

Experimental Demonstration of the Enhanced Optical User Network Interface (O-UNI) Protocol

Wenhua Jiao, Yong Wang and Ludi Zheng

Bell Labs Research China, Lucent Technologies,

*15/F, Aerospace Greatwall Building, No.30 Hai Dian Nan Lu, Beijing 100080, P.R.China
Tel. +86-10-68748088 ext. 8860, Fax. +86-10-68748360, E_mail address: wjiao@lucent.com*

Abstract: O-UNI is defined as the optical interface between the service provider (optical transport network) and user client equipment. We have built and demonstrated an O-UNI testbed between IPv6 routers through optical network in IPv6 Global Summit, May 2002. This testbed enhances UNI 1.0 with three features: TNA addressing and resolution, control plan recovery, and IPv6 support in both control and data plane. It is more robust and suitable for clients to access to the optical networks.

I. INTRODUCTION

With the development of optical multiplexing technologies such as Dense Wavelength Division Multiplexing (DWDM), optical network layer has evolved quickly to provide transport services to interconnect clients such as IP routers, ATM switches, SONET/SDH equipments, etc. User Network Interface (UNI) is defined as the interface between the service provider (optical transport network) and user client equipment. The Optical Internetworking Forum (OIF) implementation agreement on realizing UNI 1.0 has been well supported by the industry [1]. OIF UNI 1.0 defines the set of services, the signaling protocols used to invoke the services, the mechanisms used to transport signaling messages and the auto-discovery procedures that aid signaling, all of which are to be implemented by client and transport network equipment vendors to support UNI 1.0. UNI 1.0 mainly focuses on signaling (LDP and RSVP-TE) for service invocation, whereas TNA (Transport Network Assignment) addressing and resolution, control plane and data plane recovery, etc. still need to be further studied [2].

Bell Labs Research China has implemented UNI 1.0 and successfully tested it in both IPv4 and IPv6 networks. It covers the basic functions of UNI 1.0, including automatic neighbor discovery, service discovery, and connection management. In addition, we enhance it with the following features: TNA address assignment, control plan recovery, and IPv6 support in both control and data plane. In IPv6 Global Summit, May 2002, this O-UNI testbed between IPv6 routers through optical network was established and demonstrated successfully.

II. ARCHITECTURE OF O-UNI TESTBED

The architecture of the O-UNI testbed, which conforms to OIF UNI 1.0 standard, is shown in Fig. 1. The testbed architecture is divided into data plane and control plane. The interconnection model of our O-UNI is an overlay model, with indirect service invocation configuration; UNI-N (UNI Signaling Agent – Network) and UNI-C (UNI Signaling Agent – Client) are needed at both sides. Using LDP signaling defined in UNI 1.0, connections are automatically provisioned between two IPv6 edge routers with optical interface through the optical network composed of Optical Cross-connects (OXC). The LDP operation mode in UNI is downstream on demand label advertisement with ordered control.

In O-UNI demonstration, the OXC are needed for optical network. In Fig. 2, we use two optical switches in one OXC to simulate two OXC. The Source UNI-N (SUNI-N) and Destination UNI-N (DUNI-N) control the two optical switches through the OXC control computer (controller) respectively, which works in the control plane. UNI-Ns send commands to controller through TCP sockets. After receiving the command, the controller makes a decision to turn on/off the corresponding optical switch according to the message content. OSM (Optical Service Manager) is responsible for TNA address registration and resolution.

As shown in Fig. 1, when a new connection is needed, for example, for transmitting data between the client and server, the Source UNI-C (SUNI-C) sends an LDP request message to SUNI-N. If the optical network accepts the request, the DUNI-N relays the request message to the Destination UNI-C (DUNI-C) to reserve the required bandwidth. Suppose the reservation is successful, DUNI-C sends a LDP mapping message to DUNI-N with the assigned interfaces and labels; if the optical network commits the assignment, the SUNI-N relays the mapping message to SUNI-C. When SUNI-C receives the mapping message, reserve confirmation message is forwarded to DUNI-C in a fashion like the

initial request message. After completing the above signaling process, the clients can access to the FTP or video server through the optical core network. For connection release, we support all the three kinds of

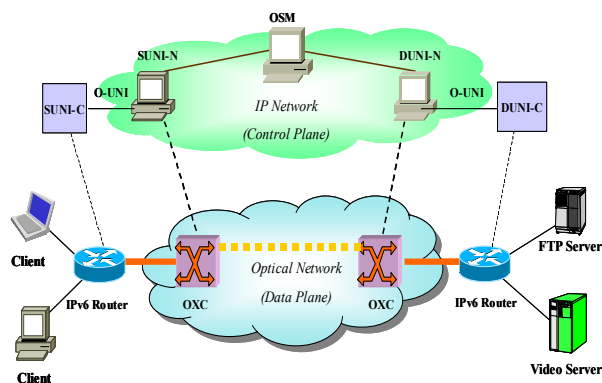


Fig. 1. Architecture of O-UNI Testbed

III. ENHANCEMENTS OF UNI 1.0

In addition to the features defined in UNI 1.0, our UNI 1.0 testbed has enhanced the following features:

3.1. TNA Addressing and Resolution

TNA is used to identify the UNI connection endpoints, which is a globally unique address assigned by the core optical network to one or more data links between a pair of UNI peers. A UNI-C must know the TNA addresses of SUNI-C and DUNI-C in order to initiate a connection request. How to assign the TNA address to a set of data links over UNI? We offer the following three approaches. The first is configuration addressing. In this way, configuration file that contains TNA allocation information in both UNI-C and UNI-N should be loaded when UNI-C and UNI-N start up. There must be some predefined rules to ensure that for the same data link, the same TNA address is assigned in both parts of UNI-C and UNI-N. Static addressing is the second approach, where a pre-assigned TNA address table exists in each UNI-N. When a UNI-N starts up, each data link is mapped with its corresponding TNA address. And when a UNI-C starts up, Neighbor Discovery (ND) and Service Discovery (SD) are conducted. In the procedures of SD, the TNA information relevant to the UNI from the peer UNI-N is obtained by UNI-C. The third method is dynamic TNA assignment. After ND and SD procedure, address registration is needed, in which the UNI-C sends UNI address registration signaling to OSM through IPCC. OSM is responsible for assigning a new TNA and sends it to the UNI just registered.

In these three methods, configuration addressing is the simplest. But it requires the cooperation and negotiation between UNI-C and UNI-N for manually configuration. Moreover, TNA address has to be assigned for those

release cases, graceful connection deletion by SUNI-C, graceful connection deletion by DUNI-C and forced connection deletion by the optical network.

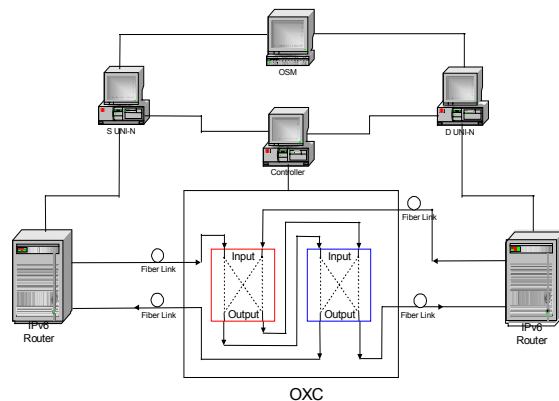


Fig. 2. The Optical Switches in OXC

interfaces that are not currently used, thus the scalability problem will arise. Both static addressing and dynamic addressing need ND and SD procedure. And in dynamic addressing, automatic addressing registration is also required. The TNA address assigned may be taken back and reused by other UNI peers if data link no longer exists between current UNI peers.

In our O-UNI testbed, dynamic TNA assignment is adopted. OSM is responsible for dynamically assigning TNA address upon a UNI is up.

3.2. Control Plane Recovery

In control plane, hello messages are exchanged periodically between UNI LDP peers to validate that the LDP peer and control link are still alive. A session will be established over IPCC between UNI peers if LDP hello procedure and session initialization succeed. KeepAlive message is sent periodically to maintain the LDP session integrity.

3.2.1 Control Plane Failure

The error cases of control plane can be divided into link failure and node failure. If one UNI peer can not get any LDP hello response from its hello adjacence after three times retransmission, hello failure event happens. In this case, hello adjacence is removed and session is torn down. UNI peers keep sending hello messages until response comes. After that, the hello adjacence is added and session procedure is restarted. If one LDP session endpoint still can not get any LDP message in the KeepAlive timeout interval, a fault event happens. A retry attempt to re-establish the original session is triggered. The failed UNI-C plays active role in setting up a new session with exponential backoff algorithm.

The control link failure should not affect the established connections between UNIs. So if link failure happens and

can not be recovered in a certain period, all established connections are kept. But all transient connection status (pending, deleting, etc.) will be aborted.

3.2.2 Control Plane Recovery Scenarios

If an IP control channel (IPCC) fails, one natural way is to switch to another protection IPCC. So if multiple IPCCs are configured, link failure in control channel will trigger an automatic IPCC switching. Over the new IPCC, the UNI-C initiates LDP session re-establishment. We have implemented this function in LMP UNI extension. But if only one control channel is configured or IPCC switching fails, we deal with this problem according to the following scenarios:

a. Control Link Failure

For the failed link between UNI-C and UNI-N, if the link can be recovered, a new LDP discovery procedure performs. And after Hello adjacencies are discovered, a new session is created. Because the transient connection is released and the stable connection information in each node is still recorded, there is no need for connection synchronization for UNI peers.

b. Control Node Failure

If a UNI-C is down (software bug, power down, hardware failure, etc), we must recover original connection information after the UNI-C is up. So after a new LDP session is established, connection status in UNI-C and UNI-N should be synchronized because the connection information in UNI-C is lost and there is no talk between them during the failure period. This synchronization procedure is done using LDP enquiry message. The response message carries the status of the specified connection and associated attributes. Any inconsistent attributes in the same connection will result in deletion of this connection.

3.3. IPv6 Support

In the O-UNI testbed, IPv6 edge routers with IPv6 protocol stacks are used as client devices. IP traffic in both control plane and data plan can be in IPv6 version. In control channel of optical network, IPv4 address is applied; and in UNI-N, IPv4 and IPv6 dual stacks are configured. Client and server computers are configured in different IPv6 subnets. Routing protocol in IPv6 edge routers can be either static routing or OSPF.

IV. DEMONSTRATION ON IPV6 GLOBAL SUMMIT

In the IPv6 Global Summit, we demonstrated the automatic provisioning for IPv6 Router with a real time high quality

Audio/Video performance on our O-UNI testbed. The source router requests to set up an optical connection between the IPv6 edge routers. After the lightpath is set up, the client can demand DVD video from the video server.

V. CONCLUSION

We enhanced UNI 1.0 with three new features: TNA addressing and resolution, control plane recovery, and IPv6 support in both control and data plane. Three TNA addressing methods are proposed, in which dynamic TNA address assignment is adopted in this testbed to be avoid of manually configurations and save the TNA address space. Control plane recovery strategies are proposed and implemented for link and node failure scenarios, which improves the control plane reliability. Moreover, this is a first O-UNI testbed in application of UNI to IPv6 router.

REFERENCES

- [1] Architecture, OAM & P, PLL & Signaling Working Groups, Optical Internetworking Forum, Multiple OIF Projects, "User Network Interface (UNI) 1.0 Signaling Specification", October 2001.
- [2] L.Andersson, et al, "LDP Specification", RFC3036, Jan. 2001.
- [3] Robert Doverspike, Jennifer Yates, "Challenges for MPLS in optical network restoration", IEEE Communications Magazine, February 2000, Vol. 39 (2), pp. 89-97
- [4] Cuangzhi Li, Jennifer Yates, Dongmei Wang, and Charles Kalmanek, "Control plane design for reliable optical networks", IEEE Communications Magazine, February 2002, vol.39 (2), pp. 90-96.
- [5] E. Mannie et al, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", IETF draft, working in progress, March 2002.