Netwatcher

ISSN 0890-5800

September, 2001
Volume 19.9
Internet Edition



Browse This Issue

Hints on Navigation

Important Notices and Copyright Information: READ THIS!

CIMI Corporation
P.O. Box 1290
Voorhees, NJ 08043
http://www.cimicorp.com

Notes on Navigation:

This document uses Adobe Acrobat™ linkage conventions throughout. On the issue cover are links to specific areas of the document, such as this navigation aid, copyright and use restrictions, and the current issue text. Clicking any of these will open the correct portion of the document at a "view" setting appropriate to the context you've selected. You can also open the Acrobat navigation panel (see the Help menu to find it) to see thumbnails of each page, for direct selection of your target pages.

In the text, we may include links to other sections of the issue, to our website, or to other websites. Links to issue sections will be coded red, while links to web material will be coded in light blue. Clicking on these links will display the relevant material; red links will navigate within the issue and blue links will launch your browser to view the web material. Please do not contact us if changes in another company's site layout invalidates a link we've given! Look around.

This document is designed for viewing with Acrobat 4.x or later. If you experience problems reading the document and have an earlier version, we recommend you visit Adobe's website at http://www.adobe.com to download the latest version of the Acrobat Reader.

Important Notices:

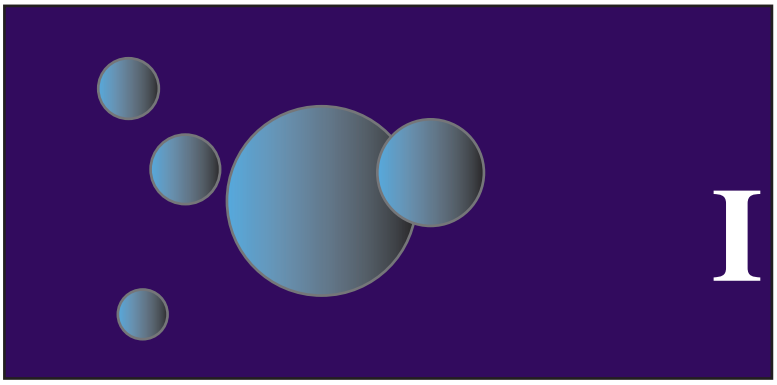
Netwatcher (ISSN 0890-5800) is a monthly publication of CIMI Corporation. Electronic subscriptions are available at $195.00 US at the time of this issue; consult us for current pricing. Printed subscriptions are not available. Subscriptions are accepted subject to qualification only. Copyright © 2001, CIMI Corporation, Voorhees, NJ. All rights reserved. No publication, reproduction, electronic storage, retrieval, or transmission of this material is permitted without the express written consent of CIMI Corporation. Journalists may quote this publication with proper attribution; quotes should be attributed to Tom Nolle, president of CIMI Corporation.

This newsletter is a copyrighted product of CIMI Corporation distributed on a use-license basis. The individual who subscribes to this newsletter may print a copy for his/her own use, and store and access this and other issues on disk to facilitate reference. That individual may also store the issues on a shared server accessible only to other employees of the same company, but those individuals may not print or make additional copies of the newsletter. No individual may e-mail copies of the newsletter to others under any conditions. No distribution of this, or any, issue outside the subscribing company is permitted. Selected portions of most issues are posted to our website, and the website material may be exchanged freely providing that we are cited as the source of all material, and that our copyright notice is included.

Company subscriptions now are available to enable all employees of a firm to print a copy for their own use, and distribute copies via email or server to other employees. Contact us for further information on these licenses.

CIMI Corporation will pay a reward of $500 US for information leading to the successful prosecution of individuals or companies who infringe on this copyright. You may contact us at the number or e-mail address found at the end of this issue if you have what you believe to be an example of such an infringement. Contact us for details on this offer.

CIMI Corporation
P.O. Box 1290, Voorhees, NJ USA
1-856-753-0004
inquiries@cimicorp.com
http://www.cimicorp.com



Inside This Issue

September, 2000 Volume 19.8

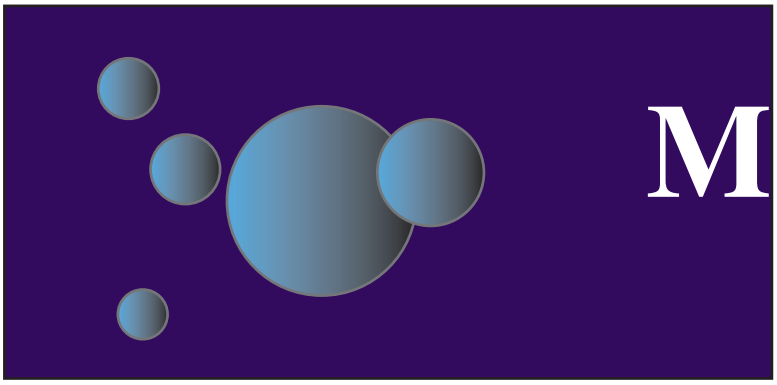
Market Area Focus will appear next in the November issue with a preview forecast of 2002 market conditions.

Management Briefing looks at the aftermath of the terrible September 11th event, and how key demand, supply, and regulatory issues might be impacted.

In the Know explores the issues of data-plane and control-plane convergence in real networks, touching on GMPLS and the MSF framework.

Strategies presents a really unique and interesting service management framework from startup player CPLANE. **Sorry, Internet readers; this is for subscribers only!**

Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved



Management Briefing

This is perhaps the very worst month we've ever had to summarize, and the toughest future we've ever had to work to foretell. We saw the awful events of September 11th play out in real time on CNN, and wondered how many friends might have been part of or impacted by the disaster. We still don't know the full extent of the losses, and we offer our condolences and prayers to those who are impacted directly.

"We mourn, but we go on" is a line from a poem that reflects the truth of the current situation. Neither we nor our readers can set national policy or determine our national course. We must do what has been asked of us by many political figures—resume as much of our normal lives as possible, and bear the burdens of our real and ongoing responsibilities. To that end, we'll summarize what we think the current industry situation is.

Clearly the US economy is taking a short-term hit from the attack and the fear of ongoing terrorist incidents. That fear is already being reflected in the stock market, which has dropped more points since the 11th than it has lost in any week in history. It is very likely that the effects of the attack will include at least one quarter in which the economy slides below zero growth, but not certain. We think that a true technical recession (two quarters of negative growth) will be avoided still.

The telecommunications industry fared better in the attack than many could have hoped. There were certainly problems with cell phone use and even with getting a local dialtone in some areas, but Verizon moved quickly to restore lost service, and much of the phone overload was caused by people calling one another about the attack and not by problems with the network itself. Whether we should increase capacity to cover that huge overrun in calling is a business and public policy issue that we'll probably need some distance and perspective to decide.

Will the events of September change the near-term future of the industry? Telecom has been in its worst slump ever, and there are both fears that this will make it worse and hopes that it will speed the recovery. We've heard them all, and also heard from people who think it ghoulish to talk about economic and business issues in the face of such personal loss. We do not mean to minimize this loss, but we cannot deal with it directly, and we believe we must each move forward to establish normal conditions in this industry, economy, and country.

There is at present no way to predict the impact that recent events will have on Congressional action on telecom reform. It's possible that the Congress will be occupied with other, larger, issues and will take no action at all this year. It's also possible that they'll see reform as an important vehicle to promote economic recovery at a time when such a recovery is most problematic, and take effective action more quickly. It's simply too soon to know what will happen.

There is a chance that the demand side of the industry will be impacted by the events as well. Companies are already facing employee concerns about working in concentrated facilities in major metro areas, and fears of flying as well. Historically these concerns have peaked within a month or so of the triggering event and faded quickly thereafter. In the current situation, the possibility that protracted action against terrorists will create a long-term state of emergency and risk must be considered. Certainly comments by the Bush Administration support that view. Thus, we expect that many firms will look closely at the idea of decentralization of facilities and employees.

Decentralization can take many forms. A company might decide to move its IT facilities out of a metro area to a more secure location. It might decide to split a big facility with lots of employees into smaller distributed facilities. It could decide to permit, encourage, or even mandate telecommuting. Whatever it does in these areas, it's likely that the result will be greater communications demand. Centralization of manpower is a reflection of the need for people to perform cooperative work, and that need won't go away just

Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved

because centralization creates risks. Communications can combat the potential loss of information access and create virtual presence to substitute for real personal contact.

The question is whether this will develop quickly, and whether the perception of the need for decentralization will persist long enough for it to impact corporate planning. Air travel was impacted by previous hijacking threats, for example, but the public accepted the incremental risk fairly quickly and there was no real change in business travel habits or policies.

We think that this time it might be different. There seems to be a good chance that the US and other countries are determined to bring the threat of terrorism to an end—at least, to end the chances of an organized mass strike of the type we just experienced. That would probably require considerable resources and a long period of action, during which the level of risk from terrorism would be higher than usual. How much higher, and for how long? That's something we just can't know at this point, but we believe that 2002 budgets will reflect a trend toward decentralization in all of the ways we've summarized here.

Communications impact of decentralization comes in two ways. First, it requires greater bandwidth between operating areas to reflect the fact that tight cooperation of employees and efficient distribution of information will now require WAN connection. Up to now, most of the information browsing in a company took place at headquarters and was thus likely to be LAN-based. If we decentralize, that has to change. Up to now, workers simply walked from meeting to meeting, and in the future they may have to meet virtually via videoconferencing. Again, more bandwidth in the WAN.

The second impact comes in the type of bandwidth to be used. It's clear to most companies that self-healing packet-based networks are a better bet to survive a major disaster than circuit-based networks or even dark fiber. Routing diversity and alternate routing may exist at the circuit level, but they require human action to invoke in most networks. Packet networks are more resilient, and their response is automatic. This isn't to say that we couldn't make circuit restoration automatic, or that packet networks couldn't be badly designed enough to be highly vulnerable, but rather that the potential survivability of packet networks is higher, all other factors being equal.

The result of these impacts is likely to be the following:

1. Companies will re-examine their networks' responses to outages of all types, but particularly to outages likely to be caused by terrorist activity. This will take some time, but will probably be complete by early in 1Q02.

2. Companies will examine dispersal of highly concentrated human and IT assets to reduce their risk. Concentrations of assets create the most attractive terrorist targets, after all. We believe that this type of planning will occur in the same timeframe as above, but that execution of decentralization plans will clearly require more time.

3. Videoconferencing will expand as companies attempt to reduce the need for employee travel to attend internal meetings. There may then be an expansion in inter-company videoconferencing, as socialization of video improves and as the number of firms with videoconference capability increases. Packet transport and broadband access reform would accelerate this broader use of videoconferencing.

4. Telecommuting will expand to deal with employee concerns about working in skyscrapers or metro centers. We expect this will be most visible in the New York metro area, for obvious reasons, but we also think Washington DC and other major cities will fall into line. Again, this is likely to be promoted by more ubiquitous broadband access. Verizon seems to be ramping up to support this very trend.

5. New-age carriers will face increased pressure, because secure communications is more credible from a major incumbent than from a new-age player. In addition, Wall Street is likely to be less tolerant of risk given the stock market trends.

We think that the demand factors listed above will impact equipment investment over the next 20 months or so, both in the enterprise and carrier spaces.

Enterprise buyers will certainly be focusing on distribution of assets and protection of service connections. Expect to see significant growth in the use of backup services, disaster recovery services and alternate site services, and redundant connection services at the access level. We've already had some inquiries from users on these issues, demonstrating that planning for network problems is already underway. Distribution or decentralization will preoccupy planners next year, meaning that other network issues will probably be receiving less attention and funding than they would otherwise.

Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved

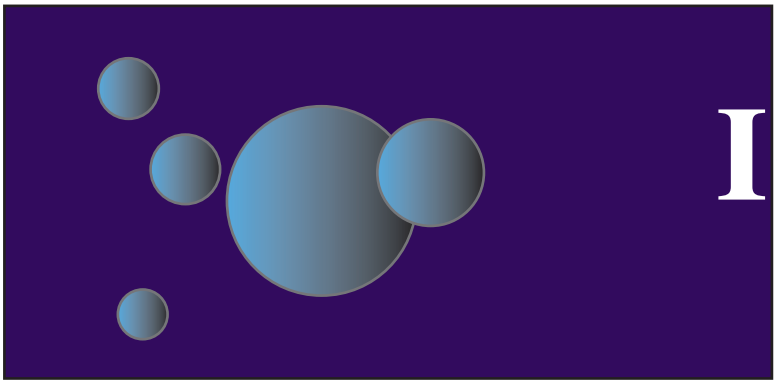
This enterprise trend will probably create a more conservative market, both because network availability is inherently a conservative issue and because it's more efficient to add resiliency to existing networks incrementally than to completely re-architect them. This incremental enhancement will then tend to solidify current technology commitments over the depreciation period of the new assets.

At the carrier level, the issue is more complicated. In one dimension, corporate demands for resilient services may encourage them to focus more on development of new-age ways of creating more traditional service frameworks. That may have a significant negative impact on the TDM equipment market. On the other hand, the fact that enterprises will probably be conservative in their consumption of new capacity may induce carriers to move away from a radical IP strategy toward one more "level 2" in nature—frame/cell or MPLS.

Overall, we think that the old adage of "following the money" will apply here. If corporate investment in networking is directed at conservative goals and to traditional partners, the likely impact will be a movement away from the "Internet-driven" period of networking and back to a conservative market. This will be exacerbated if Congress does not deal with access reform in a proactive way, because the residential broadband market would then fail to develop quickly. Given that, corporate and telecommuting interests would drive service revenue growth.

This is all going to be complicated, folks, and it will take months to sort out. We'll try to keep you up to date with our views as the situation develops.

Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved



In The Know

Behind the scenes, beneath the hype, there's still an industry trying to happen. One of the issues for that real industry is the evolution from circuit-mode networking to whatever is to come—most think it's "packet" networking. It might be possible to create a packet network today that would fulfill all the legitimate service needs of the industry, but it's far less certain that it would be possible to migrate from what we already have to any viable future packet network without excessive capital and functional risk.

One of the things that gurus here and there are pondering to reduce that risk is a form of hybridization of new-age and "old-age" technologies. This hybridization could occur in the part of the network that supports information movement (the data plane), in the part that impacts operations and management (the control plane), or both.

Data plane hybrid products are really difficult to develop. It's possible to create a device that will do TDM and packet at one time (and there are a number already available), but in truth the activities tend to be "ships passing…" rather than true integration. This is because TDM and packet are not naturally compatible.

Control-plane hybridization, on the other hand, is a real and rational process. In fact, considerable interest has been directed at the ways in which control plane technology from the packet world can be applied to the circuit-switched world. This activity is culminating in something many have already heard of—GMPLS.

This month, we're going to look at the issues of "convergence" in transport networks to see if there's any real progress being made, and whether there's a hope for a convergence of view if not of technology.

Data Architectures

Optical transport of information is currently based largely on SONET, which is a synchronous optical extension of the TDM-based T-carrier system. SONET provides for synchronous paths point to point, meaning that clock information is exchanged and the bit timings are synchronized. It also provides, through a complex system of payload pointers, for end-to-end TDM payload synchronization even if the payloads are not properly synchronized with the transport links.

It's hard to see how something like SONET could be readily displaced as long as end-to-end TDM services are a requirement, though not everyone agrees with that conclusion. On the other hand, it's pretty clear that there's a lot of baggage in SONET relating to the payload synchronization that isn't needed if there is no end-to-end synchronous path. That gives rise to some alternatives to plain old SONET.

Alternative number one is ATM over SONET. In this architecture, ATM is substituted for the SONET "network" functions, meaning the end-to-end functions. SONET is still used for framing on each point-to-point line and trunk. If there were to be a synchronous end-to-end service, it would be offered via circuit emulation services (CES, an AAL1 service in ATM). IP and other higher-layer protocols could be built on top of the ATM layer. Many IP networks already use ATM underpinnings.

To some, this is a wasteful process because of the overhead of ATM, but the use of ATM as a new binding layer creates a convenient technology bridge from old-gen TDM to the new packet age. ATM melds well with circuit-mode services as well as higher-layer packet services like IP. QoS and synchronization are already built into ATM and reasonably mature. This architecture is thus favored in a market where migration to new services (at least to new services that generate significant revenues) is likely to be slow, so existing service preservation becomes a carrier priority.

Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved

A step up from this would be a packet-over-SONET (POS) architecture. POS drops the ATM layer of the first architecture we covered, leaving packet-mode (presumably IP) traffic directly on the SONET point-to-point links. Synchronization between devices is preserved by the SONET layer, and the use of SONET add-drop multiplexers (ADMs) in the network could create a framework where existing point-to-point SONET pipes could remain in SONET form while other new packet technology simply flowed on its own SONET payloads. In most cases, though, POS supporters are envisioning a complete elimination of TDM. That makes this approach most suitable to carriers with no current or future TDM business (more properly, no non-IP business) and thus no need to sustain TDM compatibility.

Most current router vendors are supporting the "concatenated" version of the SONET interface, which allows an IP device to consume only a part of the total SONET package. This makes the POS strategy more flexible and suitable for TDM service hybridization, purists in the POS camp notwithstanding. In this form, the POS approach is arguably the equal of the ATM approach for migration from TDM. Gigabit Ethernet can also be mapped over SONET in much the same way as IP, but with some care to preserve the packet timings.

All of these approaches sustain the SONET optical protection switching capabilities that have made SONET fail-over a benchmark for other network technologies. When this is lost, as in a "packet-over-photon" scheme, the failure of an optical path is bounced up to higher layers, and that may generate considerable recovery delays as topology convergence occurs. It's to address this problem that our last technology was developed.

The final approach is the most radical, the resilient packet ring (RPR). RPR is designed to create resilient ring topologies like SONET but at the packet level. The IEEE is working on this in 802.17, and the IETF is looking at IP over RPR as well.

RPR is a kind of packet-mode SONET ring, designed to provide the same 50 ms fail-over time as SONET but to avoid wasting the bandwidth that a standard SONET ring would waste. The big questions about RPR are whether the "cost" of the SONET waste is really significant in an era of falling bandwidth prices, and whether the technology can be rolled out in a way that pays back on the investment. This latter question is really a question of how valuable RPR would be relative to the alternative—which is simply to keep SONET in place.

Which of these data plane technologies might win? Well, here's our guess. We think that ATM over SONET is probably the winner in the access and probably regional network sense, and the reason is the general economic and industry condition. Anything other than SONET would have to be deployed in a time when a lot of capital was available for up-front modernization of the network. We don't think we're in that time. Instead, we think it's likely that the market will move slowly in modernization of the network, as revenue from new services justifies the change. That would make it hard to deploy any of the alternatives to SONET at the low level, unless we did a DSL build-out based on something like RPR. That, we think, can't happen because the technology could never be proven to ILEC standards in the time we have.

In the long-haul network, we think something more like POS is the answer. We've said for some time that ATM makes little sense on high-speed trunks. The smart play for the interoffice market would be to create a big fat IP/MPLS network and share some SONET trunks between that network and the legacy voice/data/TDM stuff. As the latter could be migrated off TDM, it would migrate onto the IP side, and we'd eventually evolve to a straight POS or DWDM/packet-over-photon system.

The Control Plane

If data plane modernization isn't practical, what does that leave us with? Perhaps control plane convergence or modernization. There, we might have some logical directions to follow.

Let's look for a moment at the scenario we painted for hybrid TDM/packet networks. We've got an IP network and a TDM network (POS and standard SONET) running on common trunks. We probably have a mixture of SONET ADMs and IP/MPLS devices. To provision service in the TDM part of the network, we'd presumably use the old TL1 language and TDM craft practices that go all the way back to Bell's time. Given that we have this nice adaptive routing and discovery strategy running in parallel with the TDM network, why not capitalize on it?

That's the mission of generalized MPLS (GMPLS). In effect GMPLS provides an IP control plane in parallel with a non-IP network (it could also use PNNI or IS-IS). This control plane discovers nodes and node connectivity, so it has a map of the "real" network. If other applications like TDM exactly parallel the GMPLS structure (presumably built from the POS overlay), the routing

Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved

of connections in those applications could be done based on structural information provided by GMPLS—if all the right status and loading information were carried.

That's essentially what GMPLS is about. You first augment the topology protocols as needed to exchange information about the real network, its capacity, its status, etc. Then you add information on whatever services you want to provision. Based on all of this, GMPLS should be able to build a route map that reflects the real status of everything, and from that map select the best path for a new service connection. That connection could then be communicated across the GMPLS domain to each node, where special logic would pull the status/service data "down" out of the GMPLS domain into the real service hardware, and use it there to create the path.

The effect of all of this is to create a two-layer logical structure. At the lowest layer, which we'll call the "network plane", the real devices in the network act to switch information along the real network trunks. In the kind of hybrid network we're usually talking about, this real structure might be based on SONET ADMs and parallel MPLS LSPs (routers, for example). Above this layer is the "control plane" created by GMPLS. The LSPs' control logic creates this plane, and in it the devices exchange not only the usual topology stuff, but information on resource status relating to all the services being controlled.

This structure is similar to the one proposed by the MultiService Switching Forum (MSF), but there are some differences that are worthy of note.

In the MSF approach, networks are divided into three main functional planes—control, switching, and adaptation. The first two form a classic stack, with the adaptation plane serving as an interface between the real world and both the other layers. In effect, the adaptation plane is a ring around the "coin stack" of the other two, with links from user into the stack layers passing through the adaptation ring. GMPLS, in contrast, defines only the two layers and leaves for implementation consideration how this structure would interface to the rest of the world.

Another difference lies in what happens inside the stack of layers. With MSF, there's an attempt to define the protocols that would link the control plane functionality with that of the switching plane devices. GMPLS doesn't define how that linkage would work, leaving that to implementation consideration as well.

We could infer that the differences here could be attributed to the mission. GMPLS is really trying to develop an architecture for generalized connection routing using topology and link state information conveyed by adaptive routing protocols (and augmented selectively for this new mission). MSF is trying to develop a universal service-to-network relationship strategy, so it has to deal with the question of how the user relates to the network itself. GMPLS is also driven more by core network optical routing considerations than anything else, so it's not so involved with setting up paths that are associated with a specific service protocol like ATM or even a voice call.

The thing that makes this stuff interesting is the fact that we can look at this multi-tier structure as an evolving network blueprint. TDM/SONET forms the foundation layer, and to that layer is grafted a second service layer that presumably supports some new packet-mode service set. This new service set is created by diverting some of the SONET payloads into the packet devices, so there is a second-tier network created over the SONET trunks in parallel with the usual end-to-end SONET stuff. Since this new packet layer has adaptive routing protocols of some sort, and since it touches the SONET network at all node and trunk points, its topology is the same as that of its parent network. That means that if the packet devices are equipped with the ability to understand the resource needs of the lower-layer TDM network, the packet network map could be used by those packet devices to create instructions on how to route. That slaves the TDM process to the packet process. As more users migrate to packet services, more of the SONET capacity is diverted into the packet "parasite" network. Eventually all the trunks are supporting the packet services and no TDM is really being used end to end. At this point, the ADMs could be removed and the POS interfaces on the packet devices used instead.

We think this is the kind of evolution to be expected in the long-haul networks. As voice and TDM margins shrink, carriers migrate their users to higher-layer services. The initial investment these carriers make in the high-layer service boxes could be harnessed to create a new service provisioning system for TDM, reducing craft cost. That builds a bridge over which the carrier migrates to a pure packet environment.

And At What Pace?

The key question, of course, is when this could happen. The answer is relatively simple to explain but difficult to quantify, if you get our drift. The migration occurs faster if there are clear, highly profitable, new packet service paths for the provider to follow. It occurs faster if the packet devices can support TDM provisioning with little incremental cost. It occurs faster if the

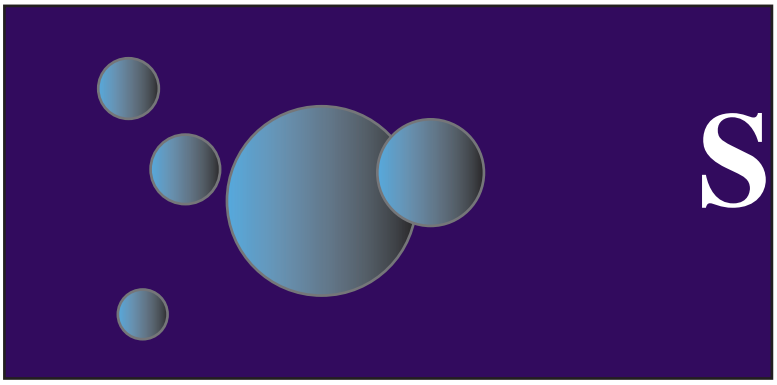
Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved

packet-to-TDM linkage needed to create the cross-provisioning features of GMPLS can be applied to the SONET devices already installed with little or no additional cost. If any of these things aren't true, the benefit case for the converging linkage between TDM and packet is weakened.

In the real world, it may be even more complicated. Because most long-haul profit margins today come from corporate data services like frame relay, we might have a kind of three-way evolution. TDM substructure might support a viable frame/cell business and at the same time support a growing IP-based service business. If the GMPLS approach can't provision both TDM and frame/cell services, the reduction in craft cost reflected by adopting GMPLS wouldn't be as large because GMPLS facilities wouldn't address the frame/cell side of the business. Since that's currently more dynamic than the TDM market, the loss of frame/cell benefits for GMPLS could be serious.

That may give the MSF people a shot at fame. MPLS has a four-year head start over MSF's approach, and the latter has more ambitious goals to boot. But MSF has at least the potential to create an answer to the control plane convergence problem that would span the list of services now offered and likely to be offered, thus maximizing the benefit case attained by adopting it. If services evolve slowly from today's TDM/frame/cell profit base to an IP-dominated base, MSF may be able to control that evolution.

Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved



Strategies

It's often true that new players in a space have an advantage over incumbents, because they have no baggage to overcome and because they can take advantage of up-to-the-minute issues and positioning. This is clearly the case with our latest, and last, service management submission, from CPLANE. But **we're sorry, Internet readers. This is for subscribers only!**

Copyright © 2001 CIMI Corporation, Voorhees, NJ. All Rights Reserved