL-IF-RS-TR-2004-303
Final Technical Report
October 2004



NATIONAL TRANSPARENT OPTICAL NETWORK CONSORTIUM (NTONC)

Nortel Networks

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-303 has been reviewed and is approved for publication

APPROVED: /s/

ROBERT L. KAMINSKI
Project Engineer

FOR THE DIRECTOR: /s/

WARREN H. DEBANY, JR., Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE OCTOBER 2004	3. REPORT TYPE AND DATES COVERED Final Sep 98 – Mar 04	
4. TITLE AND SUBTITLE NATIONAL TRANSPARENT OPTICAL NETWORK CONSORTIUM (NTONC)			5. FUNDING NUMBERS C - F30602-98-2-0194 PE - 603761E PR - G152 TA - 00 WU - 01	
6. AUTHOR(S) Paul Daspit				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Nortel Networks 2010 Corporate Ridge McLean Virginia 22102			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/IFG 525 Brooks Road Rome New York 13441-4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2004-303	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Robert L. Kaminski/IFG/(315) 330-1865/ Robert.Kaminski@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) This final report describes the National Transparent Optical Network Consortium (NTONC) efforts in dense wavelength division multiplexing (DWDM) transport, switching technologies and control strategies required to develop, deploy and operate the terabit per second optical networks needed to meet requirements of Next Generation Internet applications. The report also describes NTONC's "Flat Network Architecture" research for very high capacity networks. The objective of this work was to devise a network to serve as the Next-Generation Internet (NGI) with minimum inner-to-outer capacity ratio, linear capacity expansion, low incremental delay, no data loss, ease of access, automatic new user registration, ease of network engineering, fast recovery from failure conditions, and self-governance. The report concludes with a summary of work performed in the areas of ActiveNets/Openet research, sensor deployment and networking research and proof of concept, MEMs CAD networking research, SuperNet Analysis and fiber expansion in the Washington DC area.				
14. SUBJECT TERMS Optical networking, Next Generation Internet, Wavelength Division Multiplexing			15. NUMBER OF PAGES 39	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

Technical Achievements for Milestones 1 – 28	1
1. Introduction	1
2. Summary – NTONC Program Recap	1
Program Goals and Objectives	1
Consortium Members and Roles	1
Period of Performance	2
3. Methods, Assumptions, Procedures, Results and Discussion of Accomplishments	2
Task 1 – Deploy and Manage a West Coast Core Network	2
Task 2 – Provide Applications Integration and Support	6
Task 3 – Provide Network Management and Control Capabilities	10
Task 4 – Conduct “Flat Network Architecture” Research	10
Task 5 – Program Management	11
4. Conference Support, Demonstrations and Experiments	12
SC99 (Nov 1999, Portland, OR)	12
OFC 2000 (Mar 2000, Baltimore, MD)	14
SC2000 (Nov 2000, Dallas, TX)	14
5. Intellectual Property (September 1998 – June 2002)	18
Patent Filings	18
Technology Transfer	19
Technical Achievements for Milestones 29 - 38	20
6. Introduction	20
7. Summary – Participating Organizations	20
Task 6 – Active Networks Research	20
Task 7 – Berkeley Realtime-Application Network Demonstration (BRAND): Networked	22
Task 8 - SuperNet Maintenance and Research Support (SMRS)	23
Task 9 – 10Gbps Ethernet Connectivity Over BoSSnet	24
Task 10 – Advanced Networking Infrastructure	24
Task 11 – Program Management	27
8. Intellectual Property (July 2002 – March 2004)	27
Patent Filings	27
Technology Transfer	28
9. Conclusions	29
Other Related Documents	31
Related Documents for Milestones 1-28:	31
Related Documents for Milestones 29-38:	31
List of Symbols, Abbreviations and Acronyms	32

List of Figures

Figure 3.1	Target Core Network Configurations.....	3
Figure 3-2	San Francisco Bay Area Metro Network Migration.	4
Figure 3-3	San Francisco Bay Area Network Configuration.....	5
Figure 3-4.	Los Angeles Area Network Configuration	6
Figures 3-5A/B.	MEMs-based Optical Cross-connect at the NTON Oakland node and OMM-NTON user interface.....	9
Figure 4-1.	NTON Configuration in support of Supercomputer Conference, Nov 1999.	12
Figure 4-2.	Packet over SONET (POS) Network Configuration for Supercomputer Conference, Nov 1999..	13
Figure 4-3.	Seattle-Portland OC-192 Bandwidth support over NTON for Supercomputer Conference 1999..	14
Figure 4-4.	NTON Configuration in Support of SC2000 (Nov. 2000, Dallas, TX).....	16
Figure 4-5.	SC2000 Bandwidth Challenge Configuration - Visipult.....	16
Figure 4-6.	NTON Configuration in support of ASCI Demo at SC2000.	17
Figure 7-1	DWDM-RAM versus Layered Grid Architecture (adapted from [1])	25

List of Tables

Table 4-1.	SC2000 WAN Application Experiments and Demonstrations.	15
Table 5-1.	NTONC Report of Patents - September, 1998 through June 2002.....	18
Table 8-1.	NTONC Report of Patents – July 2002 through March 2004	28

Technical Achievements for Milestones 1 – 28

1. Introduction

This report is submitted in accordance with TIA F30602-98-2-0194 Article 29 – Final Report, and documents NTONC technical achievements for Milestones 1 through 38. Part 1A covers Milestones 1-28 (Oct 1998-June 2002) and Part 1B (starting on page 21) covers Milestones 29-38 (July 2002 – March 2004).

2. Summary – NTONC Program Recap

Program Goals and Objectives

The objectives of the NTONC program included the deployment and demonstration of a high capacity open test bed communications network to facilitate collaborative research. The central goals of the NTON program were (a) build an open research network capable of demonstrating the full benefits of very high capacity networks and the optical transport, switching and network management and control strategies on which it is built; and (b) conduct research on Flat Network architectures.

Consortium Members and Roles

2.1. Nortel Networks

- 2.1.1. Provide DWDM transport and switching equipment.
- 2.1.2. Conduct Flat Network Architecture study.
- 2.1.3. Manage the NTON Consortium

2.2. GST Telecom

- 2.2.1. Provide dark fiber between San Diego and Seattle
- 2.2.2. Provide COLO space and power for equipment
- 2.2.3. Perform network operations, surveillance and maintenance

2.3. Lawrence Livermore National Laboratory

- 2.3.1. Provide Applications Integration support
- 2.3.2. Provide San Francisco Bay Area access network coordination
- 2.3.3. Provide COLO space and power for equipment

2.4. Sprint Advanced Technology Laboratory

- 2.4.1. Provide San Francisco Bay area dark fiber
- 2.4.2. Provide COLO space and for equipment

2.5. Bay Area Rapid Transit (BART)

- 2.5.1. Provide San Francisco Bay area dark fiber

Period of Performance

The original TIA was signed in September 1998 for a 27-month period of performance. NTONC extended this an additional 18 months at no cost to the government. Performance for Milestones 1 through 28 ended in June 2002.

3. Methods, Assumptions, Procedures, Results and Discussion of Accomplishments

Task 1 – Deploy and Manage a West Coast Core Network

Principal Investigators:

Nortel Networks – Hal Edwards, Paul Daspit (pdaspit@nortelnetworks.com)

GST Telecom – Ben Peek (benpeek@comcast.net)

Lawrence Livermore National Laboratory – Bill Lennon (wjlennon@llnl.gov)

Dark Fiber contributors:

Sprint ATL – Frank DeNap (fdenap@sprintlabs.com)


Bay Area Rapid Transit – Dave Warwick, Bill Black (wblack@bart.gov)

3.1. A 2500km 16-lambda OC-192 DWDM SONET network was engineered between San Diego Supercomputer Center (SDSC) and the University of Washington/Northwest Pacific GigaPOP located at Seattle, WA with add/drop locations at Los Angeles, Oakland and Portland. The initial 10Gb/s segment between Los Angeles and Oakland was activated June 1999. The segment between Portland and Seattle was made operational in November 1999. Temporary connectivity between Portland and California nodes was established for a 10-day period in November 1999 in support of the High Performance Computing and Communications Conference (SC99). Permanent connectivity between Portland and Oakland and between Los Angeles and SDSC was not established as result of GST Telecom not delivering the dark fiber.

3.2. The 16-lambda OC-192 backbone network was engineered to provide up to 4 OC-48 interfaces at each of the 5 terminal add/drop locations. These interfaces were designed to support dedicated applications research requirements or provide switch router overlay network interconnections.

3.3. The existing 4-lambda SONET OC-48 San Francisco Bay network was converted to a 24-lambda DWDM metro network supporting any type of transport protocol at data rates from 16Mb/s to 2.5Gb/s. A 3-node linear ADM network was established between LLNL, Oakland and Sprint ATL.

Task 1 - Deploy and manage the core network



- Target Core Network Configuration ... bandwidth for both flexible and fixed-network research applications.
- Research bandwidth up to 8λ @ 10Gbps/ λ .
- 7x24 management by GST at their Vancouver network operations center.




Figure 3.1 Target Core Network Configurations

3.4. Research nodes connected via the San Francisco Bay Area metro network included Sandia National Laboratory Livermore (SNLL), UC Berkeley (UCB), Lawrence Berkeley National Laboratory (LBNL), Stanford Linear Accelerator Center (SLAC), and NASA Ames Research Center (ARC). In addition to the 3 San Francisco Bay Area metro nodes, Nortel deployed metro DWDM equipment at LBNL and OC-48 SONET terminals at NASA Ames and SLAC. In the Los Angeles area, Nortel deployed an OC-48 SONET terminal at Cal Tech for connections to the backbone node from there and from JPL.

3.5. Connections requested by researchers were typically OC-3/12/48 POS and OC-3/12/48 ATM as well as Gigabit Ethernet. Connectivity was established between researchers in the San Francisco Bay area and researchers in the Los Angeles area by interconnecting the backbone OC-192 network and the DWDM metro network.

Task 1 - Reconfigure the Bay Area Ring to 24 λ




Figure 3-2 San Francisco Bay Area Metro Network Migration.

3.6. A 16-lambda DWDM system was deployed in the Los Angeles area to solve a fiber-exhaust problem and provide connectivity between two GST Telecom facilities. This system was configured to provide connectivity for Gigabit Ethernet, OC-48/12/3 POS and OC-48/12/3 ATM access connections.

3.7. Research centers connected to the Los Angeles NTON node included Jet Propulsion Laboratory (JPL), California Institute of Technology (CalTech) and University of Southern California/Information Sciences Institute-West (USC/ISI-W).

3.8. Asynchronous Transfer Mode (ATM) networking was provided initially using Vector switches capable of OC-12 and OC-3 connections. A 3-node Passport 15000 Multifunction Switch was later provided with OC-48 switch-to-switch connectivity and OC-12/3 user interfaces.




Figure 3-3 San Francisco Bay Area Network Configuration

3.9. Avici Terabit Switch Routers, one at Oakland and one at Los Angeles provided an OC-48 switch-switch trunk and OC-12/3 Packet over Sonet (POS) user interfaces.

3.10. Network peering was established between NTON and HSCC (High-Speed Connectivity Consortium) network in July 1999. This provided POS connectivity between researchers using NTON and researchers with access to HSCC nodes at Washington, DC, Pittsburgh Supercomputer Center, University of Washington/Northwest Pacific GigaPOP and Dallas, TX for the Supercomputer Conference in November 2000. Network peering between NTON and Internet2 was established in October 2000 through the HSCC peering point in Los Angeles.




Figure 3-4. Los Angeles Area Network Configuration

Task 2 – Provide Applications Integration and Support

Principal Investigators:

Nortel Networks – Hal Edwards, Paul Daspit (pdaspit@nortelnetworks.com)

Lawrence Livermore National Laboratory – Bill Lennon (wjlennon@llnl.gov)

3.11. The network infrastructure described in Task 1 above was made available to government, university and private sector researches to enable high bandwidth Next Generation Internet applications research. NTONC assisted each research principal investigator to identify their network requirements, define the transport interfaces, make the necessary network interconnects, test the baseline functions of the access connection and monitor long-term performance.

3.12. UCB/LBNL Applications

3.12.1. The DATA GRID Project – distributed parallel storage system (DPSS) demonstrated 980 Mb/s across Gigabit Ethernet LAN and 570 Mb/s over NTON OC-12 SONET/ATM.

3.12.2. DARPA MEMs Project – very large high-resolution video of MEMs devices on the LBNL DPSS. Access provided by NTON to MEMs researchers at UCB, MIT and Carnegie Mellon University (CMU). Demonstrated 300 Mb/s end-to-end at OFC2000 conference.

3.13. SLAC Applications

3.13.1. Particle Physics Data Grid (PPDG) – BaBar experiment generates events at SLAC. 800 Mb/s high-speed file transfers of event data to Cal Tech.

3.13.2. Stanford Synchrotron Radiation Laboratory (SSRL) – High-speed readout detectors produce images of 50 GB/day. Full-motion video transferred to other institutions for control room telepresence.

3.14. QoS experiments – high-speed teleconferencing systems mixed with data distribution and distributed processing traffic.

3.15. Sandia National Laboratory Applications

3.15.1. Combustion Corridor Project – interactive, high-resolution visualization of large combustion simulations. Continuously retrieve large amounts of data from LBNL's DPSS. Compared simulation data with experimental results at SNLL's combustion research facility.

3.15.2. ASCI Strategic Alliance Program – includes Stanford, Cal Tech, University of Chicago, University of Utah, NPACI/SDSC and University of Illinois Urbana/Champaign. Facilitated Cal Tech and SNLL/LLNL resource sharing.

3.15.3. IETF MPLS Protocol Trials – Studies impact of MPLS on NTON traffic between SNLL, LBNL and SLAC. Included performance of IP over WDM.

3.16. NASA Ames Applications

3.16.1. Digital Earth Workbench – very high resolution satellite imagery of the entire planet with 3D models of buildings. Demonstrated rapid access to massive amounts of geospatial data and measured throughput, latency, and frame rate. Distributed graphic images equivalent to multiple streams of raw HDTV.

3.16.2. Distributed Collaborative Virtual Wind Tunnel – interactive, near real-time collaborative investigation of complex time-varying vector and scalar fields in three dimensions. Participants at NASA Ames, GSFC and JPL interface with abstract data as if they were in the same room interacting with an actual physical model in a wind tunnel environment.

3.17. Lawrence Livermore National Laboratory Applications

- 3.17.1. Optical Modeling – defined, implemented and validated a design simulation tool. Conducted case studies for placement of erbium-doped fiber amplifiers (EDFA) for power-limited 5 Gb/s transmission, characterization of 215 km WDM link, and computer-aided design of 430 km optical CDMA system at 2.5 Gb/s
- 3.17.2. Peregrine – in collaboration with NIH/National Cancer Institute, validated Monte Carlo method-based evaluation and planning tool for radiation treatment for multiple tumor types.
- 3.17.3. Visible Embryo – 10 PB digital library, access and visualization tools for embryology research, education and clinical planning.


3.18. SRI International Applications

- 3.18.1. TerraVision II – distributed, interactive terrain visualization system. Navigate in real-time through a 3-D graphical representation of landscape derived from topological data and aerial images. Retrieves and merges Terabytes of distributed data to include aerial and satellite imagery, topography, weather data, buildings and other cultural features. Demonstrated at OFC2000 (Baltimore, MD) and DARPA (Washington, DC).
- 3.18.2. Active Networks Backbone Coordination Center (ABOCC) – ISI and SRI joint operation for DARPA ABONE services such as secure software distribution, core node monitoring and maintenance, user and node registration for 50 CARIN and Internet-resident core nodes.


3.19. Optical Micro-Machines Demonstration

- 3.19.1. 2-D MEMs Switch with Remote Control – dual 8x8 full-duplex MEMs-based optical fiber switch were deployed at the Oakland add/drop node to groom OC-12 POS and ATM connections between the NTON metro WDM network and the NTON backbone network. (See Figures 3-5A/B). Using the Internet and telnet protocols, control of the MEMs switch at the NTON Oakland node was demonstrated at OFC2000, Baltimore, MD.

OMM OXC manages Oakland NTON connections



- First OXC in a “production” network (Feb 2000)
- Component vs subsystem
- Insertion loss range: 2-6 dB
- Modem access, remote control
- Provides some TS help but needs built-in OAM tools



Figures 3-5A/B. MEMs-based Optical Cross-connect at the NTON Oakland node and OMM-NTON user interface.

Task 3 – Provide Network Management and Control Capabilities

Principal Investigators:

Nortel Networks – Hal Edwards, Paul Daspit (pdaspit@nortelnetworks.com)

GST Telecom – Ben Peek (benpeek@comcast.net)

Lawrence Livermore National Laboratory – Bill Lennon (wjlennon@llnl.gov)

3.20. Surveillance and management of NTON was accomplished using commercial off-the-shelf network management tools. Networks managed included

- OC-192 SONET backbone network
- Optera Metro 5200 DWDM access networks in Los Angeles and the San Francisco Bay areas
- IP/POS Terabit Switch Routers (Oakland and Los Angeles)
- Passport 15000 Multi-service Switch ATM networks.

Management of the MEMs switch deployed at the Oakland node was accomplished by prototype element management software.

3.21. The GST Telecom Network Operations Center (NOC) located in Vancouver, WA provided surveillance and management of the NTON network elements listed above using the GST Telecom network management VPN.

3.22. Development of prototype software to facilitate management and control was not required. Funding for this task was reprogrammed to new tasks.

Task 4 – Conduct “Flat Network Architecture” Research

Principal Investigator:

Nortel Networks – Maged Beshai (beshai@nortelnetworks.com)

3.23. The objective of this task was to examine “novel” new network architectures that scale to very high capacities, offer controllable service quality and are adaptable to a future global Internet. The proposed architecture exploits advances in electronic and photonic switching technology to realize a network that is scalable, tractable, efficient, and almost inexhaustible. After studying several alternatives, one architecture, called the PetaWeb, appears to offer the simplest solution. It comprises several electronic edge nodes and a smaller number of bufferless core nodes. The PetaWeb is an edge-controlled complete network with an agile core that can be reconfigured rapidly in response to traffic variations. A complete network does not require traffic engineering and it allows any source node to send its entire traffic, when required, to a sink node of equal or higher capacity.

3.24. The number of edge nodes that can be supported by a PetaWeb is limited by the capacity of each of the bufferless core nodes. To extend the network coverage to several billions of edge nodes, the PetaWeb is used as a building block of a multi-dimensional structure called a YottaWeb, which scales to a capacity of the order of yottabits per second (10^{24} bits a second).

3.25. In overview, a global network that scales virtually indefinitely is realizable. In the proposed architectures, structural simplicity and high performance are realized by:

- Significant hardware reduction, due to architectural simplicity
- Very significant software-complexity reduction, due to hardware simplicity
- Very significant operational effort reduction, due to hardware simplicity and reduced software complexity

3.26. The end user benefits from a high performance communication with the propagation delay (milliseconds) being the dominant source of latency. A user can initiate a session that requires Terabits a second (Tb/s), for example, and gets it in milliseconds. The service provider benefits from the flexibility and speed of configuring edge-to-edge paths of arbitrary and adaptive capacity that matches traffic requirements. The multiplicity of routes from source to sink enables fast recovery resulting in a fault-tolerant network.

3.27. Two strategies for interworking with the current Internet are considered. In the first strategy, the new network deals with the Internet at arm's length. This allows the new network to grow freely and efficiently while still ensuring that each traffic source can reach each traffic sink in the combined network. The second strategy is based on an interwoven co-existence of the Internet and the new network.

3.28. Results of this study were documented in a report "PetaWeb – Building Block for a Yottabit-Per-Second Network" (March 2001). The Introduction, Executive Summary and Research Objectives sections of the PetaWeb report are available upon request from AFRL.

3.29. Relevant published papers:

François J. Blouin, Andrew W. Lee, Andrew J.M. Lee, and Maged Beshai *Comparison of two optical-core networks*, Journal of Optical Networking, Vol 1, No. 1 <http://www.osa-jon.org/abstract.cfm?URI=JON-1-1-56>

Jules R. Degila, Brunilde Sanso, [A Meta-Search Procedure for the YottaWeb Topology Nodal Arrangement Problem](#) (PDF) 

Task 5 – Program Management

NTONC Director – Hal Edwards

NTONC Program Manager – Paul Daspit (pdaspit@nortelnetworks.com)

Nortel Networks provided overall technical and program management for the NTON program. Nortel Networks also provided management of the NTON Consortium to include services of the NTON Consortium Financial Officer and Attorney.

4. Conference Support, Demonstrations and Experiments

SC99 (Nov 1999, Portland, OR)

NTON OC-192 backbone network segments Seattle – Portland and Oakland – Los Angeles were operational in time to support the High Performance Computing and Communications Conference (SC99). Two OC-48 2.5 Gb/s channels between Portland and Los Angeles were provided on loan from Enron. Refer to Figures 4-1 through 4-3.




Figure 4-1. NTON Configuration in support of Supercomputer Conference, Nov 1999.

Packet over SONET (POS) traffic in the San Francisco Bay Area research sites was routed through the Avici TSR at Oakland to the TSR at Los Angeles and then to the TSR located at the Portland Convention Center. Numerous high-bandwidth applications were demonstrated by NASA Ames, LBNL, SLAC, LLNL, SNLL, Cal Tech, and JPL.




Figure 4-2. Packet over SONET (POS) Network Configuration for Supercomputer Conference, Nov 1999

Agency networks such as the ESnet (DoE), NREN (NASA) and DREN (DoD) were also extended to the Portland Convention Center using NTON POS and SONET/ATM bandwidth.

UW, ResearchTV, and Sony extended previously successful LAN demonstrations to deliver five HDTV streams from five WindowsNT computers in Seattle to five machines at SC99 in Portland. This demonstrated an aggregate bandwidth for the application of more than 1 Gbps over a 300 km distance. Nortel and GST Telecom provisioned an OC48 link from the Pacific Northwest Gigapop to the Portland Oregon GigaPOP, and then to a Juniper Networks switch at SC99 show floor. This created a 300-kilometer 4-hop path from Microsoft in Redmond Washington to the Windows 2000 desktop on the show floor. The data rate averaged 1.3 Gbps (sustained) and peak rate of 1.6 Gbps.




Figure 4-3. Seattle-Portland OC-192 Bandwidth support over NTON for Supercomputer Conference 1999.

The University of Washington sent five concurrent HDTV streams from 5 workstations on the Seattle campus across the same backbone via TCP/IP to five WindowsNT systems on the SC99 show floor. So, the NTONC backbone was carrying in excess of 2.5 Gbps of TCP/IP traffic to five WindowsNT systems during SC99.

OFC 2000 (Mar 2000, Baltimore, MD)

As part of the DARPA-sponsored SuperNet participation, NTON provided program briefings and hosted the remote Internet access to the MEMs switch located at the NTON Oakland node (reference Section 3.19 above).

SC2000 (Nov 2000, Dallas, TX)

Over 20 high performance application experiments and demonstrations were conducted at SC2000 that required access to wide-area research networks. NTON provided a significant part of the total SuperNet bandwidth. All of the West Coast applications that required bandwidth in excess of 500 Mb/s were carried over NTON. Table 4-1 lists the SC2000 WAN applications. NTON-supported applications are highlighted.

Research agencies using NTON bandwidth during SC2000 included Sandia National Laboratory –Livermore (SNLL), Lawrence Berkeley National Laboratory (LBNL), Argonne National Laboratory (ANL), California Institute of Technology (CalTech), Stanford Linear Accelerator Center (SLAC) and the National Aeronautics and Space Administration (NASA) Ames Research Center.

Table 4-1. SC2000 WAN Application Experiments and Demonstrations.

#	Application	Access Network	Transport to Dallas	BW (Gbps)	BW Time Share?	UDP TCP
1	UW/ISI/Tektronix uncompressed HDTV	P/NW & Abilene	Abilene	1.5	N	UDP
2	Sandia ASCI Gbps FTP in Cluster Comp	Sandia-NTON	HSCC	1.5	Y	TCP
3	LBL Visapult	LBL-NTON	HSCC	1.5	Y	TCP
4	Caltech Particle Physics using Globus	Caltech-NTON	HSCC	0.8	Y	TCP
5	Project Data Space	Univ. Illinois	Abilene	0.25	N	TCP
6a	Data Man Climate Res - Striped FTP	ANL/LBL-NTON	HSCC	1.5	Y	TCP
6b		ANL	ESnet	0.155	N	
7	Reservoir Simulation	SDSC-CalREN	Abilene	0.1	N	both
8	Telescience Portal (Telemicroscopy IPv6)	UCSD-CalREN NPACI	Abilene	0.036	N	UDP
9	World Wide Metacomp	HLRS - Multiple	Abilene & vBNS	< 0.2	N	TCP
10	Scalable HiRez Collab	ANL - Multiple	Abilene, vBNS & ESnet	< 0.1	N	both
11	QoS Enabled Audio	Stanford-CalREN	Abilene	0.1	N	both
12	HiRez Viz Play on Tiled	ANL	Abilene	0.6	N	TCP
13	SLAC	NTON	HSCC	0.9	Y	TCP
14	Interactive Distance Viz	USC / ISI-E	HSCC	0.45	?	UDP
15	NRL HDTV/ATM	NRL-ATDnet	Qwest	1.5	Y	UDP
16	NASA Digital Sky Demo	ATDnet/NTON	HSCC	0.25	Y	TCP
17	Digital Amplitheater	DC - SuperNet	HSCC	0.1	Y	UDP
18	SRI Digital Earth	DC - SuperNet	HSCC	0.155	Y	TCP
19	SuperNet NOC	SuperNet	HSCC	0.05	N	TCP
20	NASA DISS	GSFC net	Abilene	0.3	N	TCP
21	LSR/MTU Demo	DC - SuperNet	HSCC	0.75	Y	TCP
22	Web 100	CMU	Abilene	0.1	N	TCP
23	Intrepid Network Collab	Ames - NREN	Abilene	0.1	N	TCP




Figure 4-4. NTON Configuration in Support of SC2000 (Nov. 2000, Dallas, TX).

SC2000 Network Challenge for Bandwidth-Intensive Applications used NTON 2.5 Gb/s POS connectivity between Los Angeles and Oakland as shown in Figure 4-4. Winning the “Fastest and Fattest” category for overall best performance was a team from Lawrence Berkeley National Laboratory demonstrating Visapult, a prototype application and framework for remote visualization of terascale datasets. The Visapult team recorded a peak performance level of 1.48 gigabits per second over five-second sample periods. Visipult network connectivity is shown in Figure 4-5. Bandwidth restrictions over HSCC prevented attempts at higher performance levels.




Figure 4-5. SC2000 Bandwidth Challenge Configuration - Visipult

Taking top honors in the “Hottest Infrastructure” category was a team representing the University of Southern California/Information Sciences Institute and Argonne, Lawrence Livermore and Lawrence Berkeley National Laboratories running “A Data Management Infrastructure for Climate Modeling Research.” The team demonstrated its infrastructure for secure, high-performance data transfer and replication for large-scale climate modeling data sets and achieved a peak performance level of 1.03 gigabits over the NTON 2.5 Gb/s POS path. Performance was limited due to the HSCC bandwidth restriction.

The Accelerated Strategic Computing Initiatives (ASCI) demonstration at SC2000 used a similar NTON configuration for POS traffic between Los Angeles and Oakland as shown in Figure 4-6. A number of high performance applications were demonstrated to include Interactive Remote Visualization and Redundant Array of Independent Tapes High Performance Storage System (RAIT HPSS),




Figure 4-6. NTON Configuration in support of ASCI Demo at SC2000.

5. Intellectual Property (September 1998 – June 2002)

Patent Filings

Over the 45-month program through June, 2002, Nortel Networks made 20 patent application filings, all of them regarding aspects of PetaWeb research in accordance with the intellectual property terms of the TIA. Details of the 20 patent applications are shown in Table 5-1. Enclosure G provides the final DD Form 882 Report of Inventions and Subcontracts (through June 2002).

Table 5-1. NTONC Report of Patents - September, 1998 through June 2002

CY-#	Title	Names of Inventors	Date Submitted	Patent Application Number		
				USPO (USPO Awd #)	Canada (CPO)	Europe (EPO)
99-1	Rate-Controlled Multi-Class High-Capacity Packet Switch	Maged E. Beshai Ernst A. Munter	2/4/99 (USPO)	09/244,824	2,296,923	300700.2
99-2	Self-Configuring Distributed Switch	Maged E. Beshai Richard Vickers	4/6/99 (USPO)	09/286,431	2,303,605	302888.3
99-3	Universal Transfer Method and Network with Distributed Switch	Maged E. Beshai	8/11/98 (USPO)	09/132,465 (6,356,546)	2,279,776	99306338.7
99-4	Routing and Rate Control in a Universal Transfer Mode Network	Maged E. Beshai	8/11/98 (USPO)	09/132,464	N/A	N/A
00-1	State information and routing table updates in large-scale data networks	Maged E. Beshai Richard Vickers	9/27/99 (USPO)	09/405,003	N/A	N/A
00-2	High-capacity TDM packet switch	Maged E. Beshai	9/27/99 (CPO)	09/550,489	2,283,627	308473.8
00-3	Scheduling of variable-size packet data under transfer rate control	Maged E. Beshai	10/14/99 (USPO)	09/417,769	N/A	N/A
00-4	IP address resolution methods and apparatus	Maged E. Beshai	12/23/99 (USPO)	09/471,244	N/A	311427.9
00-5	Agile optical-core distributed packet switch	Maged E. Beshai Richard Vickers	12/30/99 (USPO)	09/475,139	N/A	N/A
00-6	Global distributed switch	Maged E. Beshai	12/31/99 (CPO)	09/748,848	2,329,681	N/A
00-7	High Capacity WDM Data Network	Maged E. Beshai	3/1/00 (USPO)	09/516,938	N/A	N/A
00-8	Multi-dimensional Lattice Network	Maged E. Beshai	7/24/00 (USPO)	09/624,079	N/A	1306335.9
00-9	Courteous Routing	Maged E. Beshai Francois Blouin	8/1/00 (USPO)	09/630,190	N/A	N/A
01-1	Providing Access to a High-capacity Packet Network	Maged E. Beshai Paul F. Daspit	9/29/00 (USPO)	09/672,816 (6,771,651)	N/A	N/A
01-2	Multi-Grained Network	Maged E. Beshai Richard Vickers	9/28/00 (USPO)	09/671,140	N/A	N/A
01-3	Compact Segmentation of Variable-Size-Packet Streams	Maged E. Beshai Ernst A. Munter	12/14/00 (USPO)	09/735,471	N/A	N/A

01-4	Multi-Channel Sharing in a High-Capacity Network	Maged E. Beshai Ernst A. Munter Richard Vickers	12/21/00 (USPO)	09/742,229	N/A	N/A
01-5	Burst Switching in a High-Capacity Network	Maged E. Beshai Richard Vickers	12/29/00 (USPO)	09/750,071	N/A	1310303.1
01-6	Switched Channel-Band Network	Maged E. Beshai	7/9/01 (USPO)	09/960,959	N/A	Pending
02-1	Universal Edge Node	Maged E. Beshai Paul F. Daspit Harold G. Edwards	12/26/01 (USPO)	10/025,982	N/A	N/A

Technology Transfer

For Part 1A, technology transfer considerations are limited to Task 4 – Flat Network Architecture / PetaWeb. However, until transfer strategies mature, a determination of which technologies may be of value to transition cannot be made. The specific technology transition strategies that could be employed to effect the transition will be made once the technologies have been more fully explored.

Technical Achievements for Milestones 29 - 38

6. Introduction

This part covers the period June 2002 – March 2004, Milestones 29-38 for Tasks 6-10. The original tasks of this program were (a) build an open research network capable of demonstrating the full benefits of very high capacity networks and the optical transport, switching and network management and control strategies on which it is built; and (b) conduct research on Flat Network architectures. TIA Modification No. P00012 dated June 25, 2002, removed selected Milestones associated with Tasks 1-5 and added Milestones 29-37 associated with Tasks 6-10. Milestone 38 extended Program Management Task 5. Performance of the TIA was completed on 3/31/2004.

7. Summary – Participating Organizations

As directed, Nortel Networks established subcontracts with University of California Berkeley (UCB) and University of Southern California Information Sciences Institute (USC/ISI) for specific tasks during Part 1B of the program. The remaining tasks were performed by Nortel Networks.

Task 6 – Active Networks Research (Nortel Networks and UC Berkeley)

Task 7 – Berkeley Realtime-Application Sensor Network Demonstration (UC Berkeley)

Task 8 – SuperNet Maintenance and Research Support (USC/ISI)

Task 9 – 10Gbps Ethernet Connectivity Over BoSSNet (Nortel Networks)

Task 10 – Advanced Networking Infrastructure (Nortel Networks)

Task 11 – Program Management (Nortel Networks)

Task 6 – Active Networks Research

Principal Investigators:

Nortel Networks – Tal Lavian (tlavian@nortelnetworks.com)

UC Berkeley – Professor Randy Katz (randy@cs.ucb.edu)

The goal of the Active Nets program was to create a new network platform flexible and extensible at runtime to accommodate the rapid evolution and deployment of network

technologies and to provide the increasingly sophisticated programmable services demanded by defense applications. In order to enable programmable services network devices must have fast performance and be equipped with the networking programmability. The mission is therefore “to enable value-added services to be deployed across network elements dynamically, safely and conveniently, without degrading the performance and reliability of the network”. The Nortel Networks Technology Center’s OpenetLab was created to develop a programmable networking platform to support the mission. Openet is a service-based internetworking infrastructure that delivers such programmability to diversified network devices.

Openet comprises ORE (Oplet Runtime Environment) and hierarchical services from low-level system to high-level application, and provides a neutral service-based programmability to network devices. Moreover, Openet allows the deployment of customer network services including Active Networks-based services on current commercial network platforms. Active Flow Manipulation (AFM), a key enabling technology within Openet, enhances the control intelligence of network devices through programmability. The AFM mechanism involves two abstraction levels in the control plane. One is the level at which a node can aggregate its data into traffic flows, and the other is the level at which it can perform simple actions on the traffic flows. The abstraction allows one to think and act in terms of primitive flows whose characteristics can be identified and whose behaviors can be altered by primitive actions in real-time. With AFM, customer network services can exercise active network control by identifying specific flows and applying particular actions thereby altering network behavior in real-time. These services are dynamically deployed in the CPU-based control plane and are closely coupled with the silicon-based forwarding plane of the network node, without negatively impacting forwarding performance. The effectiveness of the AFM approach was demonstrated by four experimental applications on commercial network nodes [1]. It was also demonstrated that the AFM-based control-plane network services enhance functionality of commercial hardware without impeding performance of the forwarding plane. However, if an active service as in the case of the ORE Active Node Transfer System (ANTS) application requires processing packets on the control plane, the service performance depends strongly on the performance of the platform CPU.

Having developed the Openet and worked on active applications over the last three years, we have concluded that Openet could become a powerful platform for ActiveNetworks technology transfer if several underlying hardware limitations can be removed. Firstly, the hardware can only perform L2–L4 filtering, finer differentiation has to be done by the active service on the control plane itself. Secondly, the computational power of the control plane is limited by the CPU of the router whose optimal design does not allow much use of the CPU cycles on tasks other than routing. Thirdly, it is difficult to partition the resources and allocate them securely to different execution environments with the current structure. To overcome the first limitation, we implemented Openet on the Nortel Alteon

Webswitch that is capable of content filtering. To overcome the second and the third limitations, we deployed a new platform that has built-in filtering mechanisms and that includes extendible high performance computing planes. A detailed, final report is available upon request from AFRL. Relevant published papers:

Tal Lavian, Doan B. Hoang, Franco Travostino, Phil Wang, Siva Subramanian, I. Monga: *An Extensible, Programmable, Commercial-Grade Platform for Internet Service Architecture*. IEEE Transactions on Systems, Man, and Cybernetics, Part C 34(1): 58-68 (2004)

Siva Subramanian, Phil Wang, Ramesh Durairaj, Jennifer Rasimas, Franco Travostino, Tal Lavian, Doan B. Hoang: *Practical Active Network Services within Content-Aware Gateways*. <http://csdl.computer.org/comp/proceedings/dance/2002/1564/00/15640344abs.htm>

[1] Tal Lavian, Phil Wang, Franco Travostino, Siva Subramanian, Doan Hoang, Vijak Sethaput and David Culler, **Enabling Active Flow Manipulation In Silicon-based Network Forwarding Engines**, *IEEE Journal of Communication and Networks*, March 2001.

Task 7 – Berkeley Realtime-Application Network Demonstration (BRAND): Networked MEMs CAD and SensorWeb

Principal Investigator:

UC Berkeley – Dr. Shankar Sastry (sastry@eecs.ucb.edu)

Networked MEMs CAD The goal of this project was a demonstration of a Networked MEMS CAD system including remote operation of simulation, measurement systems and data repositories. Realtime visualization on the microscale is also an important future application of high performance networks. This demonstration addresses the potential use of high performance networks as a tool in the design and manufacture of microsystems. The rising investments required to build foundries that support successive advances in sophistication brings networks into play, with their ability to allow operators to remotely share scarce facilities. This is particularly true in the R+D community, where a community dispersed across many different organizations in varied locations require access to very centralized facilities. In the emerging MEMS area, the continued rapid development of new devices and their applications depends on having adequate CAD tools for the design, simulation, measurement, and evaluation of devices. There is no robust and widely available system like SPICE in the integrated circuit world. Our answer to this is Berkeley's SUGAR system for MEMS design. Our goal is to close the design loop by enabling design, simulation, fabrication, comparison of measurement with simulation or other data, diagnostics, and then re-design. Measurements are done by a variety of devices at Berkeley, CMU and other sites, all connected by the SuperNet, and will support a variety of outside users. The measurement devices are capable of producing nanometer real-time 3D images of operating MEMS devices, as well as simpler measurements. The model is that a user would be able to use all the facilities (simulation, measurement, and data repositories) remotely at high speed. Speed is important because of the enormous measurement files produced and the ability to control and observe the devices being measured in real-time. SUGAR will be available as a web service on the Berkeley Millennium, a multi-hundred processor cluster where jobs will migrate to least

loaded machine. We will interface the Berkeley measurement devices and simulation tools to the Matisse Resource Manager via the SuperNet. In addition to simulating and measuring devices from our own local users, we will identify and support outside users to make sure that the system supports their needs.

SensorWeb The goal of this project was to demonstrate a SensorWeb and associated communications system at a military location including sensor emplacement, sensor networking, data exfiltration, and human interface. In February, 2004, a demonstration of SensorWeb took place at China Lake Naval Air Weapons Station. A fence-line of advanced radar-based sensor nodes was tested from air-deployment through desert operation in a proof-of-concept field test. DARPA sponsored the test performed by Systima (launch platform), Advantaca (sensor network) and UC Berkeley (modeling and simulation). Sixty four sensor nodes were launched from a programmable launcher mounted on a helicopter platform. The sensors automatically self-activated, wirelessly reported their GPS position by short-range radio, then reported through a sat-link and remained on station to report test incursions into the sensor fence-line viewed on a website. Fence-line probability of detection of human targets was measured in excess of 96%. A final report is available upon request from AFRL.

Task 8 - SuperNet Maintenance and Research Support (SMRS)

Principal Investigator:

USC/ISI-E – Tom Lehman (tlehman@isi.edu)

This project contained three sub-tasks:

- Develop the monitoring infrastructure to characterize TCP and UDP flows across high bandwidth-delay product networks and compare that to the packet flow as perceived by the host protocol stack.
 - o USC/ISI recommendations: Our observations indicate that packet loss in the backbone is a rare event. Packet reordering is relatively common at the flow level: almost half of the tests saw some instance of packet reordering. Despite this, the fraction of packets reordered is low, with the overwhelming majority of flows seeing less than 1 packet in 1000 reordered. We evaluated current methodologies for classification of reordering, and demonstrate the importance of choosing an appropriate metric for the application or protocol. For example, the data shows that a flow with less than 0.4% reordering (measured by a simple percent metric) can generate either no spurious congestion events for TCP, or several thousand spurious congestion events, depending on the pattern of reordering. The resulting impact to TCP performance can be large.
- Provide support to the current SuperNet researchers as well as others who may use SuperNet during this performance period. This support is typically in the form of network build out, configuration, maintenance, monitoring, and experiment setup.

- Provide support for the planned introduction of 10 Gigabit Ethernet Switch/Routers on SuperNet. This support to include planning changes to the current fiber configuration, procurement of fiber path to support the BoSSnet connection and coordinating interconnection of BoSSnet and SuperNet.

All tasks were completed. A detailed final report is available upon request from AFRL.

Task 9 – 10Gbps Ethernet Connectivity Over BoSSnet

Principal Investigator:

Nortel Networks – Paul Daspit (pdaspit@nortelnetworks.com)

The purpose of this project was to provide two 10Gb/s Ethernet switch/routers for use on BoSSnet. In August and September, 2002, two Nortel Networks Passport 8600 Switch/Routers were shipped to BoSSnet. Each unit was configured with one 10 Gigabit Ethernet interface, 16 Gigabit Ethernet interfaces and 48 Fast Ethernet interfaces. Personnel from Nortel and USC/ISI-East installed, configured and commissioned the equipment in the USC/ISI-East laboratory. USC/ISI-East conducted performance testing for approximately 30 days. Using 3 gigabit connections, up to 2.7Gbps was successfully demonstrated over the 10 Gigabit Ethernet interface. Detailed operational performance of this equipment in BoSSnet was not a reportable requirement under this task.

Task 10 – Advanced Networking Infrastructure

Principal Investigators:

Nortel Networks – Tal Lavian (tlavian@nortelnetworks.com)

iCAIR – Joel Mambretti (j-mambretti@northwestern.edu)

This project contained two sub-tasks: DWDM-RAM and AFRL/JBI Enhancements.

- **DWDM-RAM.** The goal of this project was to demonstrate a service platform that closely integrates large-scale data services with dynamic lightpath allocation, through a network resource middleware service, using an OGSA-compliant interface allowing direct access by external applications.
 - o The proposed DWDM-RAM architecture is OGSA-compliant, supports on-demand and scheduled data management, data transport service provisioning on individually addressable (controllable) light path (wavelength) channels, within a meshed DWDM dynamically switched optical network, managed by out-of-band control plane methods for each layer of service: data services, provisioning services, and network data transport. As shown in Figure 7-1, at the highest layer, there is an encapsulated set of services defined by the GGF OGSA model. It is possible for this service to be managed also by the Grid Reservation and

Allocation Manager (GRAM), which is part of the Globus Grid Services Toolkit.




Figure 7-1 – DWDM-RAM versus Layered Grid Architecture (adapted from [1])

[1] T. DeFanti, M. Brown, Eds., "NSF CISE Grand Challenges in e-Science Workshop Report," National Science Foundation Directorate for Computer and Information Science and Engineering (CISE), Advanced Networking Infrastructure and Research Division, (Grant ANI 9980480), University of Illinois at Chicago, Dec 5-6, 2001.

- o The next service layer is a Data Management Service (DMS), which accepts application requests. A mechanism is provided to ensure that such requests can be policy-and-access filtered. The requests allow for identified data blocks to be specified for transport, either on-demand or scheduled. Through information from lower level processes, DMS has a complete knowledge of available network resources, data resources and existing schedules. A Scheduled Network Service maintains the reservations, priorities, queuing, event management, and synchronization. The Data Resources Services component has an understanding of the data resources, including potentially replication and version control. The DMS provides matching between resource requests and resources available. The Scheduled Network Service can dynamically respond to new demands and new network conditions. Similarly, the architecture allows for the creation of a Storage Resource Service that would ensure required storage resources are available for use. A Scheduled Storage Resource Service could be implemented as a parallel service that would provide for advance storage reservations. All of these processes rely on an internal system of messaging.

- The DWDM-RAM architecture illustrates how a dynamically re-configurable optical plane can provide on-demand wavelength services to a class of users for which today's Internet cannot satisfy. Data-intensive applications are being developed today that far exceed the capacity of not only legacy communication systems, but also the current generation of high performance data communication systems. DWDM-RAM addresses this requirement head-on. Specifically, DWDM-RAM demonstrates an ad-hoc GRID Data Management Service with the necessary middleware to signal path requests into the optical layer, set optical cross-connects, request layer 2/3 switch connections and VPN configurations, confirm network and device configurations, initiate the data flow, and when the session is complete, tear down the connections, cross-connects and optical path.
 - Nortel Networks recommendations: Continuation of research into the dynamic provisioning of regional and wide-area multi-domain lightpaths is critical to meeting emerging demands for data intensive applications. DWDM-RAM represents a capability that does not exist today using traditional Internet services. All tasks were completed. A detailed final report is available upon request from AFRL.
- **AFRL/JBI Enhancement.** The Joint Battlespace Infosphere (JBI) research project objective was initiated to enable AFRL researchers to explore enhancements to the AFRL JBI testbed using the Nortel Networks Alteon platform. A key part of the AFRL JBI testbed is the JBI Publish-Subscribe (PUB-SUB) system which provides information publication and subscription services.
 - This project investigated how acceleration technologies of the Alteon platform could be exploited by the JBI testbed. The Alteon platform possesses not only a capability of content switching and network processing in real-time, but also a powerful computational plane for accommodating intelligent services. The research involved incorporation of the Alteon platform into the JBI testbed for the AFRL research team to exploit Alteon intelligent software implementations controlling content-based switch manipulation and network processing. The goal was to realize JBI-specific processing and computation requirements, particularly in support of the PUB-SUB architecture. The Alteon performance boost is only realized by combining the Alteon hardware (Alteon Application Switch 2424 and Alteon Integrated Service Director (ISD) with a Nortel Application Acceleration Protocol (NAAP) which is designed and tailored for a specific application.
 - A number of previously developed NAAP applications served to baseline the performance capability of the Alteon system and provide a measure of process acceleration for the specific application. These NAAP applications include: HTTP re-direction, bandwidth adaptation, DOS blocker, TCP firewall, content replication, dynamic server load balancing; and smart storage.

- o An Alteon NAAP software development kit (SDK) was provided to AFRL for use in exploring how the Alteon system might accelerate the PUB-SUB processing. The latest NAAP SDK version 1.0.3 was provided on October 10, 2003. At about the same time, AFRL identified that the JBI PUB-SUB architecture depended on XML based content switching/routing at or near line speeds. A NAAP application with this capability has not been developed by Nortel and such development is beyond the scope of this research project.
- o Relevant published papers:
[DWDM-RAM: Enabling Grid Services with Dynamic Optical Networks,](#)
 S. Figueira, S. Naiksatam, H. Cohen, D. Cutrell, D. Gutierrez, D. B. Hoang, T. Lavian, J. Mambretti, S. Merrill, and F. Travostino, IEEE CCGRID/GAN'04 - Workshop on Grid and Advanced Networks, Chicago, April 2004

[Modeling Advance Reservation Requests in Optical Network Grids,](#)
 Sumit Naiksatam, Silvia Figueira, Stephen A. Chiappari, and Nirdosh Bhatnagar, SCU COEN Tech Report 2004-07-12A, July 2004.

Task 11 – Program Management

NTONC Program Manager:

Nortel Networks – Paul Daspit (pdaspit@nortelnetworks.com)

Nortel Networks provided overall technical and program management for the NTON program and established sub-agreements with UC Berkeley (Task 7) and USC/ISI (Task 8) as directed. Nortel Networks also provided management of the NTON Consortium to include services of the NTON Consortium Financial Officer and legal council.

8. Intellectual Property (July 2002 – March 2004)

Patent Filings

Over the 21-month program July 2002 through March 2004, Nortel Networks made 8 patent application filings, all of them regarding aspects of DWDM-RAM research in accordance with the intellectual property terms of the TIA. Enclosure G provides the final DD Form 882 Report of Inventions and Subcontracts (July 2002 – March 2004).

Table 8-1. NTONC Report of Patents – July 2002 through March 2004						
				Date Filed & Patent Application Number		
CY-#	Title	Names of Inventors	Disclosure Number	USA (USPO)	Canada (CPO)	Europe (EPO)
04-4	A Distributed Visualization Architecture	Howard Cowan, Tal Lavian, Richard Brand	16861SS	Pending	Pending	Pending
04-3	Distributed Storage Network Meta - Managers	Steven Merrill, William D. Cutrell, Howard Cowan, Tal Lavian	16836SS	Pending	Pending	Pending
04-2	Extensible Network-Application Resource Messaging	Phil Wang, Franco Travostino, Tal Lavian	16803RO	Pending	Pending	Pending
04-1	Method and Aparatus for On-line Leasing or Renting of Optical Network Bandwidth for Immediate or Future Use	Steven Merrill, William D. Cutrell, Howard Cowan, Tal Lavian	16628SS	3/30/2004	Pending	Pending
03-1	DWDM-RAM: An Architecture for Data Intensive Service Enabled by Next-Generation Dynamic Optical Networks - A	William D. Cutrell, Howard Cowan, Tal Lavian	16576SS	8/8/2003	Pending	Pending
03-2	DWDM-RAM: An Architecture for Data Intensive Service Enabled by Next-Generation Dynamic Optical Networks - B	Tal Lavian, William D. Cutrell, Howard Cowan, Franco Travostino	16578SS	10/31/2003	Pending	Pending
03-3	DWDM-RAM: Enabling Data Grid Services with Dynamic Optical Networks	Tal Lavian, William D. Cutrell, Howard Cowan, Franco Travostino	16577SS	10/31/2003	Pending	Pending
03-4	Scheduled Resource Management and Optimization for Switched Underlay Networks	William D. Cutrell, Howard Cowan, Tal Lavian	16442SS	11/21/2003	Pending	Pending

Technology Transfer

For Part 1B, technology transfer considerations are limited to Task 6 – ActiveNets, Task 7 - BRAND and Task 10 – DWDM-RAM. With respect to Task 7, UC Berkeley and DARPA have established a separate channel, independent of the NTONC, for discussions regarding transfer of the technologies that pertain to Task 7 research. For Tasks 6 and 10, Nortel intends to fully develop and deploy the technologies associated with this research. However, the specific strategies that will be employed to best effect the transition, and the development and implementation phasing for the transitions have not been determined.

9. Conclusions

The NTON Program spanned a period of 66 months; September 1998 – March 2004 and covered a broad range of optical network research and other research. Additionally, the NTON Program, particularly the NTOC test bed network, enabled applications research at various universities and government agencies.

For nearly 30 months, the NTON Consortium provided a high-performance network platform to facilitate numerous government, university and industry programs and activities for applications-oriented and network-oriented research and demonstrations under DARPA's Next-Generation Internet Program. Numerous high-performance computing demonstrations were supported over NTON enabling significant applications and network experiments and demonstrations which contributed to the advancement of eScience and general network capabilities. Research organizations supported by the NTONC testbed included Jet Propulsion Laboratory, NASA Ames Research Center, Stanford Linear Accelerator Center, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Sandia National Laboratory Livermore, California Institute of Technology, San Diego Supercomputer Center, Argonne National Laboratory and University of Washington/NW Pacific GigaPop. As part of the DARPA SuperNet, NTONC supported technology and network demonstrations and research activities at the High-Performance Computing and Communications Conference (Supercomputer Tradeshow) in 1999, 2000, 2001 and 2002 as well as at the Optical Fiber Conference (OFC) 2000. Nortel Networks, GST Telecom and Lawrence Livermore National Laboratory were key contributors. Sprint ATL and BART contributed dark fiber assets.

Nortel Networks, Inc. conducted research into Next-Generation Internet Architectures and concluded that a far-reaching and powerful packet-based, distributed TDM-DWDM Next-Generation network is achievable using existing technologies. PetaWeb research examined a new network architecture that scales to very high capacities and offers controllable service quality.

UC Berkeley conducted research into deployable sensors and in a field test demonstrated air-deployable, radar-based sensor elements that organized into a self-forming network capable of detecting and reporting incursions into and through a monitored area.

Nortel Networks conducted research and demonstrated a service platform that closely integrates large-scale data services with dynamic lightpath allocation, through a network resource middleware service, using an OGSA-compliant interface allowing direct access by external applications. The proposed DWDM-RAM architecture is OGSA-compliant, supports on-demand and scheduled data management, data transport service provisioning on individually addressable (controllable) light path (wavelength) channels, within a meshed DWDM dynamically switched optical network, and managed by an out-of-band ASTN/GMPLS-based control plane. The DWDM-RAM architecture breaks new grounds in the co-mingling of application-level data resources with DWDM optical resources, yielding high-performance and highly scalable data retrievals.

Other Related Documents

Related Documents for Milestones 1-28:

- A. “PetaWeb – Building Block for a Yottabit-per-Second Network” (Task 4 Final Report Executive Summary).
- B. “PetaWeb Architecture,” Technical Paper, *Networks 2000 Conference*, Toronto, Canada.
- C. “Emulation of A Vast Adaptive Network,” Technical Paper, *Networks 2000 Conference*, Toronto, Canada.
- D. “Routing Issues In Interconnecting IP Networks With the PetaWeb,” Technical Paper, *Networks 2000 Conference*, Toronto, Canada.
- E. “Courteous Routing,” Technical Paper, *Networks 2000 Conference*, Toronto, Canada.
- F. “A Comparison of Two Optical Core Networks,” *Journal of Optical Networking*, January, 2002
- G. DD Form 882 Report of Inventions and Subcontracts (September 1998 – June 2002).

Related Documents for Milestones 29-38:

- H. ActiveNets Final Report
- I. BRAND – Bald Camel Demo Final Report
- J. SMRS Final Report
- K. DWDM-RAM Final Report
- L. DD Form 882 Report of Inventions and Subcontracts (July 2002 – March 2004).

List of Symbols, Abbreviations and Acronyms

3-D	Three dimension
Abilene	National research network connecting 200+ colleges and universities
ABONE	Active Network Backbone
ADM	Add-Drop Multiplex
AFRL	Air Force Research Laboratory
ANL	Argonne National Laboratory
ANTS	Active Node Transfer System
ARC	Ames Research Center
ASCI	Accelerated Strategic Computing Initiative
ASTN	Automatically Switched Transport Network
ATL	Advanced Technology Laboratory
ATM	Asynchronous Transfer Mode
BART	Bay Area Rapid Transit
BoSSnet	Boston South Network
BRAND	Berkeley Realtime-Application Network Demonstration
CAD	Computer-Aided Design
CARIN	Collaborative Advanced Internet Research Network
CDMA	Code Division Multiple Address
CMU	Carnegie Mellon University
COLO	Co-location
DARPA	Defense Advanced Research Projects Agency
DCC	Dallas Convention Center
DOD	Department of Defense
DOE	Department of Energy
DOS	Denial of Service
DREN	Defense Research and Experimental Network
DWDM	Dense Wave Division Multiplex
EDFA	Erbium-doped Fiber Amplifier
ESnet	Energy Sciences Network
Exabit	10^{18} bits
GB	Gigabyte (10^9 Bytes or 8×10^9 bits)
Gb/s	Gigabit per second
GE	Gigabit Ethernet

Gigabit	10^9 bits
GMPLS	Generalized Multiprotocol Label Switching
GSFC	Goddard Space Flight Center
HDTV	High-definition Television
HPSS	High-performance Storage System
HSCC	High-Speed Connectivity Consortium
IETF	Internet Engineering Task Force
IP	Internet Protocol
iSD	Integrated Service Director
JB	Joint Battlespace Infosphere
JPL	Jet Propulsion Laboratory
KM	Kilometer
Lambda (λ)	Wavelength symbol
LBNL	Lawrence Berkeley National Laboratory
LH	Long-Haul
LLNL	Lawrence Livermore National Laboratory
Mb/s	Megabit per second
Megabit	10^6 bits
MEM	Micro Electro-Mechanical
MPLS	Multiprotocol Label Switching
NAAP	Nortel Application Acceleration Protocol
NASA	National Aeronautics and Space Administration
NCSA	National Center for Supercomputer Applications
NIH	National Institute of Health
NRL	Naval Research Laboratory
NTONC	National Transparent Optical Network Consortium
OFC2000	Optical Fiber Conference 2000
OGSA	Open Grid Services Architecture
OMM	Optical Micro-Machine
ORE	Oplet Runtime Environment
OXC	Optical Cross-Connect
PB	PetaByte (10^{15} bytes or 8×10^{15} bits)
Petabit	10^{15} bits
POP	Point of Presence
POS	Packet over SONET

PUB-SUB	Publish-Subscribe
QOS	Quality of Service
SC00	Supercomputer Conference 2000
SC99	Supercomputer Conference 1999
SCinet	Supercomputer Conference Internetworking
SDK	Software Design Kit
SDSC	San Diego Supercomputer Center
SLAC	Stanford Linear Accelerator Center
SMRS	SuperNet Maintenance and Research Support
SNLL	Sandia National Laboratory Livermore
SONET	Synchronous Optical Network
TCP	Transmission Control Protocol
Terabit	10^{12} bits
TIA	Technology Investment Agreement
TSR	Terabit Switch Router (Avici, Inc.)
UCB	University of California, Berkeley
UDP	User Datagram Protocol
USC/ISI	University of Southern California/Information Sciences Institute
VPN	Virtual Private Network
WDM	Wavelength Division Multiplex
XML	Extensible Markup Language
Yottabit	10^{24} bits