Designing a Reliable IPTV Network

Meeting IPTV's quality of service constraints (such as low latency and loss) requires designing the right combination of underlying IP-transport, restoration, and video and packet recovery methods. Carriers use link-based fast reroute (FRR) as the primary transport restoration method to achieve this goal. Although we can carefully tune the link weights in the IP routing protocol to avoid traffic overlap from FRR during single link failures, multiple failures can still cause path overlap in long-distance networks. By having FRR, Interior Gateway Protocol, and multicast protocols work in harmony and with appropriate link weight assignments, this approach can help minimize path overlap during multiple failures.

Distribution of real-time multimedia over an IP backbone has been gaining momentum with content and service providers.[1,2] (For example, see www22.verizon.com/FiosForHome/Channels/fios/FiosTV_comingsoon.aspx and www.iptvnews.net.) The challenge is to distribute real-time, linear broadcast TV while ensuring that users experience no more than a few seconds' delay. Loss-recovery mechanisms have a limited capability to recover from burst packet losses. To recover from short burst losses, providers use a combination of higher-layer methods, such as player loss-concealment algorithms, retransmissions, and packet-level redundancy mechanisms. If we combine these higher-layer loss-recovery mechanisms and protocols with the capability to rapidly restore the IP-transport network after a network failure, then we can provide end users with the necessary quality of service (QoS).

Clearly, delivering the needed QoS requires a carefully designed architecture that melds standard IPTV architectures, video and packet recovery mechanisms and protocols, and underlying network design and restoration methods. With this article, we describe such an implementation based on our hands-on experience with real networks. Specifically, we use multicast[3–5] to achieve an efficient, cost-effective network and use link-based fast reroute (FRR) in the underlying IP transport network.[6–8] The vast majority of net-

Robert Doverspike, Guangzhi Li, Kostas N. Oikonomou, K.K. Ramakrishnan, Rakesh K. Sinha, Dongmei Wang, and Chris Chase
AT&T Labs

1089-7801/09/$25.00 © 2009 IEEE

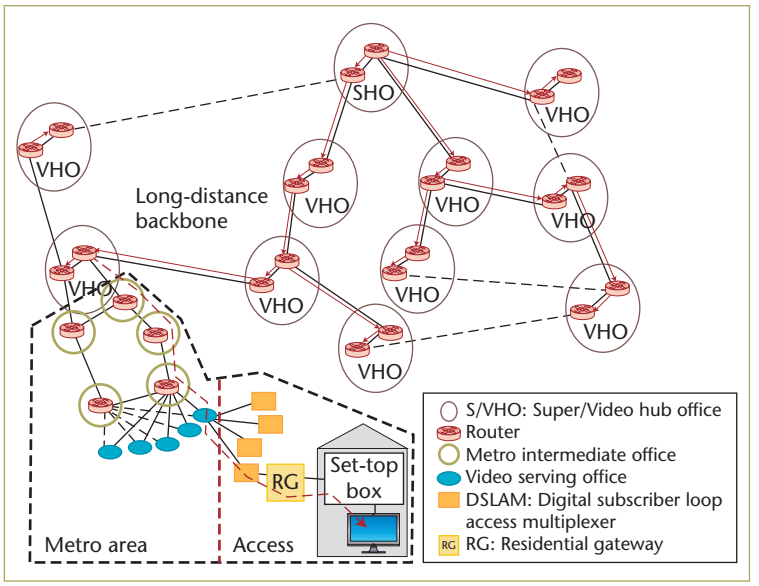

Figure 1. IPTV architecture. This architectures applies to long-distance backbone, metro, and access segments that deliver IPTV service in the continental US. The red arrows show the directed, unidirectional links of the example long-distance multicast tree.

work (physical layer) failures affect only a single link, which can be restored using FRR within roughly 50 milliseconds. The higher-layer recovery methods can handle most losses of 50 milliseconds or less without significant viewer impairment. However, FRR must be coordinated carefully with the placement of network capacity to avoid traffic overlap during network failures, which can cause congestion and packet loss. We can avoid traffic overlap during any single IP link failure by tuning the link weights of the Interior Gateway Protocol (IGP).[9]

Our remaining task then is to identify methods to mitigate the effect of multiple failures such that the resulting downtime doesn't exceed the target network unavailability objectives. By using FRR, IGP, and multicast protocols in harmony and with appropriate link-weight assignments, our proposed approach minimizes path overlap during multiple failures.

IPTV Network Architecture

To show the general network architecture and design for distribution of broadcast video and other services, Figure 1 gives an example, simplified network topology for segments that deliver IPTV service in the continental US. (Figure 1 describes a specific IPTV architecture, but other variations are possible.) The *super hub office* (SHO) gathers content from the national video content providers, such as TV and cable networks (mostly via satellite today) and distributes it to a large set of receiving locations, called *video hub offices* (VHOs). Each VHO in turn feeds a metropolitan area. IP routers are used to transport the IPTV content in the SHO and VHOs. The combination of SHO and VHO routers plus the links that connect them comprise the long-distance IPTV backbone. The VHO combines the national feeds with local content and other services (which we describe later) and then distributes the content to *video serving offices* (VSOs) via the routers in the metro IPTV backbone. Finally, the routers or Ethernet switches (or hybrids thereof) in the VSOs deliver the content down the feeder plant to the *digital subscriber line access multiplexers* (DSLAMs) located in environmentally controlled vaults or huts — sometimes called *video ready access devices* (VRADs) — arising from lawns or easements or a common space in apartment buildings.

A DSLAM often serves 100 to 200 residential gateways (RGs) that are attached to the outside of a residence. Various techniques let us transport the signal between the DSLAM and RG: typically, very high-bit-rate DSL (VDSL) for copper and broadband passive optical network (BPON)[10] and gigabit PON (GPON)[11] for fiber transport. This metro access segment uses a hub-and-spoke architecture, wherein smaller switches connect to larger ones. To reduce costs, at each stage (going from VHO→VSO→DSLAM→RG) the switch more closely resembles a pure Ethernet switch. For example, the packet switch in the VSO is an Ethernet switch with some additional capabilities found in carrier-grade routers; the switch in the DSLAM closely resembles a pure Ethernet switch.

In the IPTV environment, the VHO is the nexus for serving the bulk of residential entertainment and telecommunications. In addition to broadcast video, this includes video on demand (VoD), voice, broadband Internet service, and other multimedia services, such as music. For example, as the RG delivers video content to the set-top box (STB), voice services are either

- converted to voice over IP (VoIP) and transported as IP flows to the VHO and then to a VoIP gateway or
- multiplexed in analog form over VDSL.

In the latter case, the voice signal is demultiplexed at the DSLAM and connected to preexisting remote terminal (RT) architectures that use traditional time-division multiplexing

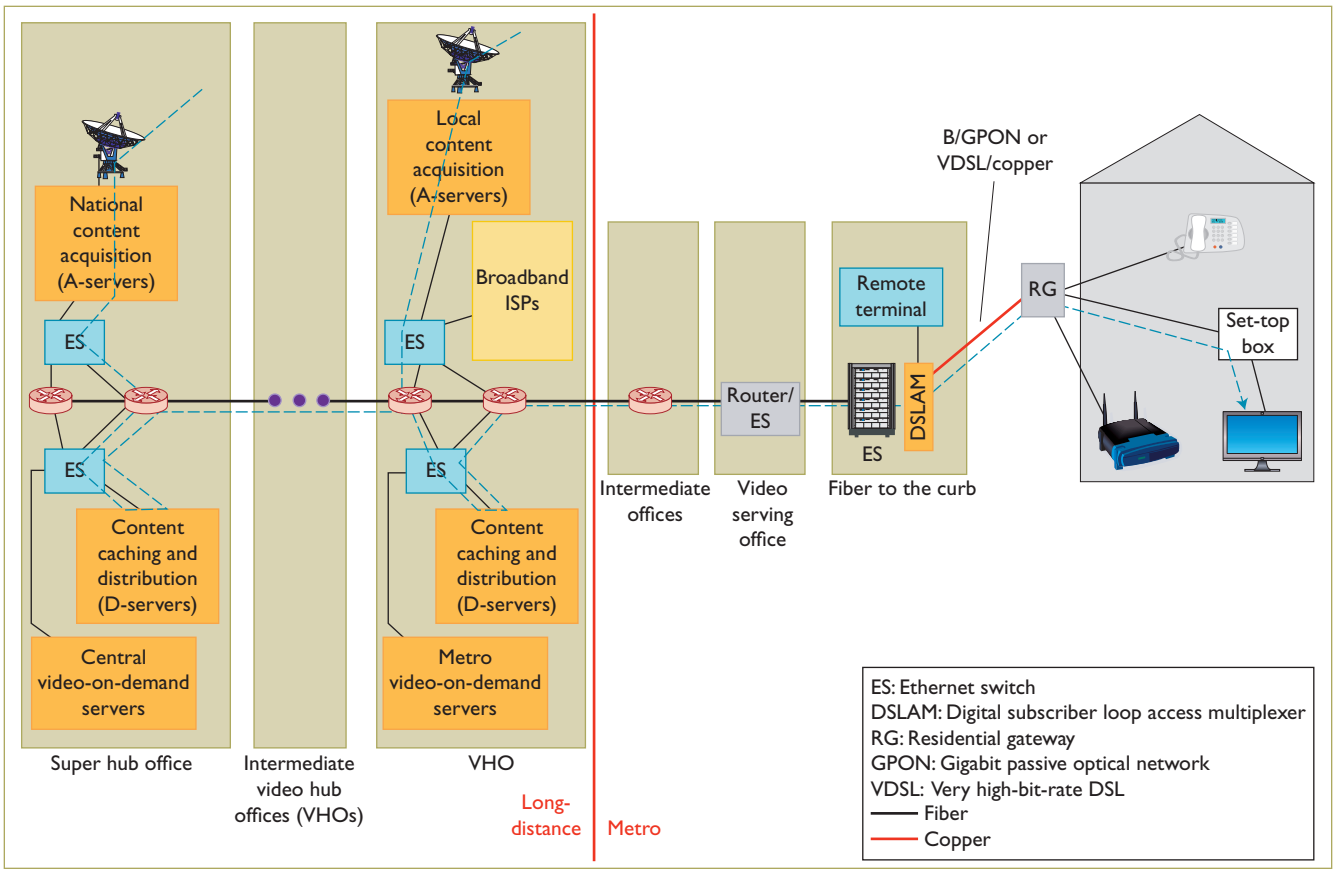

Figure 2. Flow of broadcast video through an IPTV network. The dashed line shows an alternate, simplified view of the broadcast video content's journey.

(TDM) transport technologies, such as DS-0, DS-1 (digital signal level 0/1), and SONET (Synchronous Optical Network).

When we examine Figure 1's long-distance backbone network in more detail, we see that the SHO and each VHO have two routers to provide redundancy in case of router component failure, router hardware maintenance, or software upgrade. The links of the long-distance IPTV backbone are usually 2.5 or 10 gigabits per second (Gbps) Ethernet or SONET. Because we use cost-efficient tree-like topologies, IP multicast provides economic advantages for IPTV service delivery. Note the scarcity of links between SHO routers and their neighbor VHO routers needed to form the broadcast tree (a 1-connected network). The links designated with dashed lines aren't part of the multicast tree; these extra links are alternate paths needed to restore the multicast tree in the case of a network failure or maintenance event.

Because of the less-connected nature of most metro telecommunications networks, the topologies of the metro IPTV backbones resemble rings. The links between the VSO switch and metro intermediate router are usually high speed (for example, 10 Gbps Ethernet), and automatic restoration occurs via the establishment of an alternative path between each VSO switch and I/O router, generally through another intermediate router. Typically, the network's access segment (between the VSO and RG) is only 1-connected.

The dashed line in Figure 2 gives an alternate, simplified view of the broadcast video content's journey, which originates at the SHO and ends at the user's STB. In the SHO and VHO, we see content acquisition devices (A-servers) and content and distribution devices (D-servers). The live-feed video and other multimedia flows are buffered in the D-server for retransmission and other applications. A D-server's average I/O capacity is several hundred megabits per second. The A-servers at a VHO gather local content, such as local TV stations or community programs, to add to the national content gathered at the SHO.

In addition to the broadcast feed, Figure 2 shows VoD servers. Typically, VoD is gathered

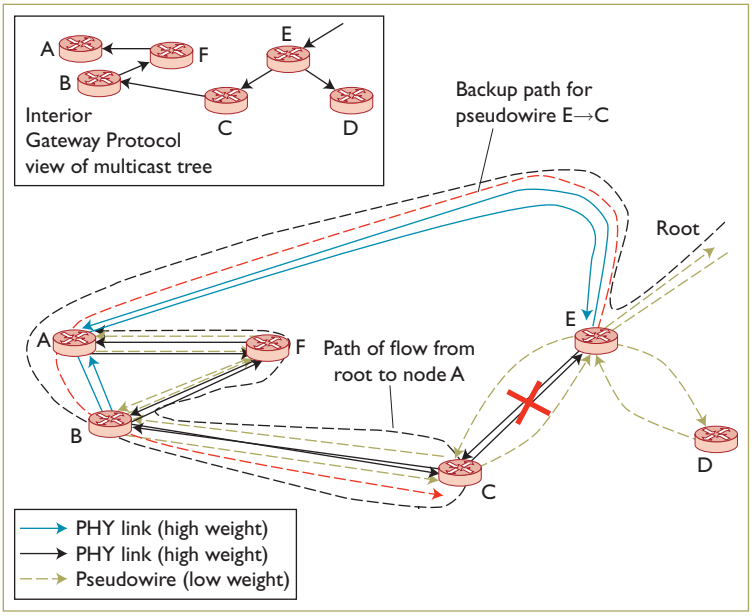

Figure 3. Example of fast reroute (FRR). The red dashed line shows the FRR backup path for the pseudowire E⟶C. The solid line E⟶C represents a single PHY link (a physical-layer link in each direction) over which the pseudowire E⟶C routes as its primary path.

at the SHO or other centralized location and distributed in the background — that is, it isn't usually live — to local VoD servers in the VHO. The servers then interact with the STBs via IP connection-oriented (point-to-point) unicast flows because of the need for individual user control. (Researchers are working on more capacity-efficient methods of VoD distribution.[12]) Servers and routers in the VHO are usually interconnected via Ethernet switches in a hub-and-spoke topology. The router in the VHO splits out all these IP flows to the various video servers and edge routers for broadband Internet service.

Many of the protocols used in IPTV transport are beyond this article's scope, but a few salient observations are relevant to our discussion. The video content uses a differential video compression technology, such as MPEG (www.mpeg.org) or the ITU-T H.264 standard.[13] The video frames are packetized and carried typically over Real-Time Transport Protocol (RTP)[14] over User Datagram Protocol (UDP) over IP.[15] The IPTV architecture uses Protocol Independent Multicast, Source Specific Multicast (PIM-SSM) to multicast the TV content over the IP network.[5] Each channel from the national live feed at the SHO is assigned a unique multicast group. There are typically hundreds of standard-definition (SD; 1.5 to 3 megabits per second) and high-definition (HD; 6 to 10 Mbps) channels plus other multimedia signals, such as music. So, the live feed can be multiple gigabits per second in total bandwidth.

This total will likely grow significantly with the number of HD channels. As the broadcast signals enter the VSO, the IPTV architecture uses the Internet Group Management Protocol (IGMP)[16] instead of PIM-SSM for the STBs to join the multicast group at the data-link layer because the VSO switch doesn't have the full capabilities of a carrier-grade router. In addition, to manage the potential explosion of control messages to the metro routers that could occur from many switches in the hub-and-spoke architecture, an IGMP snooping technique is used in the switches between the STB and the metro intermediate router to filter the IGMP messages and locally replicate the channels to the users that request them. IGMP snooping also enables the link between the DSLAM and VSO switch to have less bandwidth requirements than that between the VSO and I/O router.

Restoration Mechanisms

The network structure and restoration methods in the IPTV network we've described so far are specifically designed to deliver the needed video QoS to end users while providing an economical solution to the network provider. We can recover from relatively infrequent and short bursts of loss using a combination of video and packet recovery mechanisms and protocols, including the Society of Motion Picture and Television Engineers (SMPTE; www.smpte.org/standards) 2022-1 Forward Error Correction (FEC) standard, retransmission approaches based on RTP/RTCP[14] and Reliable UDP (R-UDP),[17] and video player loss-concealment algorithms in conjunction with STB buffering. The desired functionality of R-UDP is primarily retransmission-based packet-loss recovery. Vendors have implemented this aspect of R-UDP without incorporating the other mechanisms specified in RFC 1151,[17] such as congestion-control algorithms.

Single Link Failures

Besides preventing miscellaneous video impairments due to transmission problems, such as server hiccups or last-mile facility degradation, some combination of these methods can recover from any network failure of 50 milliseconds or less. Repairing network failures usually takes far more than 50 milliseconds,

but when combined with link-based FRR, this restoration methodology could meet the stringent requirements needed for video against single link failures.

Figure 3 illustrates how we might implement link-based FRR in an IPTV long-distance backbone by depicting a network segment with four node pairs that have defined virtual links (or pseudowires). For example, node pair E-C has a physical-layer link in each direction (often abbreviated as PHY in IP standards) and a pseudowire in each direction (a total of four directed links) used for FRR restoration. For example, the PHY link might be Gigabit Ethernet or SONET. The red dashed line shows the FRR backup path for the pseudowire E→C. Note that links such as E-A are for restoration and, hence, have no pseudowires defined. Pseudowire E→C routes over a primary path that consists of the single PHY link E→C (see the solid line in Figure 3).

If a failure occurs to a PHY link in the primary path, such as C-E, then the router at node E attempts to switch to the backup path using FRR. The path from the root to node A will switch to the backup path at node E (E-A-B-C). Once it reaches node C, it will continue on its previous (primary) path to node A (C-B-F-A). During the failure, although the path retraces itself between the routers B and C, the multicast traffic doesn't overlap because of the links' unidirectionality. Also, although the IGP view of the topology realizes that the PHY links between E-C have gone "down," because the pseudowire from E→C is still "up" and has the least weight, the shortest path tree remains unchanged. Consequently, the multicast tree remains unchanged. The IGP is unaware of the actual routing over the backup path.

Some router and switch vendors have demonstrated an FRR mechanism to switch to the backup path within 50 ms.[8] Because the primary path and its backup path are nonoverlapping at the physical layer, all pseudowires will restore in 50 ms as a result of any single fiber or wavelength division multiplexing (WDM) lower-layer link failure and thus IGP doesn't see any change in the IP topology if we configure the IGP hold-down timers appropriately. Therefore, the multicast tree will remain unaffected, reducing the impact of the single link failure from tens of seconds without FRR to approximately 50 ms with FRR.

Because IGP is unaware of the FRR backup paths during network failures, traffic overlap could potentially occur. In this case, "traffic overlap" means that the packets of the same multicast flows travel over the same link (in the same direction) two or more times. Because we need to deploy economically efficient networks and an increasing number of SD and HD channels require more bandwidth, high link utilization might occur even during the normal (no-failure) state. This could cause congestion to occur during traffic overlap. For real-time video services, loss from even mild congestion often has the same effect on the customer's perception of service as an unprotected link failure. Therefore, to exploit the cost reduction made possible by multicast, we seek to avoid this. (See related literature for a scheme that prevents traffic overlap due to a single link failure by constructing the multicast tree to be nonoverlapping from each link's backup path.[18])

Multiple Link Failures

Combining small-loss recovery mechanisms, FRR, and intelligent selection of nonoverlapping paths[18] would solve the problem of providing high network availability to achieve stringent video QoS if only single link failures occurred in the network. However, this doesn't mitigate the problem of multiple link failures. A given pseudowire's backup path is typically precalculated, and there's no real-time (dynamic) mechanism for changing the backup path due to different combinations of multiple link failures. For example, if another link failure occurs during the outage period of the first failure, then the backup path of a pseudowire might fail. In this case, the IGP routing protocol would reconverge and generate a new multicast tree.

An alternative approach builds on the FRR mechanism but limits its use to a short period. After a single failure occurs and a pseudowire's primary path fails, the traffic is rapidly switched over to the backup path. However, soon afterward, the router sets the virtual link weight to a high value and thus triggers the IGP reconvergence process — this is colloquially called "costing out" the link. Once IGP routing converges, a new PIM tree is rebuilt automatically. This ensures rapid restoration from single link failures while allowing the multicast tree to dynamically adapt to any additional failures that

Table 1. Reliability and protocol parameters for the nperf model.

Parameter name	Parameter value
Link mean time between failure (MTBF) (including the two router interfaces and a fiber)	800 hours
Router MTBF	50,000 hours
All mean times to repair (MTTRs)	4 hours
Fast reroute (FRR) convergence time	50 ms
Interior Gateway Protocol (IGP) convergence time	10 sec.
Protocol Independent Multicast (PIM) convergence time	200 ms

might occur during a link outage. It's only during this short, transient period between when FRR starts and IGP reconvergence finishes that another failure could expose the network to a path overlapping on the same link.

The potential downside of this approach is that it incurs two more network reconvergence processes — that is, the period right after FRR has occurred and then again when the failure is repaired (network normalization). Thus, if it isn't carefully executed to handle potential small losses (called hits) due to reconvergence, this alternative approach can cause small video interruptions for the more frequent single failures.

To prevent this, a key component of the method is the *make-before-break change* of the multicast tree — that is, the requirement to switch traffic from the old multicast tree to the new multicast tree with minimal traffic loss. The details of this technique are too complex to describe here, but we can summarize its steps:

1. Use the Doverspike and colleagues algorithm to set link weights.[18]
2. For each one-hop, unidirectional IP-layer link, set the primary path equal to that same single-hop link. Precompute an FRR backup path for each IP-layer link. The backup path shouldn't overlap with the multicast traffic flow over other primary paths. The link weights generated in step 1 guarantee that such a backup path exists.
3. When a primary path failure occurs, invoke the FRR mechanism to reroute the traffic to the backup path, provided the backup is operational.
4. Send out an IGP link state advertisement (LSA) with a high weight for the associated IP-layer link. During the IGP reconvergence time, forward traffic along the backup path.
5. After IGP reconvergence concludes and its shortest path tree is constructed, PIM rebuilds its multicast tree with join and prune requests, but in such a way that the new tree's branches are created before the original tree's branches are pruned and that downstream nodes aren't joined to the new tree until traffic flows to their parent nodes. This method incurs a virtually "hitless" multicast reconvergence process.
6. After the failure is repaired, execute the IGP reconvergence and PIM tree rebuilding process, as in step 5.

The make-before-break method is similar to the steps other researchers have used to switch from a shared tree to a source-specific tree with PIM.[6] In this way, each receiver is guaranteed to continuously receive packets, even when switching from one tree to the other.

Network Performability Evaluation Studies

Performance plus reliability (*performability*) evaluation analyzes a network's performance in the presence of failures. The nperf performability analysis tool represents the network under study using a multilevel model.[19] In particular, it models all failure mechanisms in its "component" level, which is closest to the real network's physical layer. For example, the failure of any line card components in routers, or optical amplifier and fiber bundle components, affects the graph's edges. nperf generates network failure scenarios systematically by assigning a working or failed mode to each component in the model. We determine the probability of a component being in each of its modes from the mean time between failure (MTBFs) and mean times to repair (MTTRs) specified at the model's reliability level. All components of a certain type, such as all line cards of a particular type, are assigned the same MTBF and MTTR. Because we are assuming the components are independent, we can determine the probability of a network state simply by multiplying together the appropriate component mode probabilities.

nperf generates these network states in order of decreasing probability. Because the total number of potential states is intractable (exponential in size), the tool uses bounding mechanisms to determine when an objective probabilistic coverage is reached (say, a total probability of 0.9999) and thereby stop examin-

ing further failure states. In each network state, nperf calculates the values of traffic lost because of no path and congestion. Both of these measures take into account various protocol-related timing parameters (see Table 1).

We used nperf to evaluate the performability of a hypothetical US backbone network with 28 VHOs and 45 links. As we already mentioned, each location contains two backbone routers for redundancy, and we consider a VHO reachable as long as at least one of its backbone routers is reachable from the source. (See related work for additional details on the network, analysis, and algorithm for weight setting.[19])

With these parameters, we used nperf to generate the most probable network states until the total probability of the evaluated state space reached the objective tolerance of 0.99999. We calculated the traffic impact due to no path, congestion, and routing protocol convergence, as well as the number of events lasting for more than 1 second per year per VHO, including all failure events that cause traffic loss due to no path, congestion, and IGP convergence.

We then evaluated three methods for dealing with failures:

- Use only IGP but not FRR.
- Apply FRR whenever backup paths are available.
- Use FRR, but only over the interval until the IGP/PIM reconvergence process completes.

For each VHO, Figure 4 shows the total traffic impact in minutes per year and the total number of service-impacting events per year. From these results, we observed the following:

- Method 1 would result in numerous service-impacting events per year for all nodes because all failures require IGP routing convergence to restore connectivity.
- Method 2 reduces the total number of service-impacting events per year significantly, but the total traffic impact from congestion due to multiple failures is still large.
- Method 3 leads to both the smallest number of service-impacting events and the least total traffic impact for all VHOs.

Our current work focuses on a detailed investigation of the protocols necessary to implement method 3.

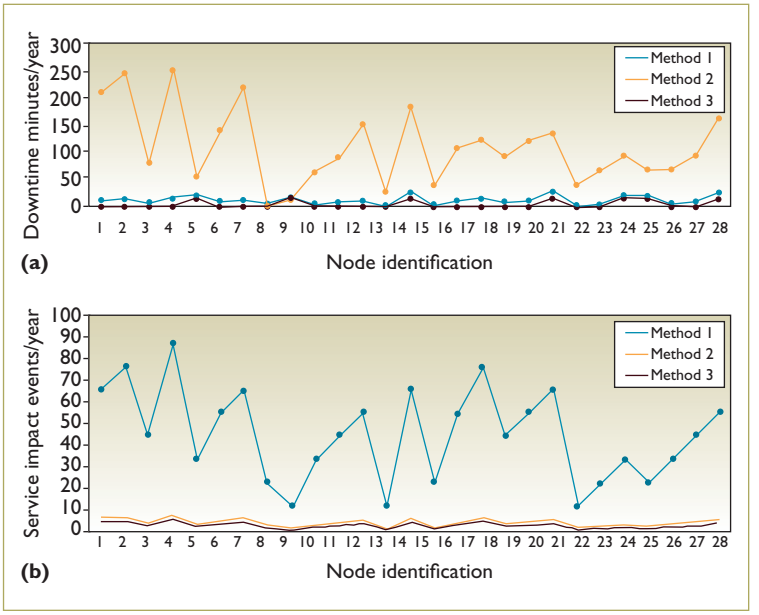

Figure 4. Comparison of total network impact. The graphs show the (a) traffic impact in minutes per year and (b) service impact in events per year.

To deliver the needed QoS for distribution of broadcast TV over an IPTV network, we must carefully design an architecture that melds standard IPTV architectures, video/packet recovery mechanisms and protocols, and underlying network design and restoration methods. We've described such an implementation based on our hands-on experience with real networks. Future work will concentrate on implementing these ideas in equipment suppliers, standards, and actual IPTV networks. Furthermore, we're exploring approaches for automatic generation of optimal backup routes and their real-time implementation in IPTV backbones to enable more dynamic adaptation of routing methods to changing network conditions.

References

1. AT&T, "Customers Love AT&T U-Verse TV," press release, Jan. 2009; www.sbc.com/gen/press-room?pid=5838.
2. AT&T, "IP Video Network Architecture," press release, Jan. 2009; www.sbc.com/Common/files/pdf/IPvideo network.pdf.
3. A. Adams et al., *Protocol Independent Multicast – Dense Mode (PIM-DM): Protocol Specification* (*Revised*), IETF RFC 3973, Jan. 2005; www.ietf.org/rfc/rfc3973.txt.
4. D. Estrin et al., *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*, IETF RFC 2362, June 1998; www.ietf.org/rfc/rfc2362.txt.
5. S. Bhattacharyya, ed., *An Overview of Source-Specific*

Multicast (SSM), IETF RFC 3569, July 2003; www.ietf.org/rfc/rfc3569.txt.
6. E. Rosen et al., *Multiprotocol Label Switching Architecture*, IETF RFC 3031, Jan. 2001; www.ietf.org/rfc/rfc3031.txt.
7. P. Pan, G. Swallow, and A. Atlas, eds., *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, IETF RFC 4090, May 2005; www.ietf.org/rfc/rfc4090.txt.
8. Cisco, "MPLS Traffic Engineering Fast Reroute: Link Protection," www.cisco.com/en/US/docs/ios/12_0st/12_0st10/feature/guide/fastrout.html.
9. J.T. Moy, *OSPF Anatomy of an Internet Routing Protocol*, Addison-Wesley, 2000.
10. *Broadband Optical Access Systems Based on Passive Optical Networks* (*PON*), ITU-T G.983, Int'l Telecommunications Union; www.itu.int/rec/T-REC-G.983.1/en.
11. *Gigabit-Capable Passive Optical Networks* (*GPON*): *General Characteristics*, ITU-T G.984.1, Int'l Telecommunications Union; www.itu.int/rec/T-REC-G.984.1/en.
12. D. Banodkar et al., "Multicast Instant Channel Change in IPTV Systems," *Proc. 3rd Int'l Conf. Communications Systems Software and Middleware and Workshops* (COMSWARE 08), IEEE CS Press, 2008, pp. 370–379.
13. T. Gary et al., "Overview of the H.264/AVC Video Coding Standard," *IEEE Trans. Circuits And Systems for Video Technology*, vol. 13, no. 7, 2003, pp. 560–576.
14. H. Schulzrinne et al., *RTP: A Transport Protocol for Real-Time Application*, IETF RFC 3550, July 2003; www.ietf.org/rfc/rfc3550.txt.
15. J. Postel, *User Datagram Protocol*, IETF RFC 768, Aug. 1980.
16. Juniper Networks, "Introduction to IGMP for IPTV Networks: Understanding IGMP Processing in the Broadband Access Network," Oct. 2007, www.juniper.net/solutions/literature/white_papers/200188.pdf.
17. C. Partridge and R. Hinden, *Version 2 of the Reliable Data Protocol (RDP)*, IETF RFC 1151, Apr. 1990.
18. R. Doverspike et al., "IP Backbone Design for Multimedia Distribution: Architecture and Performance," *Proc. 26th Int'l Conf. Computer Communications* (Infocom 07), IEEE CS Press, 2007, pp. 1523–1531.
19. K. Oikonomou, R. Sinha, and R. Doverspike, "Multi-Layer Network Performance and Reliability Analysis," to appear in *Int'l J. Interdisciplinary Telecommunications and Networking*.

Robert Doverspike is executive director of Network Evolution Research at AT&T Labs Research. His research interests include optimization and restoration in multilayered transmission and switching networks, packet transport in metro and long-distance networks, and advanced interactions of transport and IP network architectures. Doverspike has a PhD in mathematics from Rensselaer Polytechnic Institute. He's a senior member of the IEEE. Contact him at rdd@research.att.com.

Guangzhi Li is a principal member of the technical staff at AT&T Labs Research. His research interests include IP-based control plane for optical networks, restoration and protection schemes and algorithms, network simulation and performance evaluation, and network-related applications. Li has a PhD in computer science from the College of William and Mary. He's a member of the IEEE. Contact him at gli@research.att.com.

Kostas N. Oikonomou is a principal member of the technical staff at AT&T Labs Research. His research interests include probabilistic and combinatorial modeling and analysis, with an emphasis on networks. Oikonomou has a PhD in electrical engineering from the University of Minnesota. He's a member of the IEEE. Contact him at ko@research.att.com.

K.K. Ramakrishnan is a distinguished member of the technical staff at AT&T Labs Research. His research interests include networking and communications, including congestion control, multimedia distribution, content dissemination, and problems associated with large-scale distributed systems. Ramakrishnan has a PhD in computer science from the University of Maryland. He's a fellow of the IEEE. Contact him at kkrama@research.att.com.

Rakesh K. Sinha is a lead member of the technical staff at AT&T Labs Research. His research interests include networking and algorithms, with an emphasis on restoration schemes, network design, and network protocols. Sinha has a PhD in computer science from the University of Washington. He's a member of the IEEE. Contact him at sinha@research.att.com.

Dongmei Wang is a senior member of the technical staff at AT&T Labs Research. Her research interests include network-related research topics, from optical layer to application, architectures to protocols, algorithms to simulation, and provisioning to restoration. Wang has a PhD in physics from the College of William and Mary. She's a member of the IEEE. Contact her at mei@research.att.com.

Chris Chase is a lead member of the technical staff in Network Systems Engineering at AT&T Labs. His research interests include large private enterprise networks, network and protocol performance monitoring, and reliable large-scale content delivery. Chase has a PhD in electrical engineering from Princeton University. He's a member of the IEEE. Contact him at chase@labs.att.com.