

STEVE G. STEINBERG OCT 1, 1996 12:00 PM

Netheads vs Bellheads

The most vicious battle on the Net today is a secret war between techies. At stake is nothing less than the organization of cyberspace.



THE MOST VICIOUS battle on the Net today is a secret war between techies. At stake is nothing less than the organization of cyberspace.

It was a frequent observation among the laptop-toting 25-year-olds who crowded into the UC San Diego auditorium on an overcast morning last February that if a bomb were to go off right then, the entire Internet would collapse. It was the kind of braggadocio you hear among any large gathering of engineers, but, in this case, it was probably true.

The 250 engineers who filled the dark, wood-paneled auditorium during the two-day meeting of NANOG, the North American Network Operators' Group, were from America's largest Internet service providers - companies like UUNet, Netcom, and Sprint - and they possessed the self-confidence that comes from operating millions of dollars of bleeding-edge technology that the world increasingly depends on. They were the builders of a new age, and although lacking the brawn and defined cheekbones of the engineers in Soviet propaganda posters, they emanated the same heroic attitude of advancing civilization through Herculean struggles.

The first technical presentation that morning was by Warren Williams, a genial, chubby guy in charge of the Pacific Bell network access point, or NAP. NAPs are a lot like freeway cloverleaves – they allow traffic to flow between the independent networks that make up the Internet. Williams's big news was that PacBell was about to upgrade its NAP with higher-speed equipment that uses ATM – or asynchronous transfer mode – technology. By taking advantage of ATM, he explained, networks will soon be able to exchange traffic at speeds of up to 660 Mbps, instead of the current maximum of 45 Mbps. While most of the audience listened with a fair amount of interest, laughter kept erupting from a small cluster behind me.

That dissident cluster, I learned from the person sitting next to me, included some of the most important engineers in the room: people like John Curran, the chief technical officer for BBNPlanet; Yakov Rekhter, a lead engineer at Cisco; and Sean Doran, the lead engineer for Sprint's Internet services. These people, the engineer next to me said, believe ATM is a flawed technology, one that causes more problems than it solves. They see PacBell's interest in it as further proof that the RBOCs are doomed to be incompetent bumbler whenever they move away from their beloved voice networks.

Listening to my neighbor, I felt as if a curtain had been stripped away, exposing a hidden side of the industry – and revealing that much of what I thought I knew about ATM was wrong.

For years, the conventional wisdom has been that ATM will be the savior of the Internet – ending bandwidth scarcity and bringing about a new era of network reliability. As a computer science grad student, I had been taught that ATM was an elegant solution for the 21st century. The technology is routinely described by the computer trade press as revolutionary and inevitable. It has the backing of everyone from AT&T to Microsoft to Sun. Yet, I now discovered, some of the world's smartest and most powerful Internet engineers find the technology laughable. This discrepancy, it turns out, is not just a minor anomaly or the result of a few opinionated extremists. It is a critical battle in a war between two fundamentally opposed groups of engineers. This war has been largely invisible to even the most tuned-in netizen and has remained completely hidden to the world's telecommunications customers. Yet both groups will be immensely affected by its eventual outcome. Like most wars, it has ancient origins, blurred by the mists of

time, which go back, say, three decades. And like most wars, it is being fought over differences that might appear terribly minor.

It is a war between the Bellheads and the Netheads. In broad strokes, Bellheads are the original telephone people. They are the engineers and managers who grew up under the watchful eye of Ma Bell and who continue to abide by Bell System practices out of respect for Her legacy. They believe in solving problems with dependable hardware techniques and in rigorous quality control – ideals that form the basis of our robust phone system and that are incorporated in the ATM protocol.

Opposed to the Bellheads are the Netheads, the young Turks who connected the world's computers to form the Internet. These engineers see the telecom industry as one more relic that will be overturned by the march of digital computing. The Netheads believe in intelligent software rather than brute-force hardware, in flexible and adaptive routing instead of fixed traffic control. It is these ideals, after all, that have allowed the Internet to grow so quickly and that are incorporated into IP – the Internet Protocol.

The battle over whether to adopt ATM or to extend IP is likely to be the deciding fight between the Bellheads and the Netheads. The two protocols embody very different visions of communications leading to connected worlds with different social patterns, commerce, and even politics. In extreme terms, think of the difference between the chaotic world of the Web and the rigorously controlled, financially lucrative world of 900 numbers. The first reflects the technology of the Netheads, the second the technology of the Bellheads.

Today the defenders of these two visions are jockeying for jobs with power over networking infrastructure, maneuvering for control of the switches, and struggling to implement or thwart ATM. If the two sides agree on anything, it is this: As ATM goes, so goes cyberspace.

__ The man who runs the Net __

If you wanted to refute the aphorism that "nobody runs the Internet," you probably would look first to internetMCI and SprintLink. Together, these two companies handle the largest chunk of the data traffic flowing through the Internet worldwide.

Singling out a person at internetMCI, however, would be difficult. Vint Cerf, despite being the titular head, makes none of the day-to-day decisions. I wouldn't pick one of the engineers there either, because they're all hamstrung by bureaucracy from above. On the other hand, SprintLink - a small guerrilla operation within Sprint - suffers almost no interference from upper management. And there I'd pick Sean Doran, SprintLink's lead engineer and outspoken bad boy, as the person who runs the Internet.

Doran, we learn from his Web page, is 26 years old and a Canadian citizen. He is, in his own words, "a person who is deliberately cheeky toward people in authority, rude to people who think that they deserve respect that they have not earned, and who generally causes trouble on IRC." He tells us that he is "cute," "annoying," "brilliant," and "arrogant." He is "both bisexual and queer." His likes include "chocolate," "flirting," "fast computer networks," and "obscure poetry and allusions." He dislikes "Tories and other suchlike right-wing twits" and "people who try too hard to impress other people." He is, in short, completely likeable, but different from what you might expect of the person responsible for the technical administration at the heart of the Internet.

In person, Doran's manner is rather high-strung - always fidgeting, interrupting, and blurting out great huge paragraphs. He's extremely articulate, perhaps a reflection of his liberal arts education ("I'm not a science-y person," he admits). Every question elicits a 10-minute response, which, you get the feeling, is done more for some abstract notion of completeness than for your edification. When challenged on a point, he tilts his head, and his eyes begin to blink faster as he overwhelms your argument with masses of examples, quotes, and anecdotes.

There is no better way to evoke this response than to try to argue in favor of ATM. Doran believes the technology is "fundamentally broken" and "almost completely wrong." Not only does he have the technical chops to back up these beliefs, but the topic is liable to set him off on a rant about the war between the Internet community and the telecom community that at times sounds like a paranoid, revisionist history.

"How do you scare a Bellhead?" he begins. "First, show them something like RealAudio or iPhone. Then tell them that right now performance is bandwidth-limited, but that additional infrastructure is being deployed." Then "once they

realize their precious voice is becoming just a simple application on data networks," point out "that the Internet is also witnessing explosive growth combined with 45 percent returns on investment, compared to 5 percent growth and only 12 percent or so returns for voice." In short," says Doran, make Bellheads realize they are witnessing their extinction.

"How does a Bellhead react?" he continues. "Usually a 'hmm' or even a polite 'that's neat,' followed by desperate attempts to kill the Internet in any way possible. One way is to try to take control of it by controlling the transmissions layer – by trying to make ATM supplant the packet-switching fabric."

I point out that Vint Cerf, "the godfather of the Internet," says the opposite: he claims there is no conflict between the telecom and data divisions at MCI. "We don't have camps," Cerf had told me. "Our voice engineers work hand in hand with our data engineers." To that, Doran quickly retorts: "Vint Cerf is on drugs. He should start figuring out why so many MCI engineering résumés are circulating."

The battle, Doran continues, is between a community "that is used to making plans based on extrapolating from the previous 40 years" and one that is "trying to do the equivalent of modifying a 747 *while it's in the air*." It is "the detail-oriented telecom guys versus the computer yahoos." It's "hardware versus software." It's "people who need their kids to program their VCR versus people who grew up knowing how to do that stuff." And all of this, says Doran, can be seen just by looking at how ATM works.

__ The roots of war __

The inner workings of ATM are not, in themselves, terribly revolutionary. The core techniques were developed independently in France and the United States in the 1970s by two researchers: Jean-Pierre Coudreuse, who worked at France Telecom's research lab, CNET; and Sandy Fraser, then an engineer at Bell Labs. Both were interested in the same thing: designing a universal architecture that could transport data as well as voice at high speeds and that would make the most efficient use of the network's resources.

When, in the early 1980s, it became clear to the world's telephone companies that nonvoice, data traffic was going to be increasingly important, Coudreuse's and Fraser's research began to bear fruit. Although most in the telecom community were

well aware of the Internet and its use of packet switching, few felt that it was a technique appropriate for voice. (AT&T actually told the designers of Arpanet that packet switching would never work – an attitude that has only barely softened after 15 years of evidence to the contrary.)

Instead, the voice-dominated telecom industry charged the International Telecommunications Union with devising a new standard for carrying data and voice traffic at broadband rates. By 1990, after interminable debates, the group had finally settled on a solution that looked a lot like the systems designed by Coudreuse and Fraser. But for the standard to succeed commercially, the telcos needed to convince the computer industry – made up almost entirely of Netheads – to support ATM.

To win over the Netheads, the telcos advanced ATM as a supporting protocol – it wouldn't replace IP, it would run underneath it. ATM would be the low-level standard that everyone agreed on, and then the Netheads could build on it. ATM would be the standard-gauge railroad track on which both the voice and data armies could run their separate trains. The approach was successful, at least at first. Sun Microsystems was quick to announce its support of the protocol back in 1990, and in the academic world, computer scientists like David Tennenhouse at MIT began to research how to best take advantage of ATM.

But not everyone was swept off their feet. Many in the computer industry bitterly watched firms and co-workers embrace a technology they believed was second-rate. "What puzzles me about ATM," says Steve Deering, a networking researcher at Xerox PARC and a longtime Internet guru, "is that the computer industry paid attention. We always used to ignore their protocols like X.25, or ISDN. I have this fantasy that all these phone company guys sat around and said, 'How about we come up with something *completely *brain-damaged. Then maybe the computer guys will notice.'"

Today, the battle is at an impasse. In the backbone switch and router industry, a US\$1.6 billion dollar market, an equal number of companies are arrayed on both sides. For example, Fore Systems has put all its money in the ATM basket; NetStar offers grudging support while continuing to argue vehemently against it. Cisco, the industry giant, has hedged its bet: last April the company purchased ATM switch vendor StrataCom for \$4 billion. Among the big ISPs, the debate is also undecided:

MCI is now using ATM equipment in its Internet backbone, but Sprint very definitely is not.

Despite the current stalemate, it is possible to predict where the battle is headed by closely examining the three critical debates that underlie the fight between Bellheads and Netheads. The first, and most fundamental, is the battle over whether a dedicated connection should always be set up between people who are communicating.

-- Who foots the bill? --

At the heart of the ATM debate is really an older argument," says Brian Reid, the elder statesman and gnomish director of Digital's Networking Systems Laboratory. "It is the debate between packet-switching fans and circuit-switching fans; two sides with irreconcilably different points of view." It is a distinction that also underlies the most durable myths and clichés about the Internet. When people point out that "the Internet interprets censorship as damage and routes around it," they are really talking about the difference between circuit switching and packet switching.

Circuit switching goes back to the original days of the telephone network, when switchboard operators would connect two phone lines with patch cords to create an electric circuit. Today the circuits are virtual rather than physically distinct wires, but the principle of setting up - of **predetermining** - a connection remains.

When I use the phone at my office in San Francisco to call a friend at MIT, the telephone switches between us first decide on a route for our voices to travel. The route might go from San Francisco to Austin to Newark to Boston. Once that connection is established, my friend and I have full use of that conduit for as long as the connection stays up. But if one of the intermediate switches is destroyed by some calamity, say, in Austin, then our connection dies.

The Internet, on the other hand, doesn't require connections to be set up ahead of time. I can just send a packet to my friend without waiting for permission. How the packet gets there depends on the second-to-second status of the network's topography. Internet packets act like tiny autonomous agents that are able to find their own way. Both the advantages and disadvantages of packet switching are due

to the difficulty of keeping track of traffic that can come from any direction. It makes packet-switched networks harder to censor (or destroy), but it also makes them harder to manage.

"Traffic management" is a phrase with almost totemic weight among telecom engineers. So it's not surprising that when they looked at data traffic, their first thought was to make the packet switching of the Netheads act more like the circuit switching traditionally used for voice. Or as David Sincoskie, one of Bellcore's key people who helped make ATM a reality, tells it: "Our goal was to prove we could achieve the same quality of service as circuit switching, but with packets." Which is exactly what ATM does. Voice and data are both split into small, equal-sized packets. But a virtual connection, called the VC, must be established before these packets can be sent. That's because ATM packets are lobotomized: instead of knowing where they want to go, as in a true packet-switching network, they just know the ID number of the virtual circuit they belong to, and they must stick to that predetermined path.

Digital's Brian Reid compares ATM to the early days of plastic, when the magical new substance was first made to look like existing natural materials such as wood and ivory. Similarly, he says, "ATM is imitation packet switching, and the result is that you lose much of what makes packet switching special."

Nonetheless, there are some technical arguments in favor of ATM's virtual-circuit approach. For example, it takes fewer bits for a packet to carry a virtual-circuit ID number than it does to carry a complete destination address. And since the table of virtual-circuit IDs is smaller than the complete table of destination addresses, ATM switches can do the necessary lookups faster.

But the decisive argument, at least from the telecom point of view, is that it is easier to manage traffic when it's cleanly separated into unique connections. Billing makes for a good – and not unimportant – example. It's very hard to count IP packets and decide who should pay for them. But it's easy to keep track of who opens a connection and how long that connection stays open. In short, ATM would allow Internet providers to charge by usage, instead of a flat rate.

That's an ability many experts, even outside the Bellhead coterie, are quick to embrace. Economists like Hal Varian of UC Berkeley have long argued that some

kind of usage-sensitive charging system must be instituted for the Internet to remain sustainable. Flat-rate pricing, after all, doesn't provide any economic incentive to refrain from wasting the Net's limited bandwidth. One result is undergrads who, for \$29.95 a month, clog up the Internet with CU-SeeMe sessions.

But these arguments don't fly at all with Nethead zealots. As PARC's Steve Deering delights in pointing out: "IP is hard to charge for? That's not a bug, that's a *feature*!" True Netheads have an almost blind faith in the power of new technology - rather than complex charging systems - to solve bandwidth woes. (Admittedly, this "blind faith" has been borne out for the last 30 years.) But more generally, Netheads' dislike of ATM's circuit switching in drag has to do with how they think about traffic.

"Essentially, ATM is a very predictable technology. There is no scope for anything that doesn't fit the model," begins Doran, in full-power didactic mode. "ATM comes from a culture that believes it can do long-term predictions of traffic levels. The problem is that it applies a rigid scheme to what is intuitively very fluid."

Indeed, Bellheads use phrases like "rate control" or "traffic shaping" when talking about ATM. Traffic is something to be tamed, sorted, and organized. Even though hundreds of different connections end up multiplexed on a single fiber-optic pipe, every ATM connection is treated as if it has its own, narrower pipe. It's as if a group of US highway engineers went to Bombay and started building things like dividers and car-pool lanes.

Netheads, on the other hand, cheerfully admit they have no idea what traffic will look like next month. It is easier, they say, to have the packets fight it out among themselves than to try to force some kind of order on them. Traffic is unpredictable, spontaneous, liable to jump up in strange patterns. Doran sounds almost mystical when describing the traffic that flows across his network: "Each time one thinks one understands the variables, more appear."

This philosophical divide should sound familiar. It's the difference between 19th-century scientists who believed in a clockwork universe, mechanistic and predictable, and those contemporary scientists who see everything as an ecosystem, riddled with complexity. It's the difference between central and distributed control, between mainframes and workstations.

__ Fractal facts __

Only recently, however, have Netheads been able to point to analytical evidence, rather than just anecdotes and vague impressions, to support their worldview. They were in the same situation as Parisian engineers in the 1800s, when the structure of the water supply system shifted from a hierarchy to a network. One French engineer, struggling to build the newly conceived network system, wrote in 1826: "The whole problem is to know beforehand the pressures that exist at the different connection points, and it is precisely on how to get these results that we lack factual knowledge, and consequently, the necessary rules of calculation." Similarly, Internet engineers lacked solid numbers about what traffic "looks like" and the quantitative tools to model it.

For the last 20 years, most analyses of data traffic – including those ATM is based on – have assumed that traffic bursts follow a Poisson statistical distribution. The Poisson distribution is a mathematical model for describing physical processes over time; it's used to predict the frequency of auto accidents, atomic disintegrations, and, perhaps most famously, the fall of bombs and Slothrop's sexual encounters around London in Thomas Pynchon's *Gravity's Rainbow*.

One important property of a Poisson distribution is that the statistical variance between events is fairly small. It's easiest to understand this by thinking of a graph of telephone use, with the vertical axis representing voice traffic and the horizontal axis representing time. If the horizontal axis has a very high granularity – say, one tick mark represents one minute – you will see bursts of traffic as people pick up their phones. However, as you zoom out – say, one tick mark represents one day – those bursts average out, and traffic appears as a steady line. In other words, daily loads are fairly predictable.

That kind of smoothing is good news for telephone engineers, because it means they can safely optimize the network for this steady level of traffic. But in 1993, a group of mathematicians at Bellcore and Boston University published a paper that showed the Poisson distribution is worthless when it comes to describing data traffic. Their startling but widely accepted paper, "On the Self-Similar Nature of Ethernet Traffic," reported that if you graphed data transmissions over time, the result looked something like ... *a fractal*.

A fractal is a pattern that appears the same no matter what scale it's viewed at. Look at a small region of a Mandelbrot fractal at 400-power magnification, and it will look just like a duplicate version of the larger region that surrounds it. Similarly, if you graph the peaks and valleys of data traffic over time, and then vary the scale of the horizontal axis, you don't see any smoothing. Traffic remains bursty no matter what the time scale. That makes it hard to optimize data networks, because if you provide only enough capacity for the "average load," then you'll get swamped by the inevitable bursts.

According to Walter Willinger, one of the paper's authors, "Networking people were extremely happy when we published our paper. They finally had this ammunition. Now they could say, 'See, we told you that your models didn't work!'" Less happy were those who saw the theoretical underpinnings for their networking technology suddenly give way. As Willinger says: "There is no doubt in my mind that this will have serious implications for ATM."

But exactly what the implications are turns out to be harder to determine. One of the reasons for the Poisson distribution's popularity is that the method is easy to work with. Efficient tools for working with fractal-like distributions, on the other hand, simply don't exist. So until researchers like Willinger develop ways to apply this great insight, we can't describe its impact in anything but broad strokes.

We *can* specify the piece of the ATM protocol most likely to blow up when faced with massive amounts of fractal-like traffic. It's what the Bellheads call quality of service (QoS), and it represents the second of the three fault lines in the war between the telecom and data communities. It's the debate over whether to make hard guarantees about the quality of service customers can expect, or whether to promise only "best effort."

__ Malthusians versus Cornucopians __

To understand this quality-of-service issue, I went to a tall, windowless building on Franklin Street in downtown Oakland to see the PacBell NAP and the new StrataCom ATM switch that Walter Williams had announced at the NANOG conference.

Most of the cavernous building is now empty, due to years of technological progress that have shrunk giant crossbar switches to the size of microchips. The few briefcase-size boxes of blinking lights that make up the Oakland portion of the NAP are located in a gray concrete bunker on the tenth floor, along with, when I showed up, three bored, tattooed men. (None of whom, it turned out, were aware of what exactly the boxes were for – an admission that, in my mind, sealed the fate of PacBell's Internet efforts.)

During my tour of the equipment cages, Jim Diestel, the director of Internet NAP marketing, walked me through the reasons ATM is the way of the future. He spoke of how, because ATM had been designed with gigabit speeds in mind, it was much easier to build fast ATM switches than fast IP ones. (Indeed, StrataCom switches were able to support 155 Mbps connections long before Cisco equipment could.) Diestel explained that, thanks to ATM's status as an international standard, more and more manufacturers were entering the market, driving prices down to unprecedented levels. And Diestel twice mentioned how ATM technology would bring a higher quality of service to the Internet. No longer would I find myself watching the pulsating Netscape icon, wondering what went wrong.

What he didn't mention was that PacBell had tried to use ATM equipment in its NAP once before, back in 1994. That attempt ended in disaster: service quality plummeted and the equipment had to be thrown away. The paradoxical reason: ATM's vaunted ability to provide high-quality service.

An ATM switch acts a lot like a railroad switching yard. It takes packets from a bunch of different lines, then switches them onto the correct output line. A packet will arrive on a line from Los Angeles, for example, then get shot out on a line to Denver. The tricky part is when packets from different input lines contend for the same output line.

IP's solution is to handle packets on a first-come, first-served basis. That's nice and simple, but it isn't always fair. If a long stream of packets headed for New York arrives on one input line, any packets going to the same destination that arrive even a nanosecond later have to wait behind the entire stream.

ATM tries to adjudicate this process so that everyone gets an equal share of bandwidth. This is done by specifying exactly how often each connection can send

a packet. If a connection's packet arrives too early, it must wait at the switch until its time slot comes up; if the packet comes too late, it must wait for the following slot.

This vision of packets goose-stepping their way down the Internet is striking, but ultimately doomed. In order for it to work, the ATM switch must store all the untimely packets in memory until they can be sent. And there will be a lot to store. Because data traffic follows a fractal-like distribution, there will always be uneven bursts of packets – they will never follow the steady distribution that ATM desires. This is the unfortunate reality that killed PacBell's original ATM equipment: untimely packets filled up the switch's memory, forcing it to drop additional packets on the floor and hope no one would notice. But people did.

PacBell's new switch, the StrataCom BPX, has a whopping 192 Mbytes of memory on board just to buffer packets. This seems to be working, but many feel that once the traffic loads increase, it too will choke. Besides, IP advocates point out, this buffer memory ends up being the single most expensive part of the switch – sometimes as much as 30 percent of the total cost. Why not get rid of it?

Instead, Netheads say, forget about quality-of-service guarantees. Bandwidth is cheap – simply overengineer the network to provide so much bandwidth that nobody cares about fairness. If the lines are running at 100 Gbps, as they should be in a few years, nobody is going to care if their packet gets stuck behind a bunch from someone else – it will still be plenty fast. Plus, say the Nethead radicals, Internet applications should be *adaptive*. They shouldn't need a guaranteed amount of bandwidth.

Steve Deering, the PARC researcher, uses a story about the 1989 Loma Prieta earthquake in the San Francisco Bay area to make this point. "During the earthquake, a professor at Stanford was on a business trip and desperately wanted to find out if his family was OK. That's only a few bits of information, but all the phone system had to offer was a perfect 64-Kbps channel or a busy signal – and of course he, and almost everyone else trying to phone the Bay area that day, got nothing but a busy signal." Instead, Deering suggests, "An Internet telephone application that used progressively aggressive compression – or even that said, 'Insufficient bandwidth for voice: please type your message' – would have allowed the available bandwidth to carry many, many more useful bits."

This position is an absolute anathema to Bellheads, because it goes against principles that are deeply imbued in their culture. These are principles that reach back to the early part of this century, when AT&T floundered against competition from hundreds of small independent phone companies. It was only when Theodore Vail retook control of the company and emphasized high quality of service as a competitive advantage that AT&T regained its footing. Even in the 1970s, AT&T argued against allowing MCI into the long distance market, partly on the grounds that the new entrant's sound quality would be unacceptable to the American public. (Making it wonderfully ironic that today MCI's Cerf supports the need for quality-of-service guarantees on the Internet.)

Still, there are plenty of technical arguments that can be made in favor of quality-of-service guarantees, and virulent debates on the subject seem to flare up every few months on one technical mailing list or another. A representative example occurred last spring on the ip-atm list. Like all debates on this topic, it petered out after both sides had reduced their arguments to dueling historical anecdotes. On the pro-QoS side, someone pointed out that "the golden rule of disk space applies to bandwidth too: no matter how much you have, it's 90 percent utilized. Hence the need for QoS." On the anti-QoS side, Deering drew on an opposing historical precedent. "Determining the right way to provide QoS is about as interesting as the obsolete debate about what the *right *charging model is for CPU usage on a timesharing system. That debate was rendered irrelevant by the availability of cheap processors, and I predict the debate about the right model for QoS will be rendered equally irrelevant by cheap bandwidth."

In short, it boils down to the Malthusians versus the Cornucopians, with the Bellheads predicting bandwidth scarcity and the Netheads predicting abundance.

Fortunately, a middle ground is emerging. A group of IP advocates is developing a scheme called RSVP, which provides the main benefits of ATM's QoS mechanisms (essentially by locking up a set amount of bandwidth) but which is far more adaptable to different traffic patterns. This protocol could be used by special applications, such as online games, that really require bandwidth guarantees.

While ATM co-creator Sandy Fraser suggests that RSVP is simply a case of the Netheads "reinventing telephone technologies and giving them a new name," it

really is a new twist on what came before. In essence, RSVP applies a software sensibility to a problem that ATM tries to solve with rigid hardware.

-- Leave it to the diplomats --

The third and final core debate in the war between the Bellheads and the Netheads is over something terribly mundane: the number of bytes in an ATM packet. Yet this may be the fault line that ends up mattering most. Packet size is the primary reason Doran hasn't deployed ATM equipment in Sprint's network; it is also the reason MCI's Vint Cerf did.

An ATM packet, technically known as a cell, is 53 bytes long. The first five bytes (called the header) contain information about the cell and its connection; the next 48 bytes (the payload) contain the actual data. Contrast this with an IP packet, which has a 20-byte header and a variable-length payload of anywhere from 12 to 65,536 bytes. It's difficult to convey how insane ATM's cell scheme sounds to anyone in the data community, but it's roughly equivalent to Ford announcing a new car that is shaped like an upright obelisk. Sure, it could be made to work, but it's neither aerodynamic nor practical.

Computers, after all, are binary systems. That means they like to work with powers of two: 64, 128, 256.... At the very least, they like to work with even numbers, which divide easily. What they definitely don't like working with is prime numbers. But if you had to pick a prime number for cell length, you would at least want to make it large. The shorter the length, the greater the percentage of overhead. With an ATM packet, 5 of the 53 bytes are essentially worthless. That means you're throwing away about 10 percent of your bandwidth right from the start – an effect known as the "cell tax." Cerf sums up the situation: "The only good thing about 53 is that's how old I am."

So why did 53 get chosen? Because an extremely corrosive influence was injected into the technical debate over ATM: international politics.

When the telecom community wants to set a standard, it must go to the International Telecommunications Union, a UN treaty organization. Each country is represented by a delegation. Arguments are carried out formally and with great attention to protocol. And final votes are cast by government representatives. (In the

case of the US, this means someone from the State Department.) It would be hard to come up with an atmosphere less conducive to setting technical standards.

Imagine, then, the 1988 ITU meeting in Geneva, where delegations from two dozen countries met to decide on the length of an ATM packet. It was immediately obvious that consensus was going to be difficult: The Europeans wanted 32-byte payloads, because that would be best for voice, while the Americans and Japanese wanted 64-byte payloads, since that would be better for data. (Although the American delegation was firmly in the Bellhead camp, it was at least aware of Nethead concerns. The Europeans saw little demand for data and didn't expect that to change anytime soon.)

According to Richard Vickers, who was part of the US delegation, the discussion quickly turned confrontational, with the US and France becoming the main combatants. As the invective became more heated, pressure built to solve the question the diplomatic way: split the difference. And so was born the 48-byte payload, which, combined with a 5-byte header (the smallest that the ITU could agree upon), added up to a 53-byte cell.

No one was terribly pleased with the result. Sandy Fraser, forced to watch his technology haggled over like some kind of border dispute, now says, "In my view, they picked the worst of both worlds." The Netheads were far more vocal. T-shirts and buttons with a slash through a red *53* became popular at subsequent Internet conferences. Doran practically sputters when the topic comes up, claiming that the short cell size results in almost *30* percent overhead when used to carry IP traffic. That's because longer IP packets have to be distributed across multiple cells, often leaving the last cell only partially filled. "Why should I throw 30 percent of my bandwidth away?" he asks incredulously.

Carl Cargill, who has studied high tech standards for 14 years and currently handles standards strategies for Netscape Communications Corp., sees the episode as a study in the opposing ways that the telecom and data industries think about standards: "Interconnection is really important for telecom companies. So first they agree on a standard, *then* they compete. Computer companies, on the other hand, compete right off the bat and let the market pick the standard."

The advantage of letting the market decide is that bad standards are less likely to survive. (There are counterexamples, of course, like VHS.) The advantage of the telecom industry's system – standards de jure – is that manufacturers face less risk: they don't need to worry that their products will quickly be rendered worthless by a competing standard.

We are currently watching these two approaches play out in the switch and router industry. While datacentric IP router manufacturers have hung back, waiting to see where the market is headed, ATM switchmakers have leapt ahead, secure in the knowledge that their standard is backed by the world's mammoth telecom companies. Hence, it is possible to go out and buy an ATM switch today that runs at 622 Mbps. But top-of-the-line IP routers from Cisco and NetStar still go up to only 155 Mbps.

This is why some ISPs, including MCI and @Home, now use ATM in their network. "We were forced to embrace ATM," explains Rob Hagens, MCI's director of Internet engineering. "We wish we didn't have to, but the only equipment out there that can handle the speeds we need is ATM-based."

Although MCI uses ATM in name, the company certainly doesn't use it in spirit. It has, as much as is possible, lobotomized the technology. Instead of having ATM constantly opening and closing new connections each time a user links to a Web site, MCI leaves a single wide connection "nailed up" between the main hubs along its backbone. And MCI currently turns off ATM's attempt to guarantee quality of service, since it breaks down in the face of unpredictable traffic. The cell tax? That's the unfortunate cost of doing business.

__ Truce __

If you take this kind of pragmatism to its logical conclusion, you end up at an anonymous office park in Palo Alto, California. There, a small start-up called Ipsilon is building the foundation for what will come once the war between the Bellheads and Netheads ends.

Ipsilon was founded in 1994 by Tom Lyon, who had been involved with Sun's initial ATM efforts. Big and shy, prone to staring down at the floor as he speaks, he is now

Ipsilon's chief technical officer. Amid widespread speculation in the Internet community about what was going on at tight-lipped Ipsilon, a number of other well-known data networking people soon joined the firm. They were lured, it turns out, by the opportunity of doing something so innovative that it might allow tiny Ipsilon to challenge Cisco's dominance.

Bob Hinden, probably best known for his work on IP version 6, became director of routing software. Dennis Ferguson, a longtime Internet backbone guru, joined as a member of the technical staff. Brian NeSmith, a former vice president at Newbridge Networks Corp.'s Vivid Group, became Ipsilon's CEO. All of them were data guys who had lost their Nethead religion - or at least learned to soften it with a dose of pragmatism. In early 1996 the company was 50 employees strong, had two rounds of venture capital financing behind it, and had a prototype IP switch up and running.

The IP switch is Ipsilon's flagship product and a paeon to pragmatism and cost-efficiency. The switch is stripped down, built from off-the-shelf and salvaged parts, and obsessively engineered for speed. While other switches are powered by custom or top-of-the-line microprocessors, Ipsilon uses a \$500 Pentium Pro. And while Ipsilon uses ATM, the company ignores many of the complexities that make the technology expensive. ATM's ability to guarantee different kinds of QoS is thrown out, as are many of the higher-level protocols the ATM Forum, an independent standards body, has secreted like a thick crust over the last five years. Ipsilon uses ATM to quickly switch bytes between different input lines and output lines - and that's it.

According to Lyon, the decision to use ATM was purely pragmatic. "In an idealistic world, I wouldn't go near ATM," he says. "And in theory, pure IP could be as cheap as ATM. But currently the chips necessary to implement ATM are available at a lower price." And Ipsilon can use ATM without falling victim to many of the technology's flaws. The result? "A configured Cisco runs \$75K to \$130K. We cost \$46K to \$50K," says Lyon.

What Ipsilon has done, in essence, is rip out the core of ATM - its ability to quickly switch bytes - and throw away the rest. Lyon compares this with what happened to reduced instruction set computing, the microprocessing technique developed in the early 1980s. Although RISC was technically superior to the competing complex

technique, CISC. But CISC ended up winning because it enjoyed widespread market support. Nonetheless, the ideas behind RISC have stealthily seeped into the industry, so that today RISC techniques are at the core of state-of-the-art CISC chips.

"These days, RISC is pretty much dead, except for the pieces that live on inside Intel Pentiums," says Lyon. "Because Intel maintained the momentum of the software base, they were able to take their time and pick and choose the RISC techniques that made sense to integrate into their processors. So will ATM die out? Or will it live on inside IP switches like RISC features live in the Pentium?"

His answer is subsumption rather than extinction. The result will be a sort of technological sedimentation: our networks will be built on strata of technology from earlier eras. At the lowest level of the next-generation network will be the telecom community's ATM. Built on top of that layer will be the data community's IP, and above that may be the Internet's RSVP. Data networks will subsume voice networks, but the ghosts of telecom will live on in the underlying, invisible technology.

In this sense, neither the Netheads nor the Bellheads will have "won" the war. While the Netheads may come out on top, their success will rest on the technical foundations laid by the Bellheads. Each side ultimately will see its ideals incorporated into the network. At the lowest level of networking, the hardware approach favored by the Bellheads makes sense. Higher up in the protocol stack, the Netheads' software approach is better.

As the inevitability of this result becomes increasingly apparent, there are signs that the Netheads and Bellheads are ready to call a truce. Just as there are no atheists in foxholes, the converse of the maxim also holds true: religious fervor has a way of dissipating when confronted with large amounts of cash. Netheads and Bellheads are starting to realize that if they put aside their differences and simply concentrate on building better networks with whatever tools are best for the job, they will be richly rewarded.

Even Sean Doran will admit that as enjoyable as it may be to poke fun at Bellheads, they still have plenty to contribute. "There is a useful tension between the two sides," he concedes. It's an admission made easier by the fact that, according to Doran, it's largely the telecom guys who are giving ground. "We had a meeting

recently where these very senior executives from Sprint came to our building completely dressed down – trying to fit in. They actually made fun of the one guy who did wear a suit."

"Everyone is learning to work together," Doran continues with a touch of satisfaction, "so we can make lots of money and play with our big toys."

TOPICS MAGAZINE-4.10

MORE FROM WIRED

There's a Huge Covid Surge Right Now and Nobody Is Talking About It

The US is in the midst of the largest Covid surge since Omicron, but with minimal testing and good population immunity, the wave is largely being ignored.

CELIA FORD