Available online at www.sciencedirect.com
SCIENCE DIRECT
Future Generation Computer Systems 21 (2005) 199–208
FGCS
FUTURE GENERATION COMPUTER SYSTEMS
www.elsevier.com/locate/future

Active and logistical networking for grid computing: the e-Toile architecture☆

Alessandro Bassi[a,*], Micah Beck[a], Jean-Patrick Gelas[b], Laurent Lefèvre[b], Terry Moore[a], James Plank[a], Pascale Vicat-Blanc Primet[b]

[a] *LoCI Laboratory, University of Tennessee, 203 Claxton Building, 37996-3450 Knoxville, TN, USA*
[b] *INRIA LIP (UMR CRNS, INRIA, ENS, UCB Lyon 1) Ecole Normale Supérieure de Lyon, 46, allée d'Italie, 69364 Lyon Cedex 07, France*

Available online 14 November 2004

Abstract

While active networks provide new solutions for the deployment of dynamic services in the network, exposing network processing resources, logistical networking focuses on exposing storage resources inside networks, optimizing the global scheduling of data transport, data storage and computation. In this paper, we show how active and logistical environments working together can improve grid middleware, and provide new and innovative high-level services for grid applications. We have experimented with this approach by combining the Internet Backplane Protocol suite with the Tamanoir Active Node environment. Our target architecture is the French e-Toile Grid infrastructure based on a high performance backbone (Vraiment Trés Haut Débit, VTHD).

Keywords: Active networks; Logistical storage; Grid; e-Toile

1. Introduction

The emergence of grid computing as the ultimate solution for the scientific computing community has driven attention and research in the recent past, especially since the proliferation of high performance network capabilities. Within the still fuzzy borders of grid computing—as the definition of grid is still under discussion despite highly qualified efforts [1,2], many initiatives saw the light. The common purpose is to aggregate geographically distant machines and to allow them to work together to solve large problems of any nature. The most notable efforts within this area are Globus [3] and the TeraGrid [4] initiative in U.S. and the 1ST DataGRID [5], EGEE [6] projects in Europe.

The optimization of data transfers within a grid, which we define as High Performance Grid Networking, is the focus of our interest in this area. As the load,

☆ More information on e-Toile project is available at http://www.urec.cnrs.fr/etoile/.

* Corresponding author.
E-mail addresses: abassi@cs.utk.edu (A. Bassi), mbeck@cs.utk.edu (M. Beck), Laurent.Lefevre@inria.fr (L. Lefèvre), tmoore@cs.utk.edu (T. Moore), plank@cs.utk.edu (J. Plank), Pascale.Primet@ens-lyon.fr (P. Vicat-Blanc Primet).

0167-739X/$ – see front matter © 2004 Published by Elsevier B.V.
doi:10.1016/j.future.2004.09.029

the capacity, and the availability of network links used during data transfers may heavily affect grid application performances, a fully functional grid is dependent on the nature and quality of the underlying network. Appropriate performance, together with a high degree of flexibility, are therefore key factors of a successful Grid project. One of the important characteristic of data transfers within a grid environment is that the spectrum of exchanged data volume can spread over 9 orders of magnitude, ranging from a few bytes in interactive traffic up to terabytes scale bulk data transfer. Requirements in terms of reliability and delay, and instability of the network load, can also be very heterogeneous, mainly due to IP connectionless technology and unreliability of the transfers. This can lead to instability in task scheduling and poor overall performances.

To introduce our work in this area, and present our research efforts within the e-Toile [7] grid environment, we would like to step back and examine some of the fundamental shortcomings of current network services. In particular, we would like to focus attention on the *encapsulation* of current network services, while our approach allows low-level functionalities to be *exposed* to higher levels.

The architectural design of the e-Toile grid platform, based on a very high speed network using active and logistical networking technologies, allows study of the communication limits of existing services and protocols, and the validation of more efficient approaches—approaches that aim to carry gigabit performance to the grid user level and take into consideration the specific needs of grid flows. These goals are achievable through services, which are made possible by exposing the functionality of underlying equipment.

While in the context of the e-Toile project different innovative approaches are explored (such as network performance measurement, differentiated services evaluation, high performance and enhanced transport protocols, active networking technology and services, and logistical networking), this paper will focus on the last two aspects active and logistical networking.

The paper is therefore organized as follows: in Section 2 we will describe our driving philosophy about network services, Section 3 will focus on the concept of an active grid, in Section 4 we will describe the e-Toile grid platform, and Section 5 will cover the involved frameworks for active and logistical networking and their interaction.

2. Exposing network services

To the extent that the scientific computing community is already using the network as a computer, the Internet provides a ubiquitous communication substrate connecting its components (with routers acting as special-purpose elements invisible in the architecture), while network servers provide all access to storage and computation. Illustrations of such servers and services are plentiful: FTP, NFS, and AFS provide access to storage; Condor, NetSolve, Ninf provide lightweight access to processing; HTTP provides access to both; GRAM provides access to heavyweight computing resources; LDAP provides access to directory services; and so on. What is notable about these instances, and is equally true of almost all the other cases we could add to the list, is that they represent relatively *encapsulated network services*. By encapsulated network service we mean an architecture implementing functionalities that do not closely model the underlying network resource, but must be implemented by aggregating the resource and/or applying significant additional logic in its utilization.

The best effort delivery of datagrams at the IP level, on the other hand, can be taken as example of a relatively *exposed* network service. An exposed network service adds enough additional abstraction to the underlying network resource to allow it to be utilized at the next higher level, but does not aggregate it or add logic beyond what is necessary for the most common and indispensable functionality that uses it.

An important difference between the two approaches emerges when we need to extend the functionality of a given service. Encapsulated services tend to be implemented by heavyweight servers and have APIs designed at a high semantic level, interposing themselves between the client and server, without providing low overhead, transparent access to the underlying resources. As a result, it can be difficult, inefficient, or in some cases impossible to build new functionality on top of such APIs. Instead, encapsulated services tend to implement new functionality through plug in modules that extend the functionality of the server, introducing new code that has access to low level interfaces within the server. These plug-in modules are the server equivalent of microcode in CISC processors, raising a familiar set of questions about access control and security for the management of such code.

Extending the functionality of an exposed service makes different demands, because exposed services have lighter weight servers and APIs designed at a simpler semantic level. These factors are conducive to low overhead and more transparent access to the underlying resources, so it tends to be much easier and more efficient to build new functionality on top of exposed services. Exposed services promote the layering of higher-level functionality on top of their APIs, either in higher-level servers or in client code. This layering of services, which is analogous to the user-level scheduling of a RISC processor by a compiler, is perhaps most familiar in the construction of a network services stack. In the world of end-to-end packet delivery, it has long been understood that TCP, a protocol with strong semantic properties (e.g. reliability and in-order delivery) can be layered on top of IP, a weak datagram delivery mechanism. By allowing IP services to retain their weak semantics, and thereby leaving the underlying communication bandwidth exposed for use by the broadest possible range of purposes, this layering has had the crucial benefit of fostering ubiquitous deployment. At the same time, in spite of the weak properties of IP datagram delivery, stronger properties like reliability and in-order delivery of packets can be achieved through the fundamental mechanism of retransmitting IP packets. Retransmission controlled by a higher layer protocol, combined with protocol state maintained at the endpoints, overcomes non-delivery of packets. All non-transient conditions that interrupt the reliable, in-order flow of packets can then be reduced to non-delivery. We view retransmission as an aggregation of weak IP datagram delivery services to implement a stronger TCP connection.

2.1. Exposing network storage resources: the Internet Backplane Protocol

Despite the familiarity of the exposed approach for the network services stack, it may still not be obvious how to apply it to a resource such as storage. After all, almost every technology for the access and/or management of network storage one can think of (FTP, HTTP, NFS, AFS, HPSS, GASS, SRB, NAS, etc.) encapsulates the storage behind abstractions with relatively strong semantic properties. For that reason, our research in this area had to start by creating a protocol, the Internet Backplane Protocol (IBP) [8,9], that supports the management of remote storage resources while leaving them as exposed as possible. IBP is a network service that provides an abstraction of shared network storage. Each IBP server (also called a *depot*) provides clients with access to an underlying storage resource. In order to enable sharing, the depot hides details such as disk addresses, and provides a very primitive capability-based mechanism to safeguard the integrity of data stored at the depot. IBP's low level, low overhead model of storage is designed to allow more complex structures, such as asynchronous networking primitives and file and database systems, to be built on top of the IBP API. This key aspect of the IBP storage model, the capacity of allocating space on a shared network resource, is analogous to performing a C-like malloc, or memory allocation, call on an Internet resource. However, there are notable differences in the semantics that the service provides, for instance the fact that storage allocations are time-limited, which is a countermeasure to problems such as resource monopolization or storage leaks.

With IBP in place, the question becomes how easy or difficult it is to layer storage services with strong semantic properties on top of the weak underlying storage resources provided by IBP depots. Our experience shows that the answer varies between different cases. In some cases (e.g. building a file abstraction), earlier models can be followed and the design is relatively straightforward; in other cases (e.g. point-to-multipoint communication) the benefits of distributed control makes the encapsulated approach much easier to deploy.

2.2. Exposing network processing resources: Tamanoir

The growing interest in the active networking field might be seen as a natural consequence of the difficulties experienced when integrating existing and new technologies into a shared network infrastructure. In the active networking vision, routers or any network equipment (like gateways or proxies) within the network can perform computations on user data in transit, thereby exposing their computation capabilities and making them accessible to the end user by supplying programs called *services* that modify the behavior of the network. These kinds of routers, called *active nodes* (or *active routers*), show a greater flexibility towards the deployment of new functionalities and are more adapted to the architecture, the users, and the service providers' requirements than traditional networking. The price to pay to have this greater flexibility is, generally, an increased attention needed towards security and performance.

The Tamanoir [10,11] framework is a high performance active environment based on active edge routers, able to handle different applications and various data streams at the same time. The two main transport protocols, TCP and UDP, are supported by the TAN for carrying data. The Active Network Encapsulated Protocol (ANEP) [12] format is used to send data in a format understandable by other active networks nodes.

One of the characteristics of Tamanoir that differentiates this approach from other active networks is the use of exposed logistical storage for optimizing end-user services requests, especially in terms of performance. As explained in Section 5, each Tamanoir node can take advantage not only of IBP depots located on the same node, but also can leverage any depot connected to the TAN by backbones such as the Logistical Backbone.

3. Concepts of an active grid

Our attention towards an active grid paradigm was driven by the complaints of grid application designers about standard network characteristics, such as reliable packet transport between nodes using the TCP/IP protocol suite, not being suited for typical grid applications. The active grid architecture we envision is based on a virtual topology of active network nodes spread on programmable routers of the network. Active routers with logistical support are deployed on the network periphery and allow data stream processing and storage, either in an explicit way (following a request by the application) or encapsulated one. Considering that the future of WAN backbones lies in all-optical equipment, we concentrate active operations on routers and nodes mapped at the network periphery. Active Logistical Nodes (ALN) are connected between each other, and each ALN manages communications for a small subset of grid nodes. Grid data streams cross the active layer twice, before and after passing through the passive backbone.

The communications needs of grid applications can be improved by the use of an active grid architecture, especially in the following areas:

- *Application deployment*: To deploy applications widely and in an efficient way, active reliable multicast protocols are needed to optimize the source code or binary deployment and the task mapping on the grid configuration accordingly to resource managers and load-balancing tools. Active multicast optimizes the transport of applications in any form (source code, binaries, or bytecode) by minimizing the number of messages in the network. Active nodes can deploy dedicated multicast protocols and guarantee the reliability of deployment by using the ALN storage capabilities.
- *Grid support*: The active architecture can provide information to the grid framework about network state and task mapping. Active Logistical Nodes must be open and easily coupled with all grid environment requirements. ALN can implement permanent grid support services to generate control streams between the active network layer and the grid environment.
- *Wide-area parallel processing*: With the emergence of grid parallel applications, tasks will need to communicate by sending computing data streams with QoS requests. An active grid architecture must also guarantee an efficient data transport to minimize the software latency of communications. Active nodes deploy dynamic services to handle data streams: QoS, data compression, 'on the fly' data aggregation. . .
- *Coupled (meta) applications*: The active architecture must provide heterogeneity of services applied on data streams, such as data conversion services. End to end QoS dynamic services will be deployed on active nodes to guarantee an efficient data transport (in terms of delay and bandwidth).

Most of services needed by grid environments, such as high performance transport, dynamic topology adapting, QoS, on-the-fly data compression, data encryption, data multicast, data conversion, and error management could be easily and efficiently deployed on-demand in an active grid architecture.

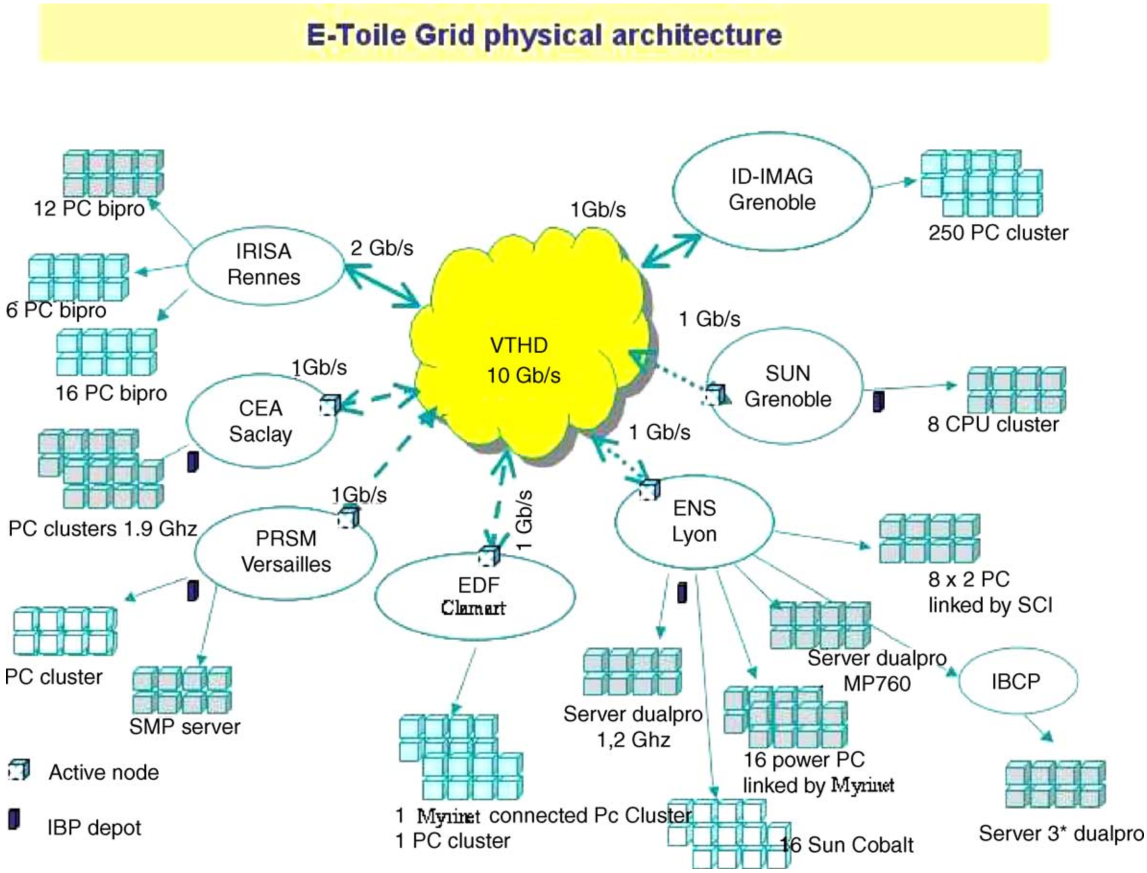

Fig. 1. Physical architecture of the e-Toile grid.

4. The e-Toile grid platform

The e-Toile project, funded by the French Ministry of Research in the realm of the Réseau National des Technologies Logicielles (RNTL) [13] initiative, focuses on three complementary objectives:

- To build an experimental high performance grid platform that scales to meet nation-wide needs and geographical distribution, providing a high-performance network and software support for the ACI-GRID [14] initiative.
- To develop original grid services, in order to go beyond the various limits of existing middleware, and to exploit completely the services and capacities offered by a very high performance network. The e-Toile middleware integrates the most recent and relevant works of the French computer science laboratories (INRIA, CNRS) on the area of enhanced communication services.
- To evaluate the deployment cost of chosen computing intensive and data-intensive applications and to measure the gain obtained thanks to the grid.

The partners of the project, providing various resources like servers, PC clusters, storage space, workstation and federating the available tools in an integrated platform, are INRIA (Reso, Remap, Apache, Paris), CNRS (PRISM, IBCP), Communication and Systems, SUN Labs Europe, EDF (Electricite De France, French Electricity Board) and CEA (Atomic Energy Center).

The e-Toile middleware relies on existing middleware, in particular Globus and the Sun Grid Engine, and integrates independent building blocks developed in the context of ACI-GRID. The most innovative grid

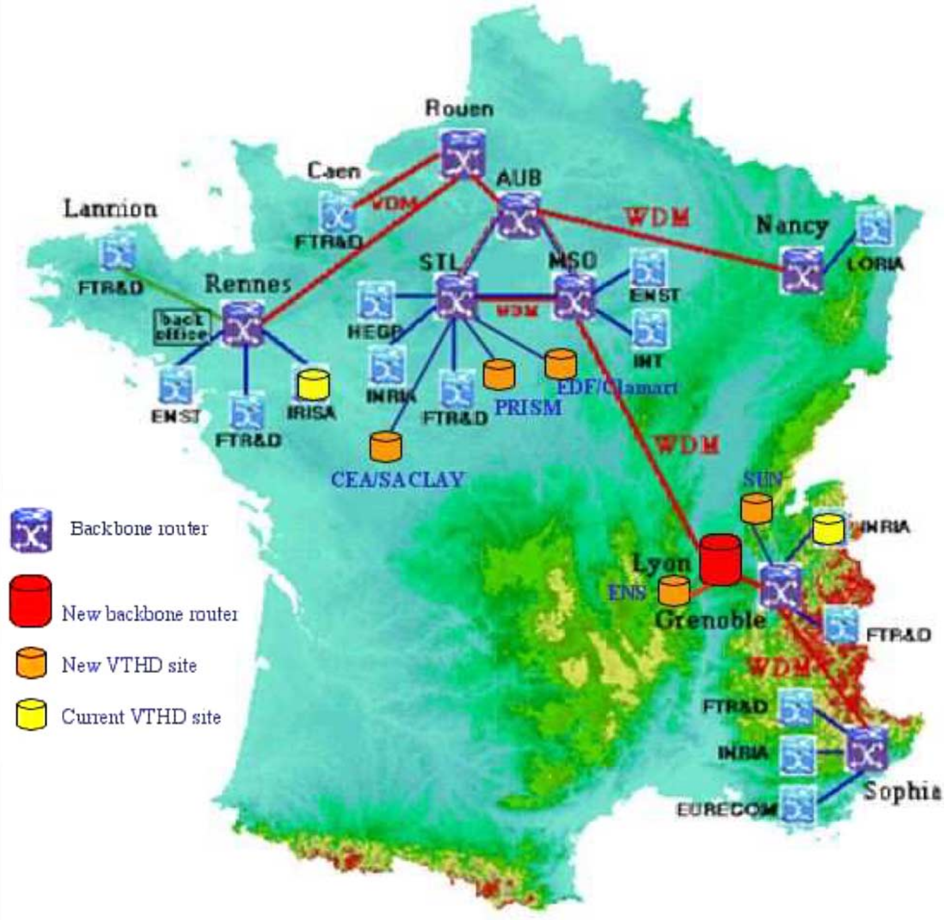

Fig. 2. The VTHD infrastructure with the nodes specifically added for the e-Toile project.

components are in the areas of grid services (in terms of resource allocation), performance measurement, security, communication paradigms, distributed storage and active networking.

Fig. 1 represents the architecture of the testbed with the resources currently interconnected. The platform is composed of two distinct testbeds: a development testbed, in order to allow testing of new middleware components, and a production testbed for grid applications. In order to provide secured access from the early system set-up, the first version of the platform activates the Globus v2.2 toolkit, with one certification authority.

The Vraiment Trés Haut Débit (VTHD) [15] network infrastructure interconnects the grid nodes with 1 or 2 Gbps links. This existing French high performance research backbone has been extended to all participant sites and also to the CERN in Geneva, Switzerland, as illustrated in Fig. 2.

Network monitoring and performance measurement tools, developed by INRIA and CNRS within the European DataGRID project are deployed in e-Toile. The MapCenter [16] Grid status visualization tool permits users and managers to visualize the available resources and services in real time. These tools allow the partners to follow the incremental deployment of the two testbeds.

4.1. High Performance Grid Networking in e-Toile

One of the original aims of the e-Toile project is to focus on the High Performance Grid Networking dimension and to evaluate the benefits grid middleware

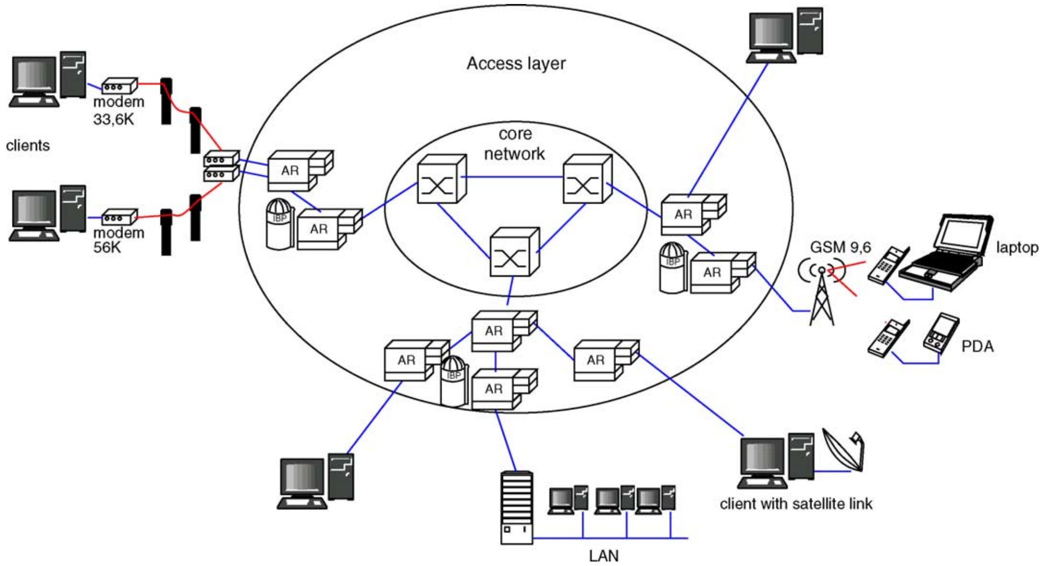

Fig. 3. Typical heterogeneous network architecture. The active and logistical nodes are located in the periphery of the network, to provide high performance and high-level services at the same time.

and grid applications can extract from enhanced networking technology. The performance problem of the grid network can be studied from different points of view:

- Measuring and monitoring the end-to-end performances helps to characterize the links and the network behavior. Network cost functions and forecasts, based on such measurement information, allow the upper abstraction level to build optimization and adaptation algorithms.
- Evaluating the advantages of differentiated services, like Premium or Less than Assured Services, offered by the network infrastructure for specific grid flows.
- Creating enhanced and programmable transport protocols to optimize heterogeneous data transfers within the grid (Fig. 3).

Active network nodes running the Tamanoir environment and IBP depots have been deployed in the e-Toile infrastructure. These generical facilities are offered to grid middleware and grid application to build more efficient and customized transport services. Several service prototypes have been studied: Reliable Multicast Transport: Dyram, Dynamic QoS control and management: QoSINUS [17]. Next section describes the Tamanoir and IBP principles.

5. Active and logistical resource deployment in e-Toile

As logistical and active networking play a fundamental role in the e-Toile architecture, the deployment of logistical and active nodes is of primary importance for the success of the project. In particular, active nodes running Tamanoir will be initially installed in five of the seven sites; those nodes will also have small IBP depots, to be used mainly by the Tamanoir services. In addition to those depots, four sites will also be hosting dedicated IBP depot machines, with an aggregate storage capacity on the order of terabytes.

Each service willing to use logistical storage has to instantiate its own IBP client classes in order to communicate with an IBP depot. These classes provide constructors and methods to create capabilities on any IBP depot, with which the Tamanoir service can write, read, and manage data remotely on the IBP depot.

A first and very basic TAN service, called *IBP_Service* uses an *IBP_store* operation to redirect the beginning of a data stream towards the IBP depot. The IBP service checks as well the presence of the required service each time that a new packet arrives, and if so, an *IBP_load* operation is done to redirect all the data cached in the IBP depot towards the service able to process, route, and forward these data efficiently. The only difference between the IBP_Service and any other service lies in the load-time, which is done at boot time for the IBP_Service, in order to be able to cache data immediately.

The use of an IBP depot on each TAN should allow for the use of redundancy to provide for reliability in the storage of data. If data are replicated on multiple servers and one of them becomes either out of order or unreachable, data should still be downloadable from another server transparently.

Explicit routing and optimized scheduling are also possible through logistical active nodes. Several experiments made by the IBP team show that, in order to transmit huge amounts of data as fast as possible, the path chosen by "classic" routers might show performance well below that of the optimal path. Adding a staging point in the middle of two fast connections, knowing the topology of the underlying network, often improves performance dramatically. Unfortunately, current transport protocols encapsulate the path, offering an end-to-end service that is not well adapted to the performance needs a grid application often has. A Tamanoir service, upon reception of a data stream, could store the stream on a number of different depots, optimizing the total transfer time by explicitly routing the data packets towards faster connections and staging data at mid-points or at the boundaries of high-speed connections. In case of a change in the application scheduling, data stored at those intermediate nodes could be re-routed towards the new processing node in a much more efficient way.

6. Conclusions and future works

A grid empowered with active and logistical networking can not only improve significantly its global performance, but, foremost, can provide new and innovative services to grid applications and middleware, giving easy and direct control over high level services, such as reliable multicast and active QoS.

We ran several experiments mixing active and logistical equipment, and the latency measured [18] when adding ANEP packet caching in an active stream shows a very small overhead. Moreover, with the deployment of the new generation of RAM-based IBP depots, we can expect that the actual overhead of caching data will be significantly smaller than the current one.

In addition to the technical challenges outlined in this article, the direction of our research is towards the integration of our active and logistical technologies with existing middleware, such as Globus and the Sun Grid Engine, giving them the ability to encapsulate active and logistical technologies for their internal needs, and to expose active and logistical services to upper layer applications.

Our plan is also to integrate these technologies further, especially by adding active and intelligent services to the logistical layer in fields such as efficient and reliable data deployment for the grid. By using the *Data Mover* (DM) plug-in module to send a request to a Tamanoir Active Node, IBP depots could take advantage of any active services and could improve data transport.

Acknowledgements

The authors wish to thank all participants of the RNTL e-Toile project. This work is supported and carried out within the framework of the RNTL e-Toile project and the RNRT VTHDH++ project. The work of the researchers at the University of Tennessee was funded by the National Science Foundation (NSF) Next Generation Software Program under grant #ACI-0204007, the NSF Middleware Program under ANI-0222945, and the Department of Energy Scientific Discovery through Advanced Computing Program under grant #DE-FC02-01ER25465; additional support for infrastructure was provided by the NSF CISE Research Resources program under EIA-022444 and by the Center for Information Technology Research (CITR) of the University of Tennessee.

References

[1] I. Foster, C. Kesselman, S. Tuecke, The anatomy of the grid: enabling scalable virtual organizations, Int. J. Supercomput. Appl. 15 (3) (2001).
[2] W. Gentzsch, Response to Ian Foster's "What is the grid?", Grid Today 1 (8).
[3] The Globus Project: http://www.globus.org.
[4] TeraGrid Project: http://www.teragrid.org.
[5] IST DataGrid Project: http://eu-datagrid.web.cern.ch/eu-datagrid/.
[6] EGEE: Enabling Grids for E-science in Europe: http://egee-intranet.web.cern.ch/.
[7] RNTL e-Toile French Grid Project: http://www.urec.cnrs.fr/etoile.

[8] J. Plank, A. Bassi, M. Beck, T. Moore, M. Swany, R. Wolski, Managing data storage in the network, IEEE Internet Comput. 5 (5).
[9] M. Beck, T. Moore, J. Plank, An end-to-end approach to globally scalable network storage, in: Proceedings of the ACM SIGCOMM 20002 Conference, Pittsburgh, PA, USA, 2002.
[10] J.-P. Gelas, S. El Hadri, L. Lefèvre, Towards the design of an high performance active node, Parallel Process. Lett. 13 (2).
[11] A. Galis, J.-P. Gelas, L. Lefèvre, K. Yang, Active network approach to grid management & services, in: Proceedings of the Workshop on Innovative Solutions for Grid Computing—ICCS 2003 Conference, Melbourne, Australia, 2003, pp. 1103–1113 (INCS 2658, ISBN 3-540-40195-4).
[12] S. Alexander, B. Braden, C. Gunter, A. Jackson, A. Keromytis, G. Minden, D. Wetherall, Active network encapsulation protocol (ANEP), rFC Draft, Category: Experimental, July 1997.
[13] RNTL: http://www.industrie.gouv.fr/rntl.
[14] ACI Grid Initiative: http://www-sop.inria.fr/aci/grid/public/.
[15] Very High Broadband RNRT VTHD Project: http://www.vthd.org.
[16] MapCenter: http://ccwp7.in2p3.fr/mapcenter/.
[17] P. Vicat-Blanc Primet, F. Chanussot, End to end network quality of service in grid environments: the qosinus approach, in: Proceedings of the First International Workshop on Networks for Grid: GridNets, San Jose, CA, USA, 2004.
[18] A. Bassi, J.-P. Gelas, L. Lefèvre, Tamanoir-IBP: adding storage to active networks, in: Active Middleware Services, IEEE Computer Society, Edinburgh, Scotland, 2002, pp. 27–24 (ISBN 0-7695-1721-8).

Alessandro Bassi despite being tempted by other disciplines decided to explore the esoteric world of computer science, because of its tension between creativity and mathematical rigor. He managed to survive Milan and its university, and becoming passionate of artificial intelligence, soft computing and software engineering. As many of us, he lent his abilities to the private sector to survive, but managed to unchain his spirit again and since summer 2000 he is involved in the Internet Backplane Protocol and other advanced networking stuff proudly developed for the University of Tennessee.

Micah Beck began his research career in distributed operating systems at Bell Laboratories and received his Ph.D. in Computer Science from Cornell University (1992) in the area of parallelizing compilers. He then joined the faculty of the Computer Science Department at the University of Tennessee, where he is currently an Associate Professor working in distributed high performance computing, networking and storage. An active participant in the Internet2 project since 1997, Dr. Beck is chair of their Network Storage Special Interest Group. In 2000, he joined with Internet2 participants drawn from industry and academia to find and serve as Chief Scientist of Lokomo Systems, a company in the area of advanced content distribution. In 2002, he joined with colleagues within the University of Tennessee's Computer Science Department to find and now directs the Logistical Computing and Internetworking Laboratory, consisting of 20 faculty, students and research staff.

Jean-Patrick Gelas received a graduate degree (DEA) of Computer Science in Ecole Normale Supérieure de Lyon (France) in 2000. He prepared next his thesis thanks to a national grant awarded in order of merit by the French Ministry of Research and Technology (MENRT) in the LIP/INRIA Laboratory in the team RESO and received his Ph.D. degree from University of Lyon (France) in December 2003. His research interests include high performance active networks, distributed computing and operating system. His main interest concerns next generation networks (like grid computing), added value network services and heterogeneous networks and terminals. He also has a large teaching experience over his three years as Ph.D. student. He is currently employed as Research Associate in the Logistical Computing and Internetworking (LoCI) Laboratory of the Computer Science Department of the University of Tennessee in USA.

Laurent Lefevre obtained his Ph.D. in Computer Science in January 1997 at LIP Laboratory (Laboratoire Informatique du Parallelisme) in Ecole Normale Supérieure de Lyon (ENS-Lyon), France. From 1997 to 2001, he was Assistant Professor in Computer Science in Lyon 1 University and in the RESAM Laboratory (High Performance Networks and Multimedia Application Support Lab.). Since 2001, he is permanent researcher in Computer Science at INRIA (The French Institute for Research in Computer Science and Control). He is member of the RESO team (high performance networks, protocols and services) from LIP Laboratory in Lyon, France. He has organized several conferences in high performance networking and computing and he is member of several program committees. He has co-authored more than 30 papers published in refereed journals and conference proceedings. He is member of IEEE and takes part in several research projects. His research interests include: high performance active networks, active services, grid and cluster computing network support, active grid, distributed shared memory systems and data consistency.

Terry Moore is currently Associate Director of the Logistical Computing and Internetworking Laboratory and the Center for Information Technology Research at the Computer Science Department at University of Tennessee. He received his B.A. in Philosophy from Purdue University and his Ph.D. in Philosophy from the University of North Carolina, Chapel Hill. His research interests include network storage and especially logistical networking, in which he has been a leader since the inception of the field in 1997.

James S. Plank received his B.S. from Yale University in 1988 and his Ph.D. from Princeton University in 1993. He is currently an Associate Professor in the Computer Science Department at the University of Tennessee. His research interests are in fault-tolerance, network storage, erasure coding and grid computing. He is an Associate Editor of IEEE Transactions on Parallel and Distributed Computing, and a member of the IEEE Computer Society.

Pascale Vicat-Blanc Primet received a Ph.D. degree in CS from INSA Lyon, France in 1988. From 1989 to 2001, she worked as a researcher and a lecturer at Ecole Centrale de Lyon (ECL), France. In 2001, she joined INRIA as a researcher. In 2001–2002, she managed the Networking Workpackage of the European IST DataGrid project. In 2002–2003, she was the Scientific Coordinator of e-Toile, the National RNTL grid platform project funded by the French Research Ministry. Based at Ecole Normale Supérieure de Lyon (ENS-Lyon), she is currently a member of the LIP Laboratory and leads the RESO team at INRIA. This team is specialized in communication protocols and software optimization for high-speed networks. Pascale's research interests include high-performance grid and cluster networking, active networks, TCP and QoS. She is a co-chair of the GGF Data Transport Research Group, a co-chair of PFLDnet 2005 and a member of the Steering Committee of the GRID5000 Project.