ROUTERLESS AGGREGATION:
CONVERGING DATA AND TDM NETWORKS

Benoit Legault
Dir. Technology Advisory Board
ADC

Abstract

MSOs interested in remaining competitive and deploying new services are faced with two network architecture options. Traditional routing-intensive metro networks offer clear benefits for data services, but fall short on the ability to converge data and TDM networks. Innovative MSOs interested in offering voice and other delay-sensitive services should deploy a network architecture with fewer router: "routerless aggregation".

INTRODUCTION

The current economy dictates that MSOs increase earnings and free cash flow. This can be done by increasing revenue through the introduction of new services and increasing the level of profitability of current services through reducing CapEx and OpEx.

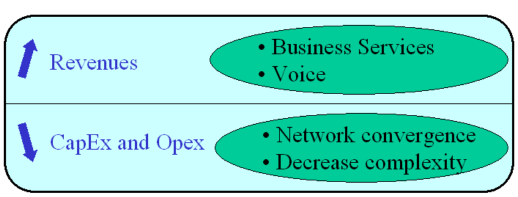

Fig 1 - MSO Opportunity

Increasing revenues

MSOs can improve revenues through increasing the market share of products and services with a potential for growth (i.e. high-speed Internet access), protecting market share of mature products and services (i.e. video), and expanding the product portfolio. Business services, historically an ILEC monopoly, are one possible area of growth or portfolio expansion, which have not been aggressively exploited by MSOs.

Reducing CapEx and Opex

An MSO's network represents the largest portion of its investment. It is composed of three different segments: the access network, which starts at the hub or headend and terminates in the subscriber's home, the metro network, which interconnects the hubs of a metropolitan or regional area, and the backbone network which interconnects the metro networks.

Today, MSOs own and operate three different metro networks: video, data and TDM. Therefore, the opportunity exists for operators to significantly reduce their network CapEx and OpEx through converging the data and TDM networks.

The Requirements

Current metro data network architectures can handle the requirements of Internet Access, both residential and commercial. To allow convergence in the future, metro networks must be able to support the requirements necessary to deliver the following services:

- Residential voice (individual POTS lines)
- Business private networking
- Business voice (T1s and T3s for PBX applications)

Voice Services

Voice services, both for residential and commercial customers, require high-availability networks with low latency, delay and jitter. The main attributes of high-availability networks are:

- No single point of failure, both in the signaling path and in the call path
- Fast network convergence to avoid dropping calls upon failures
- Transparent software upgrades to avoid downtime associated with network maintenance

Commercial private networking services require that the MSO be capable of delivering Layer 2 "pipes" across the metro network. This requirement is driven by two factors. First, many business customers still carry legacy protocols, such as IPX, SNA, LAT, DECnet, Appletalk and others, on their networks. In addition, businesses do not expect or desire that their address plan be impacted by the carrier's network.

Current Metro Network Architectures

There are two main problems with the traditional routed metro network architecture. First, it is composed of enterprise-class elements, which do not support high-availability. Second, it is exclusively composed of routers at every hop, which complicates the offering of private networking services.

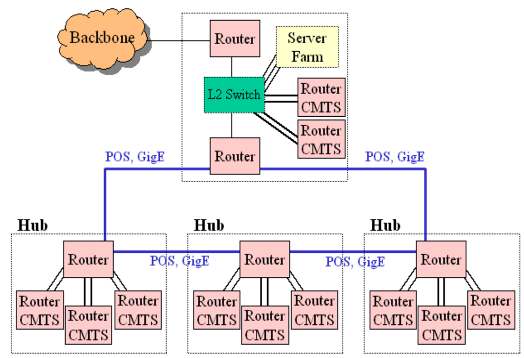

Fig 2 - Current Metro Network Architecture

Voice services, whether individual POTS lines or T1s and T3s for PBX applications, require high-availability networks that guarantee service availability and sub 50ms automatic switchover. Since the enterprise-class routers currently used in traditional metro network architectures do not offer carrier class availability, the classical method of increasing the availability of these networks consists of installing redundant edge routers at every hop. This architecture has two major drawbacks when offering telephony services:

- Since high-availability is provided at the IP layer through routing protocols, TDM services can only be offered through circuit emulation over IP.
- Since OSPF's convergence time is far above the traditional 50ms recovery time of voice networks, circuit emulation requires that operators implement MPLS-TE for its fast recovery features.

In addition, the traditional metro network routed architecture lacks native support for business private networking services, which can only be supported through the introduction of new networking protocols. The only solution is to emulate Layer 2 over Layer 3, which requires the configuration and management of a number of protocols (L2-MPLS/VPNs (Martini), MPLS-TE, OSPF-TE extensions, CR-LDP, etc.) across the metro. The result is an extremely complex and costly network to own and operate.

<u>Routerless Aggregation</u>

As stated previously, routed metro networks pose two main problems: the lack of native support for TDM services because of the absence of sub-IP layer path protection and the complexities of offering private networking services.

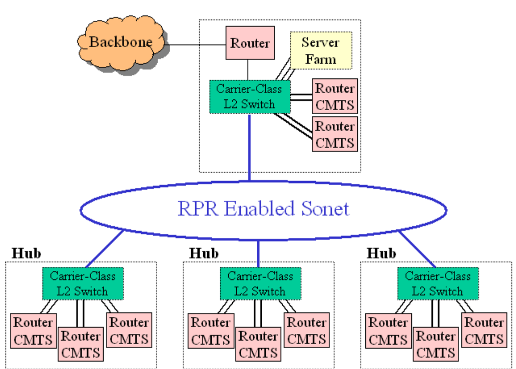

Fig 3 - Routerless Aggregation Architecture

The *routerless aggregation metro network architecture* solves both these problems. This architecture uses carrier-class Layer 2 switches as the aggregation element in each hub and pushes the routing function to the edges of the metro network. It essentially applies the networking principle of *routing at the edge and switching at the core* through the use of Layer 2 switching as opposed to the more complex MPLS switching. In a typical regional network, this architecture is composed of Layer 3 CMTSs at the subscriber edge (located in the hub), and edge routers at the provider or MSO backbone edge (located in the regional head-end or regional data center). All other elements in between are Layer 2 switches.

This architecture provides a number of benefits and, as we will discuss, a number of shortcomings that must be addressed. Lets start with the benefits.

<u>The Benefits</u>

Faster – Cheaper

Routerless aggregation is far more cost effective than enterprise-class Layer 3 aggregation for mainly two reasons. First, carrier-class switches only require half the number of elements and interfaces, compared to the traditional network architecture, by eliminating the need to duplicate the elements. A number of manufacturers now offer Layer 2 switches with built-in redundancy at all levels, and support for transparent software upgrades. Note that, when combined, these features allow a single Layer 2 switch to provide equal or better overall system availability than a pair of enterprise-class routers. Second, on a side-by-side comparison, any particular router port is generally more expensive than its equivalent on a Layer 2 device.

Simple support for private networks

VLANs have been used for years to support private networking services. Most PTTs and ILECs around the world have been using this technology for well over a decade, and continue to do so with much success. Through routerless aggregation, MSOs operate metro networks that can support Layer 2 point-to-point or multipoint-to-multipoint "pipes" from anywhere to anywhere within the region. The private networks are simple to configure and manage, and allow the MSO to support any Layer 3 protocol without getting involved with the subscribers' Layer 3 address plan or even being aware of the transported Layer 3 protocols.

PHY-layer protection

The carrier-class nature of these switches, when combined with interfaces that provide native support for TDM and packet-based services, allow the routerless aggregation architecture to provide native TDM services from anywhere to anywhere in the metro. Sonet and ATM are good examples of such interfaces. This architecture truly allows the convergence of the data and TDM networks in the metro, further reducing CapEx and OpEx.

The "Gotchas" And The "Fix-Its"

Avoiding spanning tree

Layer 2 networks usually rely on spanning tree to manage redundant paths in the network. Spanning tree provides slower convergence and is far less intelligent than routing to control and manage redundant paths in a network. Spanning tree is known to cause outages through broadcast storms, constant flapping to administrative mode, and other problems that derive from its basic operation. It is, in most cases, the main reason why many network architects have previously dismissed Layer 2 networks as viable network architectures. Most of the reasons why spanning tree was considered inadequate still exist, and therefore, the author shares the view that if spanning-tree cannot be avoided in a routerless aggregation architecture, the architecture should be considered incomplete and problematic.

On the other hand RPR, which is a new Layer 2 protocol that creates fault tolerant rings as an overlay of point-to-point GigE or Sonet links, allows the use of Layer 2 devices without resorting to spanning tree to manage redundant paths. Elements on an RPR ring are provided with a single Layer 2 path to all other elements on the ring, such that spanning-tree is never required to manage the ring's redundant paths. The RPR MAC layer handles interface and link failures transparently, such that changes to the network's links' status are never apparent to any element's Layer 2 (or Layer 3, for that matter) forwarding table.

VLAN scalability limits

The maximum number of supported VLANs on any given interface, per the standard Layer 2 header, is 4096. In some cases, this limit poses a scalability problem for MSOs, especially in medium to large size regions.

The routerless aggregation network architecture proposes to solve this problem by creating multiple Layer 2 RPR aggregation rings in the metro and to joint these rings through the use of edge routers implementing Layer 2 MPLS VPNs (Martini). This approach addresses the scalability issues of VLANs without introducing the complexities associated with implementing MPLS throughout the metro network. The result is a very scalable Layer 2 VPN solution that is manageable and has a level of complexity that grows with the services' level of success.

Impact of Layer 2 aggregation on OSPF

The routerless aggregation network architecture essentially flattens the metro network from a routing perspective. Flattening the metro has impacts on OSPF, or any other routing protocol. The most significant impact is that it increases the number of OSPF adjacencies maintained by each router in the network. If not factored into the design, an oversized growth in OSPF adjacencies will cause problems in the operation of the network. Routers will suffer from performance problems, convergence will be slow, and network stability will be negatively affected. Note that the maximum number of adjacencies supported by any given router is vendor-specific.

The scalability solution for VLANs also solves the OSPF scalability issues associated with routerless aggregation. In large metro networks, where routerless aggregation would cause too many OSPF adjacencies if a single Layer 2 network was created, the network should be partitioned into multiple sub-networks through the use of edge routers. This goal can be reached either by a physical implementation, or by creating the partitions through a logical overlay. An example of a physical implementation is to create two rings in the metro and place an edge router at their intersection point, physically separating the two MAC domains. A logical overlay uses VLANs to create two separate MAC domains over a single physical network through the use of an edge router that can be attached anywhere in the ring.

<u>RPR-Enabled Sonet</u>

This final section describes how routerless aggregation is a network architecture that can truly enable the convergence of TDM and data networks in the metro. As described earlier, RPR plays an important role in routerless aggregation metro network architectures. Also, Sonet and ATM are essential to the native coexistence of TDM and packet-based services over a single infrastructure. Given the fact that RPR can either use GigE or Sonet as its underlying physical network layer, RPR-enabled Sonet networks are, if implemented properly, the true enabler of converged networks.

A proper implementation of RPR-enabled Sonet allows the MSO to carve STS-1s out of the Sonet bandwidth allocated to the RPR interface to natively support T1 and T3 services across the metro. Products that support this feature can provide pure circuit switched T1 and T3 interfaces for telephony applications, along with GigE and 10/100BT interfaces for data applications on the subscriber side, and use a single uplink on the network side to carry all services across the metro.

CONCLUSION

The routerless aggregation metro network architecture provides a number of key benefits over current metro networks architectures, which allow

MSOs to *a la fois* increase revenues and reduce CapEx and OpEx.

The architecture benefits MSOs by enabling:

- Data and TDM network convergence, with native support for TDM services (as opposed to circuit emulation)
- Private networking services in a simpler and just as scalable manner as MPLS

The architecture avoids the age-old issues associated with link redundancy in Layer 2 networks through a new protocol: RPR.