PipelinePub.com
May, 2006 | Volume 2, Issue 10

Shifting Gears:
Migrating legacy ATM transparent LAN services to VPLS

Mustapha Aissaoui & Peter Chahal of Alcatel

Increasingly, a service provider's ability to maintain a competitive market position will depend on its ability to offer premium WAN Ethernet services such as virtual private LAN service (VPLS). For incumbents, WAN Ethernet services are no longer considered a last-resort defense but an offensive, growth play. Revenues from traditional ATM transparent LAN services (TLS) continue to shrink owing to competitive pressures and enterprise migration to a more flexible and cost-effective carrier Ethernet virtual private network (VPN) service with far higher performance.

Knowing what they need to offer, service providers are faced with the question of how legacy ATM TLS can best be migrated to VPLS.

The answer can be found with the addition of ATM bridged interfaces to a VPLS instance, which allow new applications to be deployed for Ethernet-based services. One such application provides the means for migrating ATM-based TLS over a new VPLS-based network.

What is VPLS?

VPLS, or virtual private LAN service, is a Layer 2 multipoint VPN that allows multiple sites to be connected in a single bridged domain over a provider-managed IP/MPLS network. All customer sites in a VPLS instance appear to be on the same LAN, regardless of their locations. VPLS effectively implements closed user groups via VPLS instantiation. It dedicates a separate forwarding information base (FIB) per VPLS instance in order to achieve full customer isolation. VPLS achieves a larger network FIB capacity than traditional Ethernet switching since VPLS nodes, unlike legacy Ethernet switches, do not need to learn the MAC address of a VPLS instance in which they do not participate.

VPLS allows service providers to virtualize network resources by provisioning a dedicated pseudowire (PW) mesh and a dedicated label switched path (LSP) tunnel mesh and traffic engineering paths if required on a per-VPLS instance basis.

Finally, VPLS makes use of mature multiprotocol label switching (MPLS) and pseudowire capabilities such as tunnel and PW mesh signaling, fast protection around link and node failures with fast reroute, OAM and traffic engineering.

Migrating ATM-based TLS to VPLS

Figure 1 illustrates a TLS service provided initially over an ATM network. User sites are attached to Ethernet switches or bridges with ATM uplinks. These are referred to as BR 1 and BR 2 in Figure 1. The ATM uplink is connected to an ATM switch. The Ethernet switch ports for this user form a full or partial mesh over the ATM network. In order to avoid loops over the ATM VC mesh, the Ethernet switches participate in a Spanning Tree Protocol (STP) instance and disable forwarding over VCs that create loops for this user.

Figure 1 – Extending ATM-based TLS with a VPLS

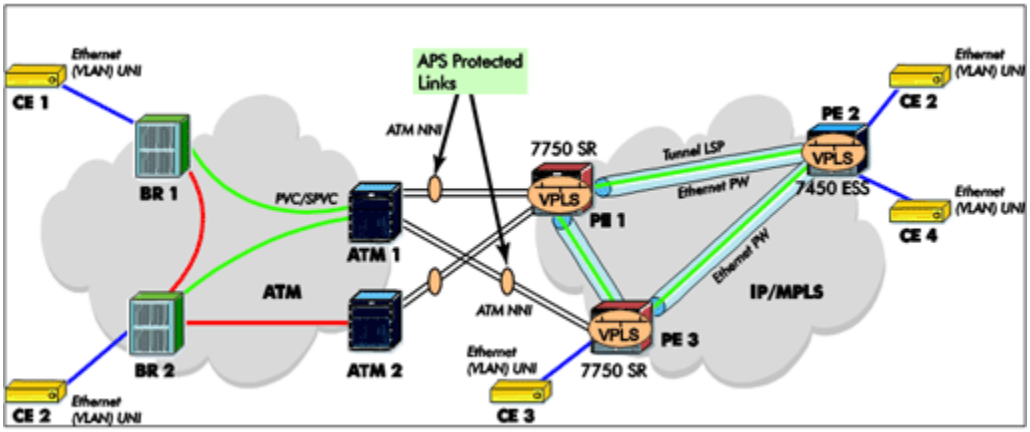

The migration of customers from the ATM-based network onto the new IP/MPLS network will occur over a transition period. During the transition period, a number of customers will have existing sites attached to the ATM network and new sites attached to the new IP/MPLS network. In order to provide connectivity between the TLS network and the VPLS network, a number of ATM VCs are configured at the boundary between the two networks. Each VC connects user ports on the Ethernet switch in the ATM network to a VPLS instance on the gateway provider edge (PE) router dedicated to this user. Multiple VPLS instances are supported to provide isolation between multiple users. Given that there is already some level of meshing in the ATM network, there is no need to configure an ATM VC for each Ethernet switch.

The Ethernet frames are sent and received on an ATM VC using RFC 2684 bridged encapsulation. The gateway PE, i.e., either PE 1 or PE 3 in Figure 1, reassembles the AAL5 protocol data unit (PDU) and processes the Ethernet frame. From that point, all other datapath operation is similar to the case of a frame received on a regular Ethernet interface terminating on a VPLS instance.

Because the ATM TLS network provides full or partial mesh with the ATM VCs, it is important to be able to configure some of the ATM VCs terminating in the PE gateway into split horizon groups to avoid loops. The VPLS instance in the gateway PE will disable forwarding between VCs that are part of the same split horizon group.

Provisioning

The provisioning steps are the same as those required for configuring a regular Ethernet interface on a VPLS instance. Let us assume that a VPLS instance is already configured in the gateway PE – PE 1 – with at least a network IP/MPLS tunnel and one Ethernet interface as a service interface. Next, the user specifies an ATM VC as a second ingress VPLS service interface and configures the slot/port/VPI/VCI identifying it.

Next, an ATM traffic descriptor should be configured and applied to the ATM VC. The user specifies the ATM traffic parameters, the ATM service category, and enables such functions as ATM policing and shaping.

Quality of Service (QoS)

Once an AAL5 PDU is reassembled by the ATM segmentation and reassembly sublayer (SAR) at the ingress ATM port in PE 1, the contained Ethernet frame is extracted. Ingress classification of the Ethernet frames received on an ATM VC allows the user to filter the received frames based on both MAC and IP criteria. An example of a MAC match criterion is the match on the 802.1 p field or a match on the source or destination MAC address field. An example of an IP match criterion is a match on the differentiated services code point (DSCP) field or a match on the source IP address field in the header of an IP packet inside the Ethernet frame. Frames that match a specified criterion can be directed to a separate forwarding class queue dedicated to this ATM VC as configured in the ingress QoS policy. Frames that did not match any of the configured criteria will be queued in the default forwarding class queue assigned to this VC as specified in the ingress QoS policy.

Frames that are received from the IP/MPLS network port are classified based on the EXP field in the shim header of the label stack.

An important aspect of providing a VPLS service over an ATM VC is the ability to achieve the desired QoS objectives for the service. When Ethernet frames are sent over an ATM VC in PE 1, the scheduling of data becomes hierarchical with two main levels: packet level scheduling and per-VC cell level scheduling.

At the first level, frames are queued on a per-CoS (or forwarding class), per-VC basis in order to achieve the proper class of service differentiation for the frames in the same VC. Each Ethernet frame is queued based on the VC dedicated forwarding class queue, as configured in the ingress classification profile. The packet level scheduling can make use of a hierarchy in order to enforce aggregate bandwidth among a group of queues feeding an ATM VC or to enforce aggregation of bandwidth across all queues of all VCs at a given customer site.

At the second level, the segmented cells are queued in per-VC queues according to the configured ATM service category and the traffic descriptor of the ATM VC. Scheduling at the ATM level enforces the priority and bandwidth sharing desired at the cell level. This architecture is illustrated in Figure 2.

Figure 2 – Hierarchical scheduling into a ATM VC

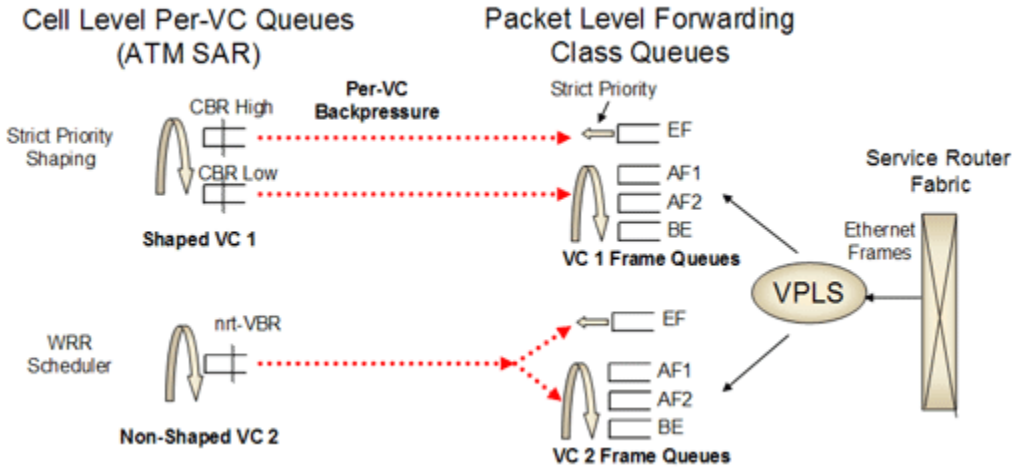

It is important to note that any discard decision should be performed exclusively at the packet level where the context for the frame forwarding class and for the 802.1p bit mapping to a forwarding class is known. When a per-VC queue backs up, a back-pressure scheme should be applied such that the frames are held in the per-forwarding class packet queues dedicated to this VC. There is, however, a condition under which a discard occurs in the ATM SAR device. This is when the CRC 32 check for a re-assembled AAL5 PDU fails. This occurs when a cell is missing or was re-ordered in the ATM network for a given AAL5 PDU.

This hierarchical scheduling of frames and cells of a given VC terminating on a VPLS instance provides the flexibility to apply policing and shaping on a per-forwarding class basis for the Ethernet frames of each VC, as well as the option to shape the aggregate cell flow into the ATM VC back into the customer site.

Resilience

Resilience of the ATM VC at the gateway PE router is provided via a SONET/SDH link automatic protection switching group. Furthermore, customer STP bridge protocol data units (BPDUs) are transparently passed from the TLS network to the VPLS network. This requires the translation of customer BPDUs from 802.1d format, which operates in the ATM TLS network, to the per-VLAN STP (PVST) format, which is typically used by customer premises equipment that is directly attached to a VPLS PE. This allows each customer to run STP end-to-end to prune the entire topology in order to select an active path and to disable loops. Resilience is further enhanced on the PEs with technology enhancements such as non-stop routing and non-stop service (NSS) to augment MPLS fast reroute (FRR) on the VPLS network to provide better than sub 50m/s availability.

Conclusion

Enterprises are transitioning from network orientation to application orientation and require an architecture that supports the convergence of sophisticated on-demand video, voice over IP (VoIP) and data applications. The best-effort performance of traditional transparent LAN services becomes insufficient: they do not scale economically, become increasingly difficult to manage, and offer limited support for multiple service classes. Service providers are seeing the writing on the wall and have already begun to structure their migration to an MPLS-enabled WAN Ethernet solution, such as VPLS.