Ethernet Transport over RPR

Vish Ramamurti

George Young

Thanks to Ralph Ballart, Kuo-Hui Liu, Mike Pepe, Eric Puetz, Yetik Serbest, Don Smith, Tom Soon, and Chin Yuan

SBC does not deploy networks this way!

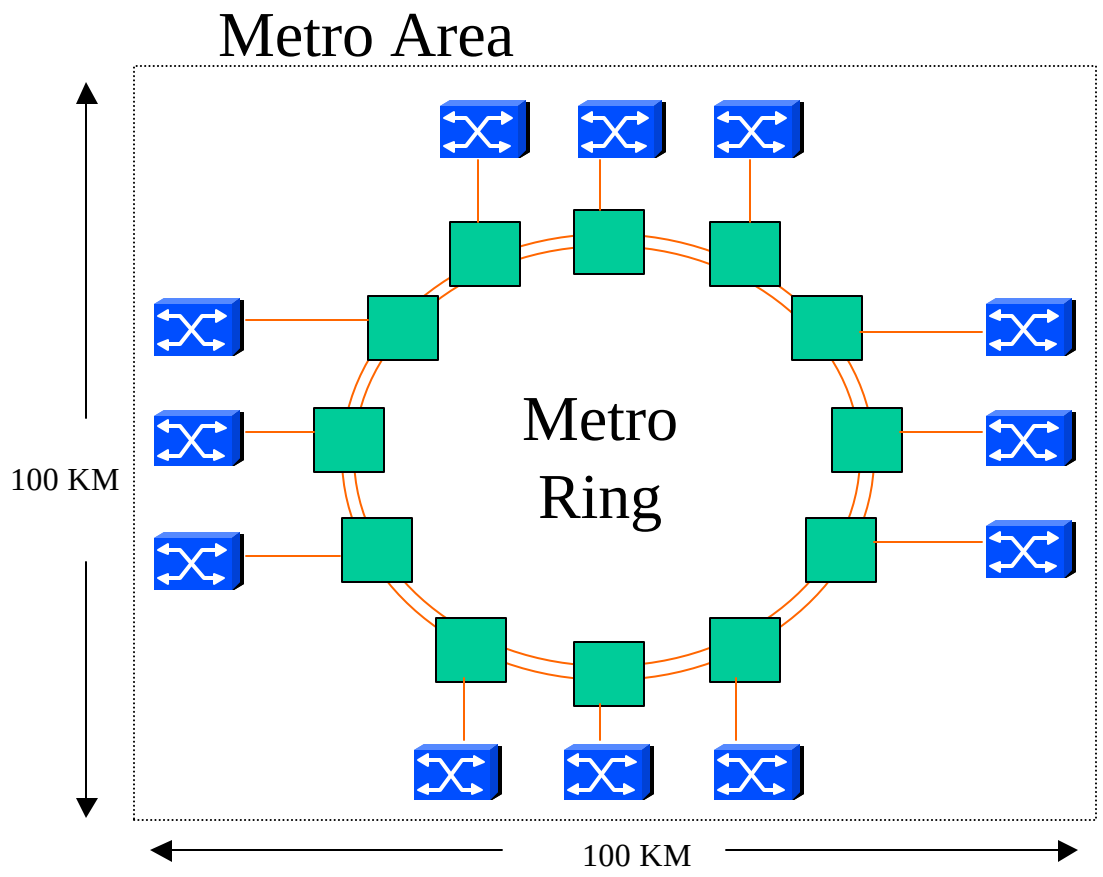

- **A lot of RPR discussions have focused on the objectives that call for rings that scale up to 1000+ KM circumference and 128+ nodes on a ring**

- **SBC thinks of RPR as a piece in an end-to-end solution.**

Typical SBC Metro TDM Transport Network

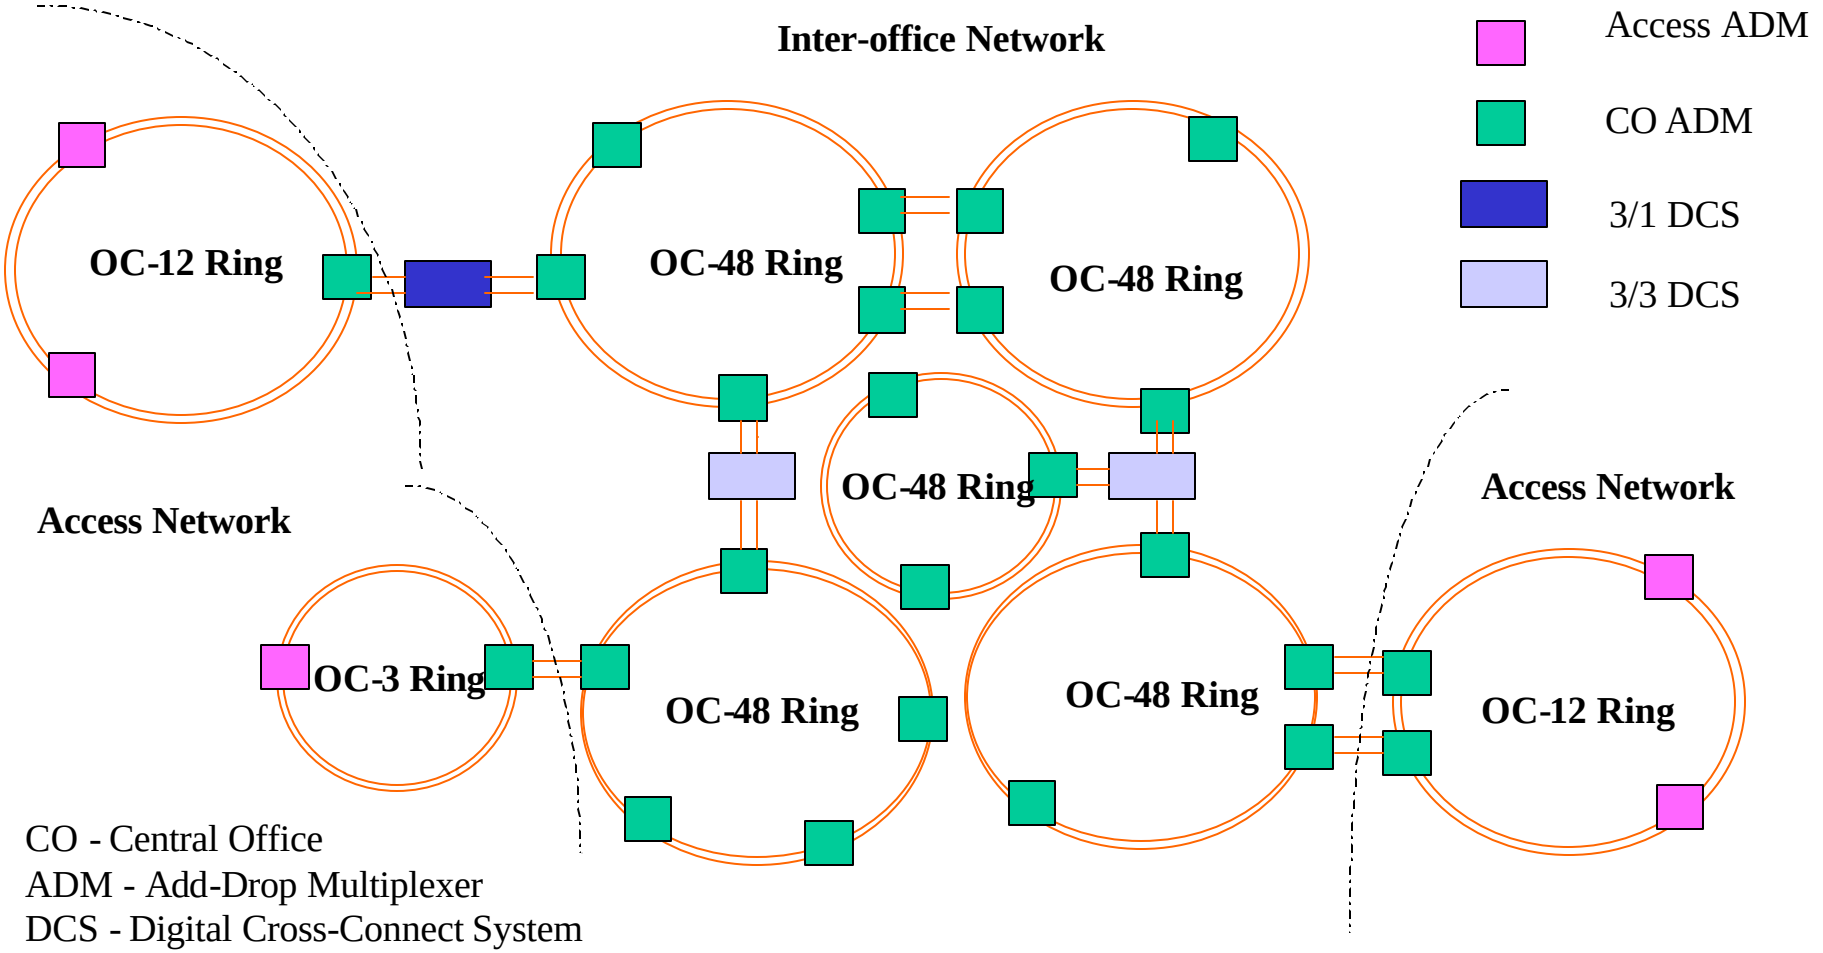

CO - Central Office
ADM - Add-Drop Multiplexer
DCS - Digital Cross-Connect System

SBC transport network

- Access rings are mostly OC-12 and OC-3, some OC-48.
- Inter-office rings are mostly OC-48, some OC-12, few OC-192.
- Access ring - 2 to 5 nodes per ring
- Inter-office ring - 4 -7 nodes per ring
- Average distance between Central Offices - 9 miles.

SBC's View on using RPR for Ethernet Transport

- Main benefits in the Access
  - Statistical multiplexing of low speed data traffic
  - Fair allocation of bandwidth to the different nodes
  - Increased bandwidth efficiencies by providing protection at layer 2
  - Reduce the number of GbE ports consumed at the central office

RPR in the Inter-office network?

- Gigabit Ethernet ports are expected to be near full capacity in the inter-office (IOF) network.
- Less need seen for statistical multiplexing or layer 2 protection.
- Expect to see more point-to-point gigabit Ethernet over fiber or WDM.

Note:

Some access applications may need full GbE transport and point-to-point “dark” fiber and CWDM could find applications there as well

Potential RPR application in the SBC network

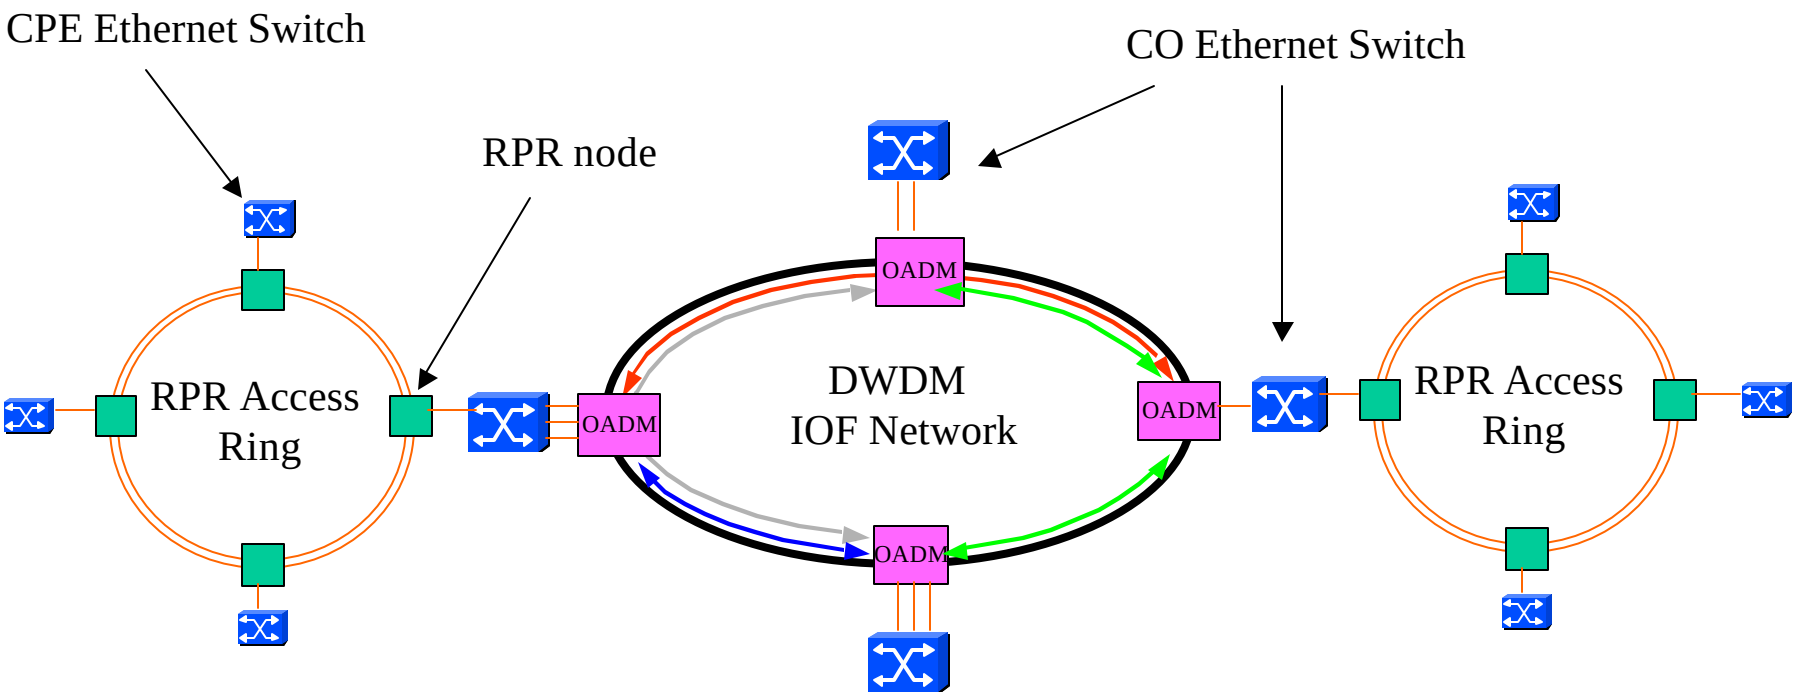

Are we done?

Need a Carrier Class end-to-end Ethernet/RPR Network

- How do we prevent looping of broadcast traffic?
  - Can use a global spanning tree in the network end-to-end and transparently pass BPDUs through RPR nodes.
  - Slow error restoration in non-RPR portions of the network
  - Can use rapid spanning tree algorithm (802.1w) to speed up restoration

- Drawbacks with the single spanning tree approach
  - Inefficient utilization of network resources
  - error restoration speed mismatch in the RPR and non-RPR portions of the network (even with the rapid spanning tree).

Towards Carrier Class Ethernet/RPR Networks

- Ethernet has a flat addressing mechanism and hence the size of the filtering database will affect the size of deployable networks
  - Encapsulated bridging in an RPR network (May 2001, July 2001 802.17 presentations) helps the filtering database to scale better as Ethernet switches connected to intermediate RPR nodes would not have to learn MAC addresses that do not originate or end in those switches.

- 802.1Q has a limit of only 4096 VLANs. This again limits the number of customers that can be supported in an Ethernet network. Stacked VLANs and VLAN tag translation capabilities implemented by several vendors alleviate this problem to some extent.
  - Support for customer traffic separation in RPR (to separate a large number of customers) and a mechanism in Ethernet switches, that bridge RPR rings, to transparently pass this RPR traffic could be a potential solution. Some discussions have taken place in the RPR group on this subject.

Towards Carrier Class Ethernet/RPR Networks (Contd.)

- Per VLAN spanning tree - 802.1s
  - The per VLAN spanning tree could theoretically help in distributing the load evenly through out the network. Need a specialized tool, that is not yet available, to study the resources available and lay out the spanning trees to most efficiently utilize network resources.

- Lack of OAM&P information in Ethernet packets - Work is being performed in RPR to have OAM&P information in RPR packets.

Towards Carrier-Class Ethernet/RPR networks

- Lack of a carrier-class end-to-end network management system that can manage Ethernet switches (as well as RPR nodes)
- Lack of a mechanism to guarantee bandwidth end-to-end in a network made up of Ethernet switches (and RPR nodes)
  - In an RPR ring, a mechanism is needed that ensures that the sum of all guaranteed bandwidths is less than half the ring bandwidth

Ethernet over MPLS has promise to address a lot of the problems mentioned above

EoMPLS Benefits

- EoMPLS can offer end-to-end network bandwidth and other QoS guarantees
- An EoMPLS network need not run the spanning tree protocol and hence could eliminate all the problems associated with spanning trees
- EoMPLS network can provide error restoration in 50 ms.
- EoMPLS network will not have the 4096 VLAN scalability problem
- EoMPLS alleviates the flat address learning problem to some extent (analogous to encapsulated bridging - intermediate MPLS nodes need not participate in MAC learning)

There is some overlap in the capabilities of EoMPLS and RPR

Are there benefits for them to exist together?

Benefits of using RPR along with MPLS

- RPR can provide fair allocation of ring bandwidth in the case where several sources of the same traffic priority class (non-guaranteed) are vying for the ring bandwidth
- RPR can provide locally fast restoration before MPLS based restoration mechanism is invoked
- RPR nodes can transparently pass MPLS packets for nodes that don't have a need to be MPLS enabled or to reduce MPLS processing in MPLS enabled nodes (in the event MPLS processing turns out to be expensive).
- Reduced operations cost for non-MPLS traffic due to the plug-and-play nature of RPR.
- Providing OAM&P capability to Ethernet transport.