OUR 15TH YEAR

PACKET®

CISCO SYSTEMS USERS MAGAZINE

FOURTH QUARTER 2003

Full Service Branch 24

Equal Access for the Remote Office User

CISCO SYSTEMS

cisco.com/packet

High Availability Voice over Cable

A Realistic, Affordable Goal for Cable Operators Who Take a Network End-to-End Approach

BY JANET KREILING

CABLE OPERATORS VIEW IP TELEphony over cable as a tantalizing new revenue stream—a natural next move for the massive residential IP data networks they have built. That is, until they start thinking about the availability standards "everybody" knows to be true of the traditional voice network—the widely touted "five nines reliability and 50-ms failover"—and the cost of matching that level of availability.

But in this case, it so happens that five nines for the traditional voice network is a myth. Yes, certain devices in the network can achieve it, particularly the core switches and routers. Certain types of systems and transmission pathways achieve 50-ms failover. But not every system is available 99.999 percent of the time, and not every failure is fixed in 50 ms. The public switched telephone network (PSTN) actually runs at 99.94 percent availability end to end, according to the PacketCable high availability model. That's the target for cable operators to hit, and it can be done at a reasonable cost.

To accomplish the target, cable operators must look at their network end to end—from the cable modem to the hybrid fiber coax plant to the cable modem termination system (CMTS) through the regional network, the IP backbone, and the PSTN gateway on the other side. This network perspective allows cable operators to create an availability budget for specific network segments that maximizes return on investment, and helps them choose the specific techniques and technologies that increase service availability at the lowest cost. Note that the availability that counts is the availability of service to customers, not the individual percent availability of every network element.

"The three important drivers of end-to-end network availability and end-user service availability are the products in the network and their individual availability characteristics, the end-to-end network architecture in which they are deployed, and the operational environment that surrounds them," says Navin Thadani, manager for cable business development at Cisco.

EMILY BURCH

Products might have redundant route processors, line cards, and the like, and these affect the overall service availability. Individual product features such as Route Processor Redundancy (RPR+), Stateful Switchover (SSO), and Nonstop Forwarding (NSF) also play a crucial role in end-to-end network availability.

Also, the availability of a product is intrinsically tied to the architecture in which it is deployed. "Consider, for example, a product whose line card fails," explains Thadani. "In one architecture, the traffic starts flowing in less than one second, and in another architecture, the traffic starts flowing only after two to three minutes. Architectural considerations are of paramount importance, and the routing architecture specifically plays an important role in end-to-end availability."

The operational environment includes practices such as sparing policies, ensuring that maintenance contracts match your Mean Time To Repair (MTTR) standards, scheduling of maintenance, lab testing of new systems, operator training, troubleshooting guides, and change management. Each affects availability.

"We've found that in certain cases, cable operators can achieve about an 80 percent reduction in downtime for a 20 percent increase in capital expenditure, and this level of expenditure is often enough to give them availability equivalent to the PSTN," says Thadani.

Cisco's work on availability in cable networks is part of a broad program to help cable operators build a voice solution that attracts and retains customers, adds Tarun Loomba, director of the Cable Business Development Group at Cisco. "New services, such as voice, drive new requirements in the network, such as QoS [quality of service], security, and reliability. Cisco is developing programs to achieve those at reasonable cost as well."

The Availability Equivalent for Cable

The 99.94 percent figure for the PSTN was derived by PacketCable, an initiative within CableLabs that develops specifications for real-time services over two-way cable plant. Using Telcordia specifications for traditional circuit-switched voice telecommunications systems, PacketCable analyzed the PSTN segment by segment to arrive at the end-to-end percentage. (The data are published in PacketCable's report, *VoIP Availability and Reliability Model for the PacketCable Architecture,* PKT-TR-VoIPAR-V01-001128.)

In developing its specifications for voice over cable, however, PacketCable did not simply superimpose the PSTN architecture and metrics appropriate for it onto a cable network. "An IP network is fundamentally different from a fixed-path, circuit-switched one," Loomba points out. "It uses a different signaling protocol, different network components, and a different architecture. Intelligent systems throughout

MEASURING UP: PacketCable allocates percentages of unavailability for each component of a voice-over-cable network.

PACKETCABLE AVAILABILITY MODEL: IP INFRASTRUCTURE AND THE PSTN

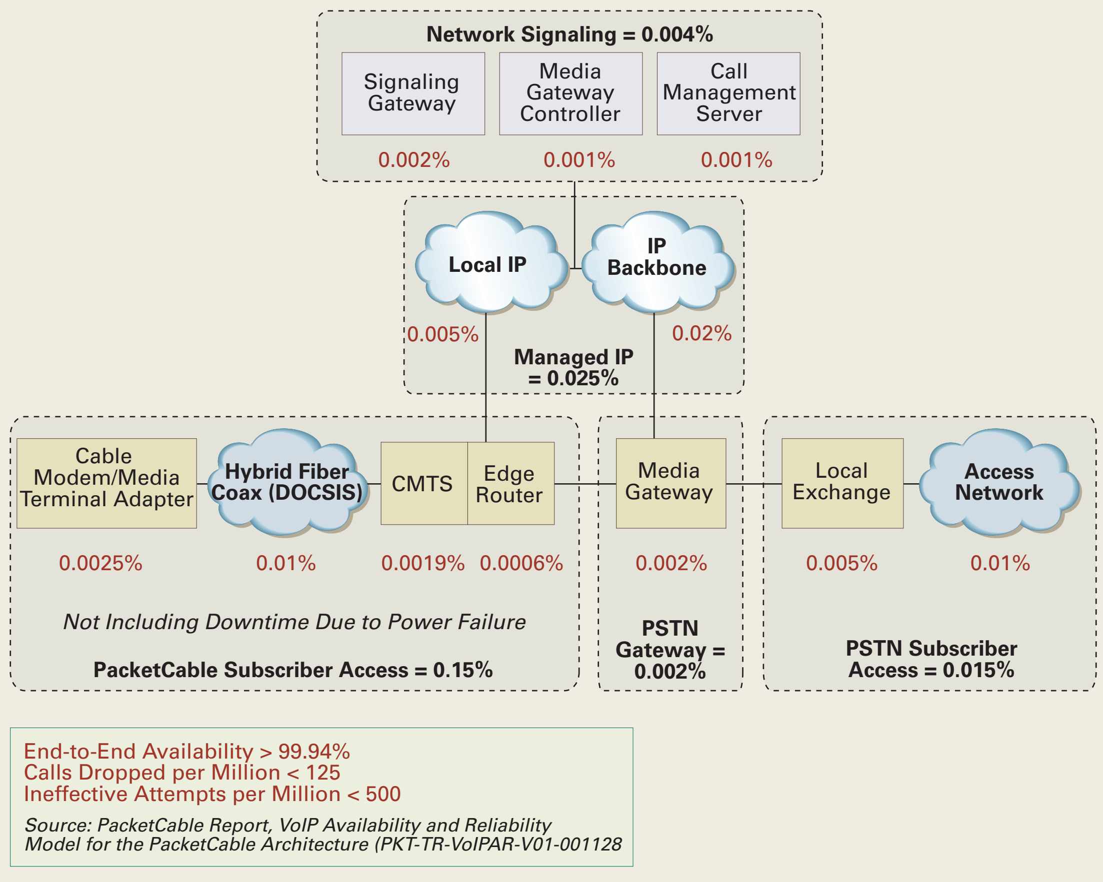

the network produce path diversity—if at any point in the network, a system goes out, the next packet is simply routed around it. This difference means that we need to approach availability in a different manner."

The overall measure is *service availability—availability from the user's point of view.* PacketCable's assessment includes three metrics that comprise its definition of service availability: network downtime, the number of calls dropped, and the number of ineffective attempts. The network end to end should be up at least 99.94 percent of the time, which translates to a maximum of 315 minutes of downtime per year. In addition, there should be no more than one in 8000 calls dropped, and no more than five in 10,000 ineffective attempts.

These two metrics are defined in the user's terms. A call is deemed dropped when the caller experiences dead air for three seconds or longer. That is the amount of time most callers are willing to wait for a call to come back up. Dropped calls generally result from a failure in the bearer path of the call for more than three seconds. They can also be expressed as defects per million—calls dropped, or DPM(CD)s. Dropping one in 8000 calls would yield a DPM(CD) rate of 125.

Ineffective attempts are usually caused by failures in the call's signaling path. The industry has reached a consensus that the time limit for an ineffective attempt is 30 seconds. The user tries once, gets nowhere, hangs up, and tries again successfully. If this happens within 30 seconds, it is counted as a successful attempt. Anything longer is counted as an ineffective attempt. Ineffective attempts can also be expressed as DPM-ineffective attempts, or DPM(IA)s. Five ineffective attempts in 10,000 translates to a DPM(IA) rate of 500.

To achieve these metrics, PacketCable has construed an availability budget with a maximum availability loss of 0.06 percent across the IP network. The limit on downtime is 315 minutes per year; on DPM(CD)s, 125; and on DPM(IA)s, 500. After assessing the characteristics of network components and subsystems, PacketCable assigned the subscriber access portion a maximum loss of 0.15 percent; the managed IP network, 0.025 percent; the PSTN gateway, 0.002 percent; and the PSTN subscriber access portion, 0.015 percent (see figure, page 62).

The availability budget has been broken down even further to assign maximums for specific components and subsystems within the network segments. For example, the hybrid fiber coax network should be out no more than 0.01 percent of the time; the CMTS and edge router, no more than 0.0025 percent; the local IP backbone, no more than 0.005 percent; the IP backbone, no more than 0.02 percent; and the signaling components, no more than 0.004 percent (includes the call management server, media gateway controller, and the signaling gateway).

Cisco was a major participant in developing these specifications; for example, its engineers developed the availability budget for the IP network and negotiated maximum downtimes among a number of systems vendors. The availability budget accomplishes two things, says Cisco Distinguished Engineer John Chapman. "It gives us both a way of breaking up the architecture into smaller blocks, and it helps us measure availability against a budget so we can get an end-to-end picture."

The Network End to End

Building a network based on the specifications in the availability model takes more than just looking for devices with the best-stated availability ratings and putting them together. Network elements or devices need to *work* together. Hardware and software redundancy; resilience features such as the aforementioned NSF and SSO; routing optimization techniques; fiber transport technologies with rapid failure recovery such as Dynamic Packet Transport, Resilient Packet Ring, and wavelength protection switching; and many other technologies are crucial to overall network availability.

Routing optimization is one technique that alone can eliminate many dropped calls and ineffective attempts. Thadani cites, for instance, an Ethernet network in which the CMTS is directly connected to redundant aggregation routers, A and B, which in turn are each connected to redundant regional routers, C and D, that lead to a national backbone. Assuming Open Shortest Path First (OSPF) routing, if the Gigabit Ethernet line card on aggregation router A fails, the directly linked regional router C will detect the problem within 100 or so ms. Router C then takes less than one second to recompute the shortest path and redirects downstream traffic to aggregation router B. Downstream traffic heading for the CMTS is redirected in less than one second.

Upstream, however, upon receiving a Link State Advertisement (LSA) about the card outage, the CMTS (or any router) waits the default amount of five seconds (SPF computation hold timer) before recomputing the shortest path. So, the total delay in this direction is more than five seconds.

"This timer was set 12 or 15 years ago, when people were transmitting just data," Thadani says. "That timing doesn't work for voice." In some cases, depending on the underlying architecture, it might be possible to lower the CMTS SPF computation timer down to one second, which permits reconvergence of the Layer 3 signal in less than the three seconds desired.

For a network that uses a Layer 2 multi-access network between the aggregation and regional routers, rather than direct links, optimization has to be done slightly differently. In this case, it might be necessary to reduce the hello and dead timers in addition to the SPF computation timer to achieve the less-than-three-second outage, according to Thadani.

Continued on page 65

Cable High Availability, Continued from page 63

An analysis of optimized routing in a typical end-to-end cable network "that employed reduced SPF computation timers, dead timers, fast hellos, and improvements in SPF calculation algorithms showed a 32 percent reduction in dropped calls," says Thadani. It is important to note that cable multiple service operators (MSOs) must carefully scale test these Layer 3 optimization techniques in their labs before rolling them out. If not implemented correctly, they might cause instability in the network, leading to additional downtime.

Other crucial pieces of the system-level, end-to-end approach, says Cisco Distinguished Engineer Mod Marathe, are availability modeling and operational practices. "Cisco has developed a set of tools that model availability in our own systems and the networks we put in place and also in customers' networks, to help them understand the availability implications of what they are doing," says Marathe. Among the sophisticated analyses these tools perform are the different types of redundancy and their effects in the network, repair times, and switchover and various switchover success probabilities. Cisco works with cable MSOs to analyze the availability characteristics of their networks, determine whether they can meet the PacketCable high availability specifications, and recommends network design enhancements to help MSOs exceed PacketCable specifications—thereby helping MSOs get ready for "primary line" telephony/voice-over-IP deployments.

> "The three important drivers of end-to-end network availability and end-user service availability are the products in the network and their individual availability characteristics, the end-to-end network architecture in which they are deployed, and the operational environment that surrounds them."
>
> —NAVIN THADANI, MANAGER FOR CABLE BUSINESS DEVELOPMENT, CISCO

The Economics of Availability

Achieving the PacketCable specifications might require some capital outlay depending on the availability characteristics of the existing network. Starting with a basic cable network that includes no redundancy, Thadani has analyzed the cost and return of various amendments to it, looking at their effects on the specifications for downtime, DPM(CD)s, and DPM(IA)s. The basic network—including the CMTS, aggregation switch, the IP backbone consisting of two SONET/SDH optical devices, the distribution switch, a softswitch, and a PSTN gateway—would cost about US$66.50 per subscriber. This network would have approximately the following characteristics: 325 minutes of downtime, 36.9 DPM(CD)s, and 500 DPM(IA)s—more than the PacketCable specifications.

Adding a redundant distribution switch or router reduces downtime by 11 percent, DPM(CD)s by 17 percent, and DPM(IA)s by 10 percent. As a second step, Thadani suggests adding a redundant Gigabit Ethernet card to the SONET/SDH devices, resulting in 19 percent less downtime, 11 percent fewer CPM(CD)s, and 18 percent fewer DPM(IA)s. Third, add a redundant aggregation switch or router, for an additional 31 percent reduction in downtime, 28 percent in DPM(CD)s, and 32 percent in DPM(IA)s. Fourth, add one redundant line card for every seven in the CMTS, another redundant route processor, and a backup Gigabit Ethernet uplink. Benefits of this step are further reductions of 70 percent in downtime, 32 percent in DPM(CD)s, and 73 percent in DPM(IA)s.

These four steps would reduce downtime by a total of 84 percent, to 49 minutes; DPM(CD)s by 64 percent, to 13.5; and DPM(IA)s by 86 percent, to 70. The cost would be only an additional 20 percent of capital expenditure per subscriber, or US$13.06. That, Thadani says, "is very far from the 100 percent outlay many providers fear." And for that modest additional outlay—which is less than about one-third the average monthly revenue per subscriber—the network's availability is considerably better than the PacketCable specs.

Doing better even than the specs isn't such a bad idea. Cable providers are already looking at video on demand and regular broadcast cable over IP—both will require high levels of availability. As Chapman points out, "All networks need to achieve high levels of availability for all types of traffic. We're moving into an era of real-time communications, in which the Internet itself is becoming mission-critical."

"What started as an availability strategy for voice for the cable industry has become an availability strategy for multiple applications," adds Chapman. "It's the way we have to go." ▲▲

FURTHER READING

- **Cisco cable voice solutions:**
  cisco.com/packet/154_8c1
- **White paper: *Cisco Cable IP Solutions for High Availability Networks*:**
  cisco.com/packet/154_8c2