Available online at www.sciencedirect.com
ScienceDirect
Optical Switching and Networking 5 (2008) 150–158
Optical Switching and Networking
www.elsevier.com/locate/osn

# Quality-of-service-aware fault tolerance for grid-enabled applications

Luca Valcarenghi[a,*], Filippo Cugini[b], Francesco Paolucci[a], Piero Castoldi[a]

[a] *Scuola Superiore Sant'Anna, Pisa, Italy*
[b] *CNIT, Pisa, Italy*

Received 31 August 2007; received in revised form 30 December 2007; accepted 24 January 2008
Available online 13 February 2008

**Abstract**

This study first reviews how grid-enabled applications can be provided with fault tolerance. Existing methods, implemented either in the grid application/middleware or in a Generalized Multi-Protocol Label Switching (GMPLS)-based network, are outlined. Then, the paper shows the advantages of integrating application/middleware fault-tolerant schemes, such as service replication, with GMPLS network-layer fault-tolerant schemes, such as path restoration. An integrated fault-tolerant scheme is capable of providing flexible QoS-aware fault tolerance while minimizing the necessary computational and network resources. In the end, the implementation of the proposed integrated scheme in a Video-on-Demand (VoD) application is experimentally validated.


*Keywords:* Fault tolerance; Grid; Replication; Restoration; Multimedia

## 1. Introduction

The current communications network evolution envisages a single network infrastructure to support fixed and mobile voice, video and data services (i.e., quad-play services). In this scenario legacy (e.g., phone calls) and emerging (e.g., multimedia streaming and grid-enabled) applications share the same network resources while requiring different Quality of Service (QoS) in terms of bandwidth, delay, jitter, and fault tolerance (i.e., resilience). This study focuses on a specific class of emerging applications, such as grid-enabled applications, and on the problem of guaranteeing fault tolerance and preserving the application requested interconnection bandwidth while minimizing the required resources. The proposed approach leverages the integration of application/middleware layer resilience schemes, typical of grid-enabled applications, with current network layer resilience schemes.

Commonly, the term grid-enabled application is utilized to refer to distributed computing applications (e.g., the SETI@home project) in which computational resources, available in the Internet, are shared to perform computations that would otherwise be too time consuming if run on stand-alone systems. In this case, however, because of the low bandwidth requirement and interactivity of the application, the

* Corresponding author. Tel.: +39 0505492138; fax: +39 0505492194.
*E-mail addresses:* luca.valcarenghi@sssup.it (L. Valcarenghi), filippo.cugini@cnit.it (F. Cugini), fr.paolucci@sssup.it (F. Paolucci), castoldi@sssup.it (P. Castoldi).
*URL:* http://www.ircphonet.it (L. Valcarenghi).

1573-4277/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.osn.2008.01.004

best effort IP protocol is sufficient to provide the application with the required QoS. However other grid-enabled applications, such as distributed visualization sessions, large data streaming coordinated with job execution, and networked supercomputing [1,2], require to be supported by a network capable of providing specific QoS in terms of bandwidth, delay, jitter, and resilience. This latter type of grid-enabled applications are considered in this study and, in particular, the focus is on applications involving a large real-time data streaming that has specific bandwidth and resilience requirements, such as distributed visualization sessions involving data rendering and display.

Currently, grid-enabled application fault tolerance is assured by schemes already utilized in cluster grid, implemented in the application or in the middleware, or by common IP layer fault-tolerant schemes, such as IP dynamic rerouting. However, current fault-tolerant schemes do not take into account the need for recovering application failures while maintaining application QoS guarantees, i.e. they are not QoS-aware. In particular, they do not consider network status parameters, such as network capacity occupation, and potential contention for network resources during the failure recovery procedure. Thus, they might be capable of recovering connectivity among the sites involved in the grid application but they might not recover also pre-failure communication QoS guarantees, such as minimum guaranteed bandwidth.

High-capacity optical networks based on Generalized Multi-Protocol Label Switching (GMPLS) offer the means for implementing QoS-aware resilient schemes, such as bandwidth-guaranteed Label Switched Path (LSP) dynamic restoration. However, the current absence of integration between GMPLS and application/middleware fault-tolerant schemes leads to potential inefficiencies. If service replication is the utilized application/middleware fault-tolerant scheme, several scenarios in which the missing integration undermines a successful recovery attempt and an efficient network resource utilization can be pinpointed.

First of all, if just replication is utilized for recovering from all the failures (i.e., both network and application level failures), a set of static QoS-guaranteed LSPs should be set up in advance for assuring QoS-guaranteed connectivity. However, in failure-free conditions, most of the LSPs would be idle, thus potentially wasting useful network resources, as in the case of $1 + 1$ protection schemes. On the other hand, if just LSP dynamic restoration is utilized, application level failures (e.g., system crash) cannot be recovered. In addition, the set up of QoS-guaranteed backup LSPs to the original service location might fail due to the lack of resources shared with other applications requiring the same QoS-guaranteed connectivity. This would make the QoS-guaranteed connectivity recovery attempt unsuccessful. In case both replication and LSP dynamic restoration are available but not integrated, several inefficiencies may arise. For instance, if after switching to a service replica upon failure, the set up of QoS-guaranteed LSP to the replica location is not automatically triggered, the pre-failure QoS-guaranteed connectivity might not be recovered. Moreover the switching to a service replica could be triggered to recover an LSP from a network failure instead of resorting to the more natural QoS-guaranteed LSP dynamic restoration. Thus, the recovery procedure would potentially require longer time, waste network resources, and imply service synchronization and restart.

In this study, typical fault-tolerant schemes, separately utilized by grid-enabled application and GMPLS, are reviewed. Then, both a numerical and an experimental evaluation of the proposed fault-tolerant scheme that integrates service replication and LSP path restoration is performed. To perform the evaluation a video-on-demand service (VoD) between one client and multiple video streaming servers is considered. This allows also to show the applicability of the proposed integrated scheme to generic multimedia streaming applications supported by quad-play networks. The integrated recovery scheme numerical evaluation is based on a modified version of the Mixed Integer Linear Programming (MILP) model presented in [3]. It shows that, if the number of replicas is limited, placing them in a network-aware fashion (i.e., by taking into account the nodal degree) allows to efficiently exploit them and to recover a high percentage of failed connections. The experimental validation shows that the integrated scheme is capable of overcoming heterogeneous failures without affecting the application quality perceived by the client. In addition, the considered scheme can be implemented in networks based on commercial routers supporting (G)MPLS without the need of any router protocol or software modification and applied to generic streaming applications.

## 2. Fault tolerance

Grid-enabled application failures include not only physical network failures, such as fibre link fault, but also computational resource failures, such as application crash [4]. In grid computing, fault-tolerant schemes