January 2022

# MULTILAYER RESTORATION OVER A CONVERGED IP AND OPTICAL NETWORK

Randy Zhang

Gabriele Maria Galimberti

Errol Roberts

## Recommended Citation

Zhang, Randy; Galimberti, Gabriele Maria; and Roberts, Errol, "MULTILAYER RESTORATION OVER A CONVERGED IP AND OPTICAL NETWORK", Technical Disclosure Commons, (January 17, 2022)
https://www.tdcommons.org/dpubs_series/4847

# MULTILAYER RESTORATION OVER A CONVERGED IP AND OPTICAL NETWORK

AUTHORS:
Randy Zhang
Gabriele Maria Galimberti
Errol Roberts

## ABSTRACT

In a converged Internet Protocol (IP) plus optical multilayer network that provides an optical restoration capability, techniques are presented herein that enable differentiated forwarding for different classes of traffic (such as, for example, high-priority traffic, low-priority traffic, and best-effort traffic) through restoration awareness at Layer 3 (L3). Aspects of the presented techniques work without the need for any inter-layer protocols, such as the Generalized Multiprotocol Label Switching (GMPLS) User Network Interface (UNI) (GMPLS-UNI). This is important because a converged network, such as a routed optical networking (RON) environment, does not dictate protocols. Aspects of the presented techniques encompass restoration-enabled differentiated forwarding (RDF) at L3. Further, aspects of the presented techniques encompass enhanced interfaces on a routed optical gateway (ROG) – transitioning from the current binary states (of Up and Down) to four states (comprising {Up, Restoration Capable}, {Down, Restoring}, {Up, Restored}, and Down) – along with an associated interface state machine.

## DETAILED DESCRIPTION

Optical restoration is a widely implemented feature that supports restoring traffic at the optical layer. It can be established either through an on-demand need or it may be constructed in advance. Currently there is no mechanism for routers to provide the enhanced routing and forwarding capability that is enabled by optical restoration.

In a traditional IP plus optical network, the Generalized Multiprotocol Label Switching (GMPLS) User Network Interface (UNI) (GMPLS-UNI) provides a dynamic means for routers to initiate and request a Layer 0 connection with constraints and requirements. However, there is no explicit coupling of the optical restoration capability through GMPLS-UNI to L3 routing and forwarding (e.g., an interior gateway protocol (IGP), segment routing, etc.) thus optical restoration cannot be directly leveraged at L3.

In a converged IP and optical network, such as a routed optical networking (RON) environment, there is no dynamic inter-layer protocol such as GMPLS-UNI between the routing and optical layers.

There is a strong desire from network operators to leverage the existing optical restoration capability to enhance the L3 routing and forwarding. A focus of the techniques presented herein is the enablement of multilayer restoration (without the need for an inter-layer protocol) and the creation of additional and more nuanced L3 interface states (that trigger differentiated L3 forwarding). In particular, under aspects of the techniques presented herein, interface states are enhanced with a restoration state, which may be used in L3 routing protocols, segment routing, or network controllers. Such an enhanced capability provides preferred L3 forwarding for high-priority traffic and more nuanced forwarding for low-priority and best-effort traffic during an optical restoration.

In a converged IP and optical network, a routed optical gateway (ROG) may connect to an optical network through a dense wavelength division multiplexing (DWDM) optical module (such as a 400 gigabit Ethernet (400G) ZR or ZRP). Optical networks are commonly constructed through reconfigurable optical add-drop multiplexer (ROADM) nodes, which may provide on-demand or pre-provisioned optical restoration. When the optical transport is attempting restoration, such as in a 1+R (i.e., working plus restoration) setup, the L3 interface is down. When restoration is successful, the L3 interface comes up. There is no difference in routers between the normal up state when the working path is up and when the restoration path is up. Traffic is being switched back and forth between these state transitions. There is a strong desire from network operators to leverage the existing restoration capability to enhance the L3 routing and forwarding. The techniques presented herein address such a desire.

A key benefit that is provided by aspects of the techniques presented herein encompasses restoration-enabled differentiated forwarding (RDF) at L3. Network operators have the choice in a policy for how RDF is implemented and for timer customization. As a result, traffic may be classified into three major categories in terms of forwarding – high-priority traffic, low-priority traffic, and best-effort traffic. High-priority traffic is treated as it is today without change. RDF allows for a more nuanced treatment of low-priority traffic and best-effort traffic. RDF is implemented through controller policies and it is enforced through quality of service (QoS) rules on the routers.

In particular, low-priority traffic and best-effort traffic may be tagged for a restoration-enabled path. When optical restoration is being attempted, these classes of traffic are candidates for not being switched (and thus unprotected by the network) as allowed by an operator's service level agreement (SLA) with customers. According to an SLA, an operator may tag both classes of traffic or only tag best-effort traffic for a restoration-enabled path.

The use of RDF provides a number of benefits. For example, RDF reduces state transitions and traffic switching for lower classes of traffic without a guaranteed QoS. Further, the switching of low-priority and best-effort traffic during optical restoration may congest or overload other parts of the network. Additionally, RDF ensures that high-priority traffic is preferred during switching and provides operators an additional tool with which to offer a more nuanced QoS and SLA. Timers may be customized to affect the wait time before switching.

Under aspects of the techniques presented herein, the interfaces on an ROG that directly connect to an optical transport network are enhanced, transitioning from the current binary states (of Up and Down) to four states (comprising {Up, Restoration Capable}, {Down, Restoring}, {Up, Restored}, and Down). Figure 1, below, depicts an exemplary interface state machine that is based on the above description.

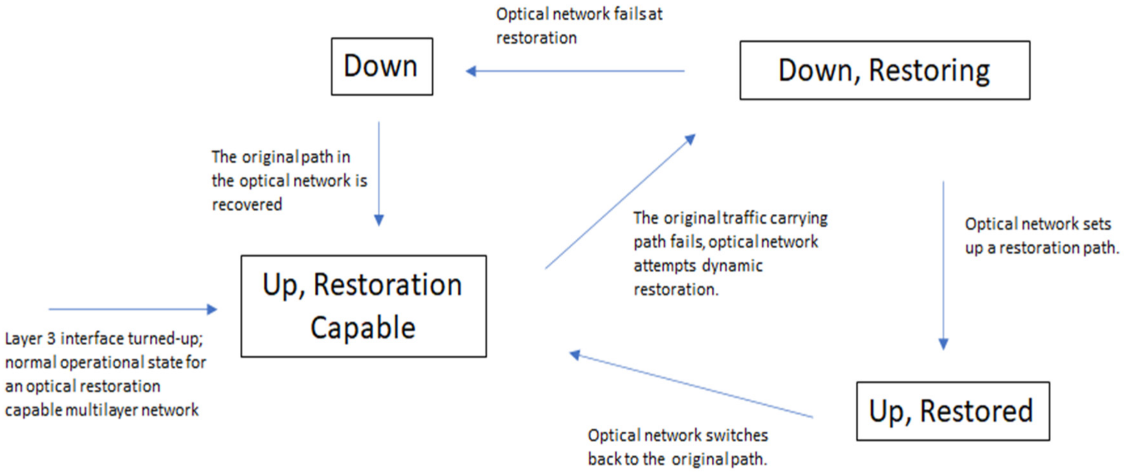

*Figure 1: Restoration-Enabled L3 Interface State Machine*

With reference to Figure 1, the {Up, Restoration Capable} state is the normal operating state when the interface is up, traffic is carried over a working path, and an optical network is restoration capable. All of the L3 traffic flows normally as it does today, similar to the current Up state. The interface may come into this state through the {Up, Restored} state or the state Down.

The {Down, Restoring} state encompasses when the original traffic-carrying path in the optical network goes down and restoration is being attempted. This state is different from the current Down state in routers, where all of the L3 traffic will be switched away from the interface. Given that restoration is ongoing, network operators now have a nuanced choice to provide differentiated services through a restoration policy. The high-priority traffic (as defined by the operator) can and should be switched as it is today. The low-priority and best-effort traffic can, according to an operator policy, stay with the interface without switching and wait for the restoration. Given that low-priority and best-effort traffic have little or no QoS guarantee, again according to operator policy choice, this minimizes unnecessary traffic switching back and forth potentially causing congestion or latency increases in other parts of a network. Traffic policy, together with timer settings, may be used to customize the deployment for how traffic is classified and handled.

In some instances, a Wait to Restore timer may be set so that if restoration does not occur within the timer's interval the policy allows for an additional setting – i.e., switch over the low-priority and best-effort traffic or continue to leave that traffic on the interface.

With the additional state information as described above according to aspects of the techniques presented herein, a dedicated low-priority path within segment routing or traffic engineering may be created. Network operators, according to policy choices, may further distinguish between the low-priority traffic and best-effort traffic and only leave the best-effort traffic in a down state.

If restoration (as described above) fails, the interface transitions, as depicted in Figurer 1, above, to the Down state. Existing policy may then be used to decide how traffic should be handled.

For example, the {Up, Restored} state encompasses when the optical network successfully creates a restoration path. All of the low-priority and best-effort traffic are immediately flowing over the interface as they were never switched away. For the high-priority traffic, network operators, according to policy choice, may decide to wait with a timer (e.g., wait to Revert-Restore) before switching the traffic to the restoration path or they may wait until the original traffic path is up. The policy can also take into account a shared risk link group (SRLG) when considering switching the priority traffic to the restoration path. Priority traffic may have an SRLG requirement, and the restored path may or may not satisfy such a requirement.

A set of external controllers may be created to provision and coordinate the above-described states. In particular, a packet network controller (PNC) and an optical network controller (ONC) are controllers for their respective network domains. A hierarchical network controller (HNC) or a software client may set policies and control a PNC and an ONC.

Figure 2, below, presents elements of an exemplary environment according to aspects of the techniques presented herein (and reflective of the above discussion).

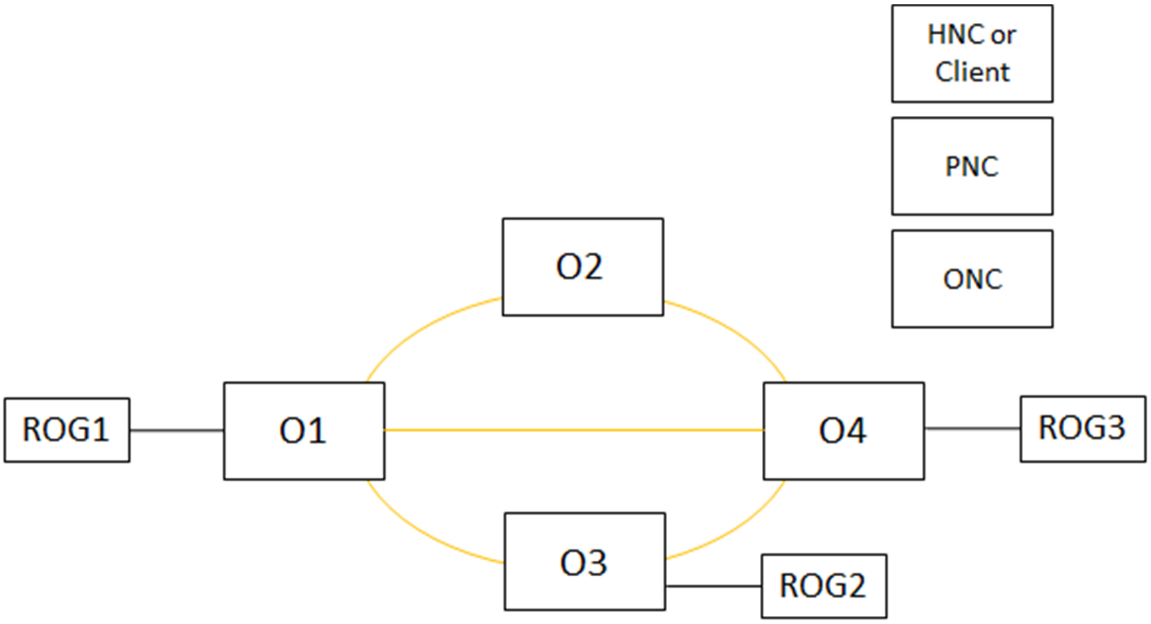

*Figure 2: Example of a converged IP and Optical network*

As depicted in Figure 2, above, routers (which are identified as ROG1, ROG2, and ROG3 in the figure) are connected over an optical transport network that consists of four nodes (which are identified as O1 through O4 in the figure).

Figure 2 illustrates two major steps or components. A first major step or component encompasses registration and capability discovery. This may include a number of activities. First, the optical network (i.e., nodes O1 through O4 in the figure) is capable of wavelength restoration and registers its capability to an external optical software-defined networking (SDN) controller (e.g., the element ONC in the figure) over its management network. Next, a ROG is a packet optical device and registers itself with an external packet SDN controller (e.g., the element PNC in the figure). Further, a hierarchical SDN controller or client (e.g., the element HNC in the figure) discovers the multilayer network through the domain controllers. According to aspects of the techniques presented herein, any existing inter-layer interface mapping discovery mechanism may be employed.

Finally, an optical capability, including a restoration capability, is discovered by the controllers and client and then advertised to the ROGs and conventional routers. (Note that conventional routers are the devices that do not directly connect to an optical transport network). There may be multiple ways in which such a capability is advertised. In cases where ROGs and conventional routers are running dynamic routing protocols, the restoration capability may be advertised into these protocols. For example, an ROG that is connecting to a restoration-capable optical network may advertise such a capability in its routing protocol updates for each such interface. L3 protocols, such as a link state protocol, may be enhanced to carry restoration capability as part of the link states. It is important to note that the interface state is {Up, Restoration Capable} during normal operation.

Further in connection with Figure 2, above, a second major step or component encompasses restoration in action. For example, in the event of a traffic carrying path failure in the optical transport, the optical nodes may trigger restoration according to instructions from the controllers, and thus the ROGs and the client may transition through the interface state machine. The ROG interface state becomes {Down, Restoring}.
Next, the ROGs may send link state updates to other routers in the case of IGP. Additionally, RDF is started. Further, once the ROGs are notified that restoration is up, the interface state is transitioned to {Up, Restored}. All of the low-priority and best-effort traffic immediately flows down the restoration path. Policy and timers allow for various

customizations. If restoration fails, the interface state transitions to Down and traffic is switched according to policy and timer settings.

Finally, when the original path becomes available again from the optical network, all of the traffic is reverted according to policy and timer settings. The interface state is transitioned to {Up, Restoration Capable}.

Thus, aspects of the techniques presented herein support creating restoration-aware interface states in the routed optical gateways and populating such states into other routers (through, for example, routing protocols or network controllers). Aspects of the presented techniques further support leveraging the enhanced interface states with restoration awareness for differentiated traffic forwarding (e.g., RDF) for different classes of L3 traffic.

Accordingly, aspects of the techniques bring awareness of the optical restoration to L3 without the need for an inter-layer protocol. Additionally, the techniques do not use the protection bandwidth for low priority traffic. Further, aspects of the techniques create an enhanced L3 interface state machine reflecting the optical restoration awareness. These states may be communicated and shared among routers through L3 protocols. While a link state protocol is not required, the techniques can leverage a link state protocol to communicate such awareness. The awareness of the optical restoration may be used by routers to create differentiated forwarding. While differentiated forwarding that is based on QoS is widely implemented in routers, the presented techniques add the ability to use the restoration awareness to enhance such forwarding. Thus, RDF as presented herein may be beneficial to network operators as it reduces network churn and congestion during optical restoration.

As noted above, aspects of the techniques presented herein enhance L3 interface states with restoration awareness of the optical layer without the need for a complex inter-layer protocol (such as GMPLS-UNI). The enhanced states provide a capability to create differentiated forwarding. More concretely, if a L3 interface is aware of the restoration capability and state in the optical layer, it can transition into the enhanced state machine as depicted in Figure 1, above. The enhanced restoration-aware interface states allow routers to create forwarding policies to handle lower classes of traffic differently from a higher class of traffic. One example would direct that lower classes of traffic are not switched to alternate paths while a restoration is being attempted. The restoration capability awareness allows for some loss of traffic with the understanding that restoration may succeed. There is a benefit to enduring some loss of low-priority traffic as this reduces network churn and congestion, and protects the high-priority traffic, thus yielding a better network SLA as a result. While differentiated forwarding based on QoS is widely implemented in routers, aspects of the techniques presented herein add the ability to use restoration awareness to enhance such forwarding.

Optical protocols such as bidirectional line-switched ring (BLSR) can use reserved protection bandwidth to carry lower priority traffic during normal operation with such traffic being dropped during protection switching. Aspects of the techniques presented herein do not make use of the protection bandwidth in the optical network for extra traffic during normal operation. All classes of traffic flow as implemented today during normal operation. But when a working path in the optical network fails and restoration is being attempted by the optical network, the restoration-aware L3 has a choice – i.e., switch all of the traffic to alternate paths (as is done today) or only switch high-priority traffic to alternate paths. According to aspects of the techniques presented herein, the low-priority traffic is not switched and will wait for the optical restoration. The benefit is that high-priority traffic is given preferential treatment and there will be less network churn and congestion in the alternate paths.

While aspects of the techniques presented herein do not require a link state protocol, they do allow for the enhancement of a link state protocol to distribute optical restoration state information in the link state update (LSU) packets. Further, the state update may be used in any routing protocols beyond link state protocols. For example, it may be added into a Border Gateway Protocol (BGP) advertisement. Aspects of the presented techniques support enhanced L3 states (incorporating restoration awareness) and use such states for differential L3 forwarding. The techniques work best in an SDN network as they bring awareness to the multilayer controllers in coordination with domain controllers.

As one example, consider an ROG that has one or more L3 interfaces that are connected to an optical network that has restoration capability. The same ROG may have one or more L3 interfaces that do not connect optical networks. The ROG may have multiple paths to a destination. In today's networks, an ROG has no awareness of an optical restoration capability (i.e., they are treated the same as the other networks) and such an awareness is not being leveraged for enhanced forwarding. But with such an awareness (as enabled through aspects of the techniques presented herein), routers (through L3 routing protocols (such as link state, distance, vector-based, etc.) or other mechanisms) have additional insight into what is happening (or to what may happen) in the optical network. That insight may be leveraged for better forwarding decisions. It is not that alternate paths are necessarily more bandwidth constrained, but during failures, routers will be adding additional traffic to the alternate paths, which may cause congestion. When the optical network once again becomes available, traffic may be switched back. These activities cause network churn. If some of the traffic is of a low priority anyway, such churn is not beneficial to network operators. Because those operators already have optical restoration, they are asking why not use it to achieve a better forwarding advantage.

In the rare case that an ROG only has interfaces to one optical network and if that optical network is down and performing restoration, then all of the traffic will be down and differentiated forwarding has no value. However, most operators design their networks with alternate paths to create diversity, even if some of the alternate paths may go through another optical network. In such a case, the techniques presented herein still apply and bring value.

There is a great deal of industry interest in leveraging the existing restoration capability for enhanced traffic forwarding. Aspects of the techniques presented herein couple the enhanced states of restoration to differentiated forwarding, where lower classes of traffic do not need to be switched while a restoration is ongoing. This reduces the network churn and congestion. As the industry is pushing for more converged IP and optical networks to reduce cost and increase agility (such as in a RON solution), such an approach will help network operators tighten the integration and achieve better network SLAs. Under such an approach centralized SDN controllers may be leveraged.

In summary, techniques have been presented herein that enable differentiated forwarding for different classes of traffic (such as, for example, high-priority traffic, low-priority traffic, and best-effort traffic) through restoration awareness at L3. Aspects of the presented techniques work without the need for any inter-layer protocols (such as, for example, the GMPLS-UNI). This is important because a converged network, such as a RON environment, does not dictate protocols. Aspects of the presented techniques encompass RDF at L3. Further aspects of the presented techniques encompass enhanced interfaces on a ROG – transitioning from the current binary states (of Up and Down) to four states (comprising {Up, Restoration Capable}, {Down, Restoring}, {Up, Restored}, and Down) – along with an associated interface state machine.