Alcatel·Lucent

Q What access and network diversity enhancements are envisioned to strengthen VPLS resiliency?

A Ethernet is rapidly emerging as the Layer 2 technology of choice to support the transformation to IP, and Carrier Ethernet VPNs based on VPLS are key in this transformation. Enterprise Chief Information Officers (CIOs), who are focused on corporate governance and business process automation, expect the network to always be available to support their business-critical applications. On the residential side, similar expectations now also exist thanks to triple play offerings of video entertainment and telephony and an increasing dependence on the Internet.

MPLS is being used to put essential "carrier" attributes into Ethernet, but operators remain focused on further improvements and cost optimization of end-to-end networking and the elimination of any residual failure conditions. For the MPLS network to fully meet these expectations for Ethernet services including VPLS, it must:

- Provide resiliency against catastrophic node failures in the core of the MPLS network
- Provide resilient access to Ethernet services delivered by the MPLS network

Techniques used in today's MPLS networks to provide resilient Ethernet services can be broadly categorized as:

- Node level redundancy
- Network level redundancy

Node level redundancy minimizes the impact of failures of particular components of a node from impacting the network or service. Non-stop routing and non-stop services on the Alcatel-Lucent 7750 SR and 7450 ESS are examples of node level redundancy.

Network level redundancy protects against failures of links or nodes in the core of the network. At the Ethernet layer, the use of link aggregation groups (LAGs), can provide both redundancy and extra capacity between two systems by combining multiple Ethernet links into a group and representing the group as a single bundle, a LAG, on the connected systems. The problem with a simple LAG is that it provides no protection against a catastrophic failure of a provider edge (PE) node. Spanning tree protocol (STP) and rapid spanning tree protocol (RSTP) could also be run between the customer premises equipment (CPE) and the provider network where multi-homing at the Ethernet layer is used. However, performance concerns have meant that service providers are reluctant to use STP or RSTP for virtual leased line (VLL) and VPLS services. At the MPLS layer, network level redundancy can be provided by MPLS fast reroute to achieve sub-50 ms protection to all of the Ethernet pseudowires (PWs) carried by an LSP. However, this is insufficient to protect against failures of the PEs or attachment circuits (AC), or the switching PEs when PW switching is used.

Two features that address these gaps, providing end-to-end protection for Ethernet services, are:

- Pseudowire redundancy.
- Access and PE redundancy using multi-chassis link aggregation (MC-LAG).

Pseudowire Redundancy

PW redundancy enables one or more standby PW to be configured to protect the traffic on an active PW. Each redundant set of PWs is associated by configuration to a single Ethernet service at each end. PW redundancy relies on extensions to PW signaling that use LDP status messages to indicate the active or standby state of a PW. When a PE signals to a remote PE that a given PW is active, and other PWs in the redundant set are signaled standby, the remote PE should use the active PW to forward packets from the AC to which it is bound. Note that although this discussion focuses on the use of PW redundancy for Ethernet VLLs, PW redundancy can also be used for other VLL types such as ATM or frame relay.

Multi-Chassis LAG

MC-LAG extends the concept of LAG sub-groups such that one end of the LAG is split between two systems instead of, for example, two router blades. MC-LAG thus provides redundant Ethernet access connectivity that extends beyond link level protection by allowing two systems to share a common LAG end point. Link aggregation control protocol (LACP) is used to manage the available LAG links in “active” and “standby” states such that only links from one PE node are active at a time, to and from the Ethernet edge device. A further MC-LAG control protocol runs only between the redundant pair of PEs. This is an IP-based protocol that synchonizes the LAG state between the MC-LAG peers. Connectivity is maintained even on complete failure of a single system. The CPE still uses standard LAG (as per IEEE 802.3-2005).

Figure 1. Multi-Chassis Link Aggregation

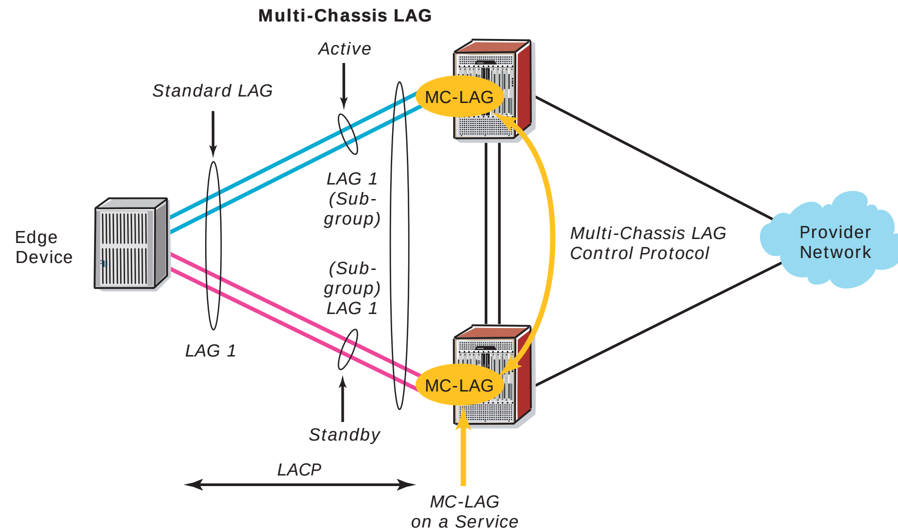

Combining PW Redundancy and MC-LAG for Enhanced Reliability

End-to-end protection for MPLS-based Ethernet services relies on the effective combination of pseudowire redundancy with MC-LAG.

Figure 2 illustrates how this works for a VLL service. CE1 and CE2 are dual homed to PE1/PE2 and PE3/PE4 respectively by Ethernet attachment circuits, and the four PEs are interconnected using Ethernet PWs. PE1/PE2 and PE3/PE4 form redundant pairs, whose active/standby state is coordinated using an IP-based inter-chassis MC-LAG protocol. Standard LACP operates between each CE and its connected PEs, such that one of the LAG sub-groups is active and one is standby at any given time. The LAG sub-group status is reflected in the PW status that each PE signals to its far-end peer PE. This determines which PW is used for forwarding. Thus, the MC-LAG state at either end of the service effectively drives the forwarding state of the PWs; the end-to-end path where the MC-LAG status of both ACs is active and both PW endpoints are active.

Figure 2. Pseudowire Redundancy and Multi-Chassis LAG on a VLL

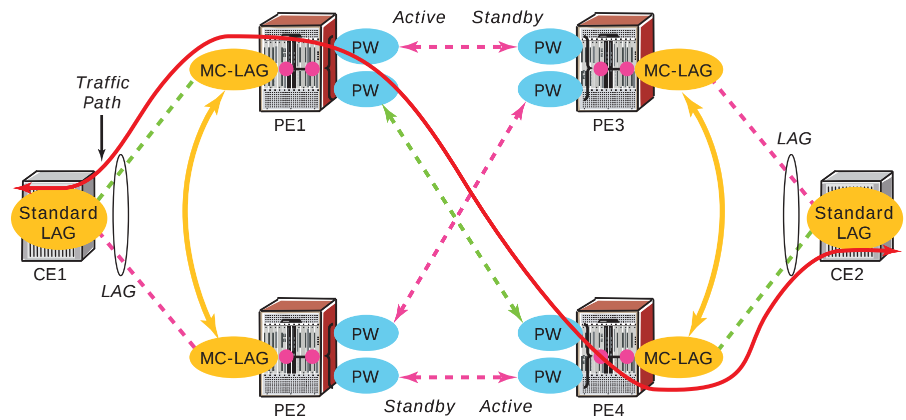

If one of the active LAG links fails, the attached PE will signal standby on its PWs. The MC-LAG protocol will also switch over to the redundant PE, which signals active on its PWs. LACP then switches the formerly standby LAG link to active. A new active path between the customer edge (CE) devices is thus created that avoids the failure.

Similarly, if an active PE fails entirely, the MC-LAG control protocol will trigger the redundant PE to become active. The newly active PE will signal active on its PWs, and use LACP to switch over the state of the LAG links to its attached CE. As above, a new active path between the CEs is thus created that avoids the failed PE.

Two key objectves are achieved through the use of MC-LAG and PW redundancy in this manner:

- The Ethernet service stays operationally up, despite the failure of an attachment circuit.
- The failover operation is transparent to the far end CE. This is important for large scale deployments where it is desirable to localize failover operations to minimize network loads and failover times.

MC-LAG can also be used to protect the access to a VPLS (see Figure 3) and for redundancy interconnection of VPLSs between metro networks. As before, standard LACP is used to select which LAG sub-group is active and which is standby. A failure of the link between CE1 and PE1 is detected by PE1 through physical layer, LACP, or Ethernet OAM mechanisms. This triggers the MC-LAG control protocol to make PE2 the active PE, which then uses LACP to make the lower LAG link active.

Figure 3. Using MC-LAG to Protect Access to a VPLS

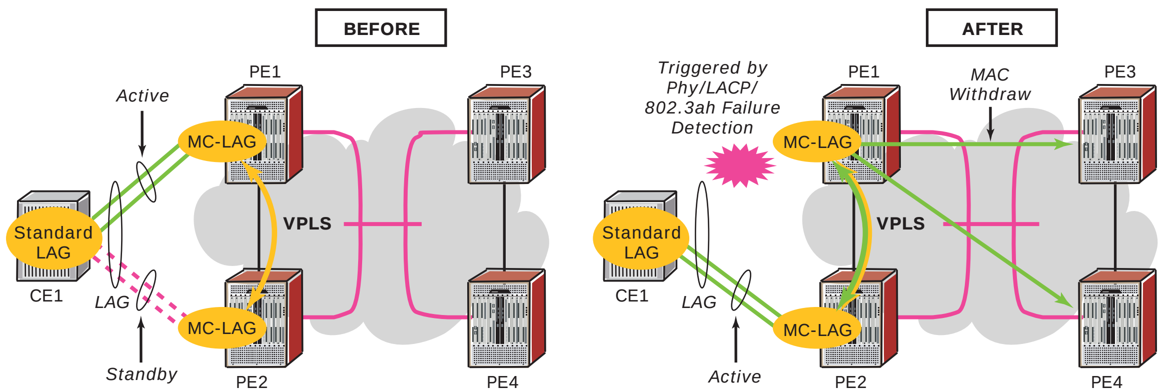

Note that stale MAC address information may be contained in the VPLS PEs following such a failover, and so PE1 sends a MAC withdraw message to its connected PEs. In VPLS, this message is carried in the LDP signaling used for the constituent PWs, and causes the PEs to remove those MAC addresses from their forwarding tables. PEs participating in the VPLS will then learn the identity of the new PE to which frames should be forwarded for CE1.

MC-LAG and pseudowire redundancy are based on standard protocols. Alcatel-Lucent is driving the standardization of the required extensions to pseudowire status signaling at the IETF in draft-muley-dutta-pwe3-redundancy-bit-01.txt and draft-muley-we3-redundancy-01.txt. Standard LAG is run on the CE, meaning that no changes should be needed to CEs that support the active/standby operation of LAG.

MC-LAG and PW redundancy are the latest additions to the Alcatel-Lucent Service Router portfolio's comprehensive range of resiliency features. These features ensure services – both business and residential – are always on.