A NATION GOES ONLINE

INSTITUT CA*net INSTITUTE
A PROGRAM OF CANARIE
UN PROGRAMME DE CANARIE

A NATION GOES ONLINE

INSTITUT CA*net INSTITUTE
A PROGRAM OF CANARIE
UN PROGRAMME DE CANARIE

TABLE OF CONTENTS

FOREWORD

A NATION GOES ONLINE

More Canadians are connected to the Internet than any other country. This should come as no surprise, since we are global leaders in information communications technologies and Internet development. We did not get there by accident – we got there by innovation and establishing world class design expertise. Canada is proud of its advanced networking history. As this publication illustrates, we have built an Internet infrastructure which links Canadians to each other and reinforces the economic and social underpinnings which define a modern nation.

Canada's networking success is one based on partnership and co-operation between the academic and research community and the public and private sectors. The story told in these pages is a testament to this successful approach. It is not the work of a single group rather that of a series of grass-roots efforts that took shape at universities and other institutions in regions across the country. These pioneers worked to connect a population scattered over immense distances, to create opportunity from potential isolation, and to develop regional collaboration and cohesion. That determination spurred much of the early networking research at Canadian universities and ultimately the national partnerships that led to the creation of CA*net, Canada's first information highway.

Canadian networking pioneers built a communications infrastructure that would not only link Canadians to each other but also give them access to the global society. Their efforts allow us to share our culture, our creativity and to access the global marketplace – in turn, stimulating economic growth and investment toward ensuring prosperity for Canadians.

This historical account of the Canadian Internet is told by the people who worked on the front lines – the believers, the researchers, the innovators, the businessmen and the politicians. It was their commitment and hard work that led to a national network. And now, Canada has the world's first national optical network. I encourage all Canadians to continue to build, to innovate, to lead.

The story that follows helps define our history as Canadians. It provides a legacy for our children and demonstrates a commitment to continued online innovation designed to serve Canada and Canadians for future generations.

A NATION GOES ONLINE represents an important part of 21st century communications in the making. A history of which we can all be proud.

DAVID JOHNSTON, *PRESIDENT, UNIVERSITY OF WATERLOO*
CHAIR, NATIONAL BROADBAND TASK FORCE

ACKNOWLEDGEMENTS

Many people devoted their time and talent to the preparation of this document. On behalf of the CA*net Institute, I want to take this opportunity to thank them for their efforts. It's my sincere hope that all those who helped with interviews, consultations and research will be proud of what's been accomplished in these pages.

The Institute asked Omnia Communications Inc. of Toronto to assume responsibility for the management of the project and the preparation of the document. The core Omnia team, comprising David Ellis, Simon Cheesman and Mark Czarnecki, worked diligently over a period of more than two years to collect and sift through interview material, write several drafts and respond to queries and critiques from the Institute's editorial board. My thanks go to the members of the Omnia team and to the members of the editorial board: Ken Fockler, Dave Macneil and Alan Greenberg.

No fewer than 24 of the builders and founders of Canada's Internet were interviewed in person or by email. Their stories and analysis form the heart of this history, and their words will be found quoted verbatim at many places in the text. The 24 interviewees were:

Andy Bjerring	David Macneil	Richard Sexton
John Demco	Mike Martineau	Eugene Siciunas
Paul Gilmore	Gerry Miller	Henry Spencer
Steven Goldstein	Mike Parker	Roger Taylor
Warren Jackson	Kent Percival	Vincent Taylor
Ken Fockler	Joseph Reid	Bernard Turcotte
Richard Lacroix	John Robinson	Roger Watt
Jack Leigh	Peter Scott	Rayan Zachariassen

I want to single out the special contribution made by Ken Fockler. In typical fashion, he managed to play several roles in the process, not only as interviewee and a member of the editorial board, but also as a special advisor to the Omnia team.

Several other people and institutions deserve our gratitude for looking after supporting tasks of various kinds. My thanks to Industry Canada for their generous support in helping to fund the French translation of the original document. Credit for the excellent French text goes to Freynet-GagnÈ Translation Services of Winnipeg.

Thanks are also due to Bell Canada, who, together with CA*net Networking Inc., created and financed the Institute in 1997.

The Institute now works closely with CANARIE Inc. and two of my colleagues there deserve mention for their patience and unfailing support of this project. In his role as President and CEO of CANARIE, Andy Bjerring has kept an eye on our progress and offered help at many points, despite a very busy schedule. In addition, Nancy Carter, Assistant Director of Finance, has carried us through many a crisis, looking after our collaborators, finding resources and working hard to keep the whole team on schedule.

A word of thanks to Juanita Brandt, who helped get the document finalized and into shape for public release.

The story told by A Nation Goes Online, like the Internet itself, is a tapestry of interwoven narratives. While telling the complete story of the birth of Canada's Internet may be unachievable in any single account, it is nonetheless the ideal to which we should aspire, and we apologize to any and all whose stories should be in this account and aren't, or who contributed to this history and have not been explicitly thanked for their contribution. Needless to say, should any readers find errors or omissions in the text, they are the responsibility of the Institute and the editorial board, and not of Omnia or of any interviewees or other contributors.

*– **GERRY MILLER,** CHAIR, EDITORIAL BOARD, **CA*NET INSTITUTE***

INTRODUCTION

A NATION GOES ONLINE

This is the story of a handful of dedicated Canadians and how they worked for two decades to bring computer networking, and ultimately the Internet, to Canada. The narrative revolves around the launch in 1990 of CA*net, Canada's first coast-to-coast, IP-based production network – one designed for practical rather than experimental purposes and running on the Internet protocol suite. But this milestone event in Canadian communications can only be understood and appreciated in the context of what happened in the years before and after 1990. That was the challenge established in the commissioning and writing of this story.

And an improbable story it is.

While travelling the Internet is not yet as widespread or easy as watching TV or making a phone call online activities are becoming part of daily life for the great majority of Canadians. At this writing more than 50% of Canadian homes are online – more than one in four connected through a broadband or "high-speed" device. Canadians are becoming known as the world's most avid users of the Internet. The remarkable thing is that all this has happened less than a decade after the World Wide Web saw the light of day. No other communications medium in history has been adopted so quickly.

A national network was for many years nothing more than a distant dream to Canada's research community. Just as telephone systems began on a local or regional scale long before operating nationally and internationally, so computer networks were at first very modest and highly localized affairs. Building a national network presented additional and not insignificant challenges.

Unlike telephones, computers and computer operating systems as originally designed were not intended to communicate with each other. In fact, for many years computer manufacturers placed strategic importance on the development of competing networking standards and products based on proprietary technology that actually prevented the kind of open network that we accept today as the norm. Even years after researchers had developed open protocols that allowed for efficient data transmission over great distances, few in business and government circles saw any practical applications for them outside academia. It took the microcomputer and, above all else, e-mail, the original "killer ap", for the value of open networks to begin to be felt. At the root of these debates was the Internet protocol suite, TCP/IP. In spite of its highly regarded technological attributes, the eventual acceptance of TCP/IP was anything but a foregone conclusion for the networking community – and even less so for government officials, telephone companies and standards-setting bodies.

In Canada, this combination of circumstances played out against a backdrop of grass-roots efforts defined by geography and a sense of isolation among those having a professional interest (some might say obsession) in getting computers to communicate. Most of those involved were working at universities, a few were in government and the private sector. Despite differences in goals and methods, they had one overriding idea in mind that kept them going, and talking among themselves, as they searched for answers: how to create an open network that would allow computers from around the world to communicate and exchange data.

Our story thus unfolds not on the national stage but at a dozen or more locales scattered across the country, most of them university campuses. On one coast, the University of New Brunswick had developed one of the first regional academic networks – precisely to overcome the sense of isolation felt by those working on campuses far from colleagues in major centres. On the other coast, the University of British Columbia marched to the beat of a different drum, building a network called CDNnet to test a particular kind of application, using protocols not widely embraced by other Canadian efforts.

Almost everywhere in between, computer service centres and academic departments were trying to solve their special problems in ways that often had little in common with parallel efforts in other parts of the country. In Quebec, for example, the networking effort involved not just the widely scattered campuses of the Université du Québec, but community colleges and government departments, culminating in the creation of Edupac, which linked some 200 public institutions. While Canada's military establishment developed DREnet, a highly restricted, government-controlled facility, more independently-minded Canadians were linking up with the maverick Usenet community in the United States and laying the foundations for much of the Internet's "Wild West" culture.

Many readers will be familiar with the development of the Internet in the United States, starting with the early experiment in computer networking known as ARPAnet. Canadians could be forgiven for thinking that our progress in creating sophisticated networks followed a well-worn path forged by our neighbours to the south.

Canadian researchers began with similar motives. Like their American counterparts, they were concentrated in universities and determined to find more efficient ways to communicate with their colleagues – using computers that had never been intended for transmitting data or email, let alone the audio and video material that is now becoming commonplace on the Web. For many of those involved in both countries, the desire to communicate and share knowledge over distances turned into a concerted effort to develop networking protocols and applications that would eventually serve a far larger community of users, though that was the farthest thing from their minds when they first sat down to work.

Nevertheless, Canada's networking pioneers faced a dilemma encountered many times before in our history, the desire to build domestic links along the east-west axis, while resisting the powerful pull of north-south forces. Canadian researchers followed the work of their American counterparts closely – making every effort to get connected to networks that would link them not only to US centres of activity but international ones as well. The challenge was to become part of the larger international evolution towards packet-switched, IP-based networks, particularly as developed in the United States, while sticking determinedly to the goal of creating a made-in-Canada solution.

Slowly but surely, these regional efforts coalesced into a national network, though not before numerous false starts and dashed hopes. CA*net placed computer networking in a well-deserved position of prominence in the eyes not only of the academic community, but also among influential decision-makers in government and the private sector, not the least of which were telephone and high-tech companies, who were finally awakening to the full potential of this new medium.

CA*net had been operating about two years when a group representing both public and private sector interests came together to form CANARIE Inc. Its mission was and is to develop state-of-the-art, high-speed networks and networking technologies, as well as special educational and other applications serving the broader public interest. Once CANARIE was in place and subsidizing the CA*net operation, the Internet as we know it today was taking shape – attracting a larger and larger number of non-specialists and eventually going commercial. In a twist of fate, the very excitement instilled by the work of CA*net and its sister networks in the regions created a demand for services that a group of busy academic volunteers simply couldn't meet. In 1997, CA*net Networking Inc. ceased operating as an independent entity and passed the torch for that aspect of its mission to the private sector, while other aspects were picked up by a range of not-for-profit organizations, including CANARIE.

THE PIONEERS SPEAK

Part of CA*net's legacy is being carried out by the CA*net Institute. The Institute was created to raise funds and distribute them to worthy research projects, including this history. It decided in early 1999 to retain Omnia

Communications Inc. to research and write the story of how computer networks were built in Canada and gradually became part of the global Internet. Two fundamental approaches were agreed on from the outset.

First, the narrative was to be couched in language that, while still somewhat technical, would be accessible to the largest possible audience, and not merely to specialists. The rationale for undertaking the project was to provide a service in the public interest – to tell all interested Canadians about an exciting and ground-breaking part of their history.

Second, while some information came inevitably from archival and Web-based documentation, the majority of the story emerged from a long series of consultations and in-depth interviews with the principal pioneers and builders themselves. This, in other words, is the story as told by those who worked on the front lines – finding the resources, volunteering personal time, developing the applications and workarounds, staying in touch with the Americans and others working in the field, right up to the point at which Canada finally had a national network backbone to call its own. The extensive interview transcripts have been quoted verbatim in many places, to give the narrative a direct and personal quality.

While the work reflected here is comprehensive and a first of its kind, it had to be undertaken with realistic goals. In many passages, brief explanations are provided of larger historical developments, as well as of technical concepts such as packet-switching and the TCP/IP suite that underlies the Internet's infrastructure. Nevertheless, limited space and resources prevented the writers from exploring all aspects of the development of the Internet in Canada, particularly the growth of Internet service providers (ISPs) and other commercial interests that worked in parallel, sometimes in cooperation, with the academic networking community.

As noted above, this part of the story of networking in Canada ends at the point where CA*net, Canada's first national, IP-based production network, ceased to operate and left the field open to the private sector. That turning point marked the end of the golden years of regional experimentation and the struggle to put computer networking on the national stage. It remains for others at some future date to bring the story up to the 21st century by describing the growth of institutions like CANARIE and explaining how the Internet moved from academic obscurity to the centre of Canada's economic and social life.

1

UNCERTAIN BEGINNINGS

"THERE IS A GREAT NEED FOR A COMPUTER COMMUNICATIONS NETWORK ABLE TO SELECTIVELY CONNECT ANY USER TO ANY SERVICE, QUICKLY, RELIABLY AND CHEAPLY. WE HAVE CALLED SUCH A NETWORK THE TRANS-CANADA COMPUTER COMMUNICATIONS NETWORK."

~ SCIENCE COUNCIL OF CANADA, AUGUST 1971

In 1970 the Université du Québec hired systems engineer Joseph Reid to design and implement an inter-campus computer network. At that time, such a task was daunting enough in itself. But within months Reid found himself co-leader of an attempt to build not just a provincial network but Canada's first national, packet-switched data network. In other words, to build the backbone for an Internet in Canada before the word was even invented.

Reid faced a truly formidable challenge in commandeering scarce resources and overcoming Canada's geographic obstacles. Looking back on it, Reid can't help identifying with the "mad English engineer" and POW in The Bridge on the River Kwai who doggedly persisted in building his bridge joining one side of the river to the other – only to have it blown up before his eyes.

Although Canada's universities were far from POW camps, they did present the computing community with harsh realities. When Reid started on his quest, desktop computers didn't exist – they were barely even dreamed of. The most common device at the time, the mainframe computer, had to be housed in its own spotlessly clean room and tended to by lab-coated technicians. In the academic world, even the largest universities could not populate their campuses with these huge machines typically costing as much as $3 million. Normally a university would invest in a single mainframe to form the core of a computer services department providing both administrative and academic services for the entire university.

The specialized needs of computer science departments, however, led them to embrace the minicomputer – which was still as big as two full-sized refrigerators. Cheaper than a mainframe, computers like DEC's PDP series provided the still-fledgling discipline with resources that operated outside the school's centralized computer infrastructure. Other university departments that needed computers for their own special applications but could not afford their own hardware had to rely on the university's centralized computer services.

When Reid started to work, computer networking of the kind envisaged by the Science Council of Canada (SCC) was in its infancy. For the managers of most computer services departments, the basic problems of how to get multiple users to access one mainframe were still in the process of being resolved. Multiple monitors and keyboards connected to a centrally located time-sharing mainframe were state of the art. The technology had evolved so that one mainframe could accommodate hundreds of people using these keyboard-based dumb terminals. Soon it became possible for these users to interact with each other through the centralized mainframe, giving researchers a taste of real computer networking – including email.

PACKET-SWITCHING: A REVELATION

Reid was familiar with the dumb terminal systems developing rapidly on many campuses. Quebec's ministry of education had already implemented an elaborate and expensive remote terminal network, leasing dedicated phone lines to connect terminals at elementary and secondary school boards to the main computing centre at ministry headquarters. A parallel network linked each CÉGEP (similar to a community college) in the province to the ministry as well.

These networks were used primarily for administrative tasks, not teaching. One drawback of their star-network design was the expense involved in connecting every point on the network back to the central computer. Reid faced the same problem: because the Université du Québec served the entire province, its campuses were small and widely dispersed across a huge geographical area.

Computer networks in use at the time were limited to a particular set of technologies and functions that typically served a small, single community. Although dumb terminal systems were a possible albeit limited solution, Reid was looking ahead and searching for a more sophisticated technology that would provide true networking capability for the university's various campuses.

Only a few months after taking on the job, Reid found his solution – or so he thought. In October 1970, he attended a conference organized by the Inter-university Communications Council. This council was founded in 1964 for the purpose of encouraging universities to share computer resources and expertise, and to build actual computer links between universities. Reid hoped to find some networking insight and inspiration from the other attendees and speakers. The solution he stumbled on exceeded his most optimistic expectations.

A computer scientist from the Advanced Research Projects Agency (ARPA), a research arm of the US Department of Defense, delivered one of the presentations. Larry Roberts was there to disclose the recent accomplishments of his group, which was successfully implementing a networking technology he referred to as "packet-switching." The technology was currently being tested in a network called the ARPAnet.

Although the ARPAnet began life as the world's first large-scale implementation of a packet-switched network, the direct motivation behind the ARPAnet project had nothing to do with packet-switching. The goal was at once more basic and more complex: to get different kinds of computers talking to each other on a single network.

Computer science was still a young discipline by the mid-1960s, yet already fragmented, since the computers that engineers designed to tackle specialized tasks used unique and proprietary hardware and software. The balkanization of computer science was slowing progress, creating a great deal of duplication of effort and wasting resources. By 1966 the head of ARPA's Information Processing Techniques Office, Bob Taylor, had no less than three computer terminals in his office connecting him to three different remote computing facilities. The collection of terminals was so obtrusive that they were housed in their own room.

Although the ARPAnet was always intended to serve as a common interface between all these machines, the decision to design the network as a packet-switched one was made later in 1966 when the ARPAnet team stumbled onto the work of Paul Baran, a US computer scientist, and a UK research scientist named Donald Davies. Their strikingly similar work on packet-based computer communications convinced ARPA's designers that packet-switched networking provided the most flexible, durable and efficient model for the ARPAnet.

Roberts explained the significance of packet-switching and the ARPAnet by pointing out that standard analog telephone networks, which were "circuit-switched," were too slow, crude and inefficient to meet the requirements of computer network communications. In a telephone network, sound is transformed into an electrical signal carried across a temporary circuit established by mechanical switching devices within the network (hence the term “circuit-switched”). In 1970, it took 20 to 30 seconds to connect two mainframe computers over a circuit-switched network – whereas the optimal connect time for computers exchanging data was well under a second. Circuit-switched analog transmissions were also far too riddled with errors to meet adequate data transmission standards. Finally, the “bursty” nature of data transmissions, as opposed to the continuous use made of a telephone line during a typical human-to-human conversation, made for highly inefficient use of a telephone network’s capacity.

The most revolutionary aspect of the ARPAnet lay in the mechanics of packet-switching, which Bob Kahn, one of the ARPAnet’s inventors, compared to distributing electricity by hand-delivering each electron. Unlike a circuit-switched network, a packet-switched network breaks down data (for example, a text message, an image or a sound) into individual data packets, each one identified and addressed with its final destination. Distributed through the network are specialized computers that read the packet addresses and send them towards their destination. Once arrived, the destination computer reassembles and then processes the data. In 1970, short text messages could be sent over the ARPAnet as packets and received on the other side of the country in under 0.2 seconds.

The way computer networks were currently set up, Roberts concluded, didn’t give even a hint of how they might be used. That’s where the ARPAnet came in. It was built from the ground up as a computer-centric network. In this sense, it’s fair to say that the ARPAnet was the first real computer network, rather than simply a collection of networked computers. Though still in the testing phase, the network had proven itself to be much faster and more reliable than conventional circuit-switched networks.

Another advantage of the ARPAnet was that it wasn’t an application-specific network. The only job ARPAnet was expected to perform was to distribute data of whatever kind to its final destination. Thus it served as an all-purpose interface between people and any computer-based resource or application made available on any computer on the network.

Sitting in the audience Reid felt lightheaded as he absorbed what Roberts was saying. He knew that ARPA’s achievement was significant, though he couldn’t begin to imagine the full impact packet-switching would have on a host of existing and future applications. But he was certain of one thing – it would be a real coup to use packet-switching to run the Université du Québec’s computer network.

Not surprisingly, Reid wanted his university to join the ARPAnet in the very near future. In fact, as word of the ARPAnet spread, the question "When do we get to join?" was on everyone's lips. Some proponents talked about the possibility of the ARPAnet being commercialized and offering access to any organization willing to meet the cost of joining.

But the managers of the ARPAnet had different views on the matter. Reid pursued his hopes of joining the ARPAnet directly with Larry Roberts himself at ARPA headquarters. Despite the fact that US universities working on contract for the Department of Defense were being added to the network at a steady rate, Reid recalls that "Roberts wouldn't play ball because at the time [the ARPAnet] was a network of US defence establishments and letting a foreign tap into it wasn't acceptable." Quickly changing gears, Reid began investigating the possibility of building a similar network in Canada instead.

A HOME-MADE SOLUTION

Whatever ARPAnet's own policies were about letting Canada in, not everyone at Canada's universities was ready to embrace the idea. Among them was Leon Katz, a professor of engineering and physics at the University of Saskatchewan and director of its nuclear research lab. Katz was dissatisfied with the notion of Canadian organizations depending on a network entirely owned and operated by American interests, especially the US government.

Katz was part of a growing community who believed that computer-mediated communications would one day reach into Canadian homes and become as important as traditional phone and television networks. In Katz's mind, if the ARPAnet, or its successor, became Canada's incumbent network provider, the nation might eventually forfeit much of the sovereignty it had gained through the construction of its own transportation and communication networks.

Katz was certainly in a position to make his views known. As chairman of the committee on computer applications and technology at the Science Council of Canada (SCC), he oversaw publication in August 1971 of a report titled "A Trans-Canada Computer Communications Network." The document describes a digital backbone traversing the country, connecting independently owned and operated regional and local networks. In many ways the report is both remarkably prescient and a projection of the rigid status quo in the early 1970s.

It describes the construction of a national network with no pre-ordained purpose – in sharp contrast to telephone and cable networks. It would provide whatever services an organization (public or private) imagined its users might need or want. This open-handedness was a radical notion at a time when no one was allowed to attach any third-party device to the established telephone networks, on the grounds that their structural integrity would be put at risk (an assumption later judged to have more to do with protecting the financial integrity of the telephone companies, rather than the technical integrity of their networks).

On the other hand, the Council balanced its entrepreneurial vision with the idea that the owner-operators of this network would be highly regulated, required to provide universal service and barred from providing any of their own services directly to end-users. The report ended on this note:

> In summary, we have a long tradition of creating nationwide links for national survival and unity. In the past we have created the Canadian Pacific Railway, the Canadian Broadcasting Corporation and Air Canada. Some of these links carry goods, others carry information, ideas and values. Each required a sustained national effort and substantial leadership by the Government of Canada, often in the form of a crown corporation. Each has been a major factor in building national unity and in balancing the north-south pulls of continentalism with a heavy concern for national goals, and at the same time has built up a strong and valuable Canadian industry. The Trans-Canada Computer Communications Network [TCCN] is in fact the newest potential member of these national links.

Despite the perceived urgency, the report insisted that actual network construction should not begin until a full five years of studies, forecasts, trials and tests had been completed. With classic bureaucratic reserve, the report concluded:

> Immediate action is required. However, "immediate action" does not mean "start building a [network]."

ENTER THE FEDS

While Katz and the SCC were preparing their report, Reid spent the winter and spring of 1971 mobilizing forces for a new network. The University of Waterloo then as now was one of the leading computer science universities in Canada (these days having proven to be a hotbed for digital start-up companies like Research in Motion and Wi-Lan) – and made an obvious choice as partner. Waterloo agreed and together they prepared what Reid calls a "modest proposal" submitted to the federal Department of Communications (DoC). Essentially, Waterloo and Quebec asked for enough money to buy several PDP-11s and lease dedicated phone lines from the telephone company to connect them.

The 1971 proposal met with enthusiastic acceptance – perhaps, in retrospect, overly enthusiastic. John Demercado at the DoC liked the idea so much he didn't want to see a mere replica of the ARPAnet built in Canada. He wanted to use the ARPAnet technology as the starting point and then use the TCCN concept as the catalyst for developing an even more advanced packet-switched network. It would be much larger than Reid's proposal and built on the scale of a TCCN, a digital backbone running from one end of the country to the other.

Assigned to lead the project at the DoC were Terry Shepard and Doug Parkhill. Along with Reid, they were in charge of contacting prospective participants and orchestrating the development of a network proposal for their superiors at the DoC. They decided to keep the project within the university community since Canada's university computer centres had much to recommend them as potential partners.

Universities shared an obvious need to pool their computer resources more efficiently. They were non-profit organizations that wouldn't be looking for an immediate return on whatever investment they were asked to make. The universities were very much in the thick of computer and networking innovation and since there were at least one or two major universities in every province, geographic representation wasn't a problem.

Reid also thought it would be a good idea to bring Bell Canada into the project. After all, Bell Northern Research (BNR) was one of the largest communications research organizations in North America. They certainly had the resources necessary to help design the network, and Bell would have to supply the network's phone lines in any case. But Demercado quashed the idea, wanting to keep the project between the educational community and government – specifically the DoC. According to Reid, Demercado felt that Bell was a little too well resourced to bring in as just a partner. If Bell eventually wanted to use packet-switching technology commercially, bringing the telco in at this stage could generate a conflict of interest down the line.

Since the project was explicitly centred on the universities, it was dubbed the Canadian University Network, or CANUnet. As Reid flew across Canada from Quebec to Victoria drumming up interest in the network, he stumbled on existing regional networking initiatives in Saskatchewan and Ontario. While these projects were encouraging, they added an element of competition and territoriality to the discussions. Some provinces had already invested considerable time and money in the planning and development of networks that would probably not be compatible with the proposed national backbone. Nevertheless, all of the universities were willing to help investigate the possibilities, and DoC's donation of $35,000 to help with the investigation didn't hurt.

SELLING BELL ON PACKETS

During this phase of the project, Reid had another opportunity to try and bring Bell into the loop but from the opposite direction. At a BNR symposium he found himself debating the merits of packet-switching with Dave Horton, a senior VP at Bell Canada. Horton claimed that Bell didn't have to worry about Reid's packet-switching ideas because the carrier had already implemented a national data service called Dataroute for Canada's banks and other large corporations.

That's fine for them, returned Reid, who went on to paint the scenario of a librarian at the Université du Québec needing to search resources located at several different locations, each using a different application. Could Dataroute easily and economically handle that kind of use? This gave Horton pause.

Soon afterwards, Walter Inkster, BNR's VP of computer communications engineering, contacted Reid saying he wasn't sure what Reid had said to Horton but Horton had just told him to make packet-switching research a priority. Excited, Reid suggested that Parkhill from the DoC be involved with BNR's research, but this time the idea was quashed by Bell. Some years later Bell released Datapac, a packet-switched service delivered over the public network, which became widely available throughout Canada (though hardly commonplace outside of large organizations). Ironically, Inkster later credited the post-symposium chat between Reid and Horton as the origin of the service.

During this same period, Roberts re-contacted Reid to say yes, the Université du Québec would likely be welcome to join ARPAnet. But by this time simply linking Canadian universities with the American network was not in the cards. Reid and his partners at the DoC focused instead on the CANUnet proposal.

The Science Council's report came out the same month that Reid's new CANUnet group met for the first time – August 11, 1971 in Quebec City. Leon Katz was in attendance, as were the various computer centre directors Reid had met on his travels. Parkhill chaired the meeting and the universities shared their regional experiences to date in creating computer networks. Each university committed to preparing a letter of support for the project in which they would outline the value of CANUnet to their institution, as well as the people and resources they would provide for the project. They then broke down the network proposal into discrete issues that could be dealt with by smaller expert groups. They settled on four: network design (topology), utilization, communication protocols and the institutional framework.

Although Reid eventually authored a strong report summarizing the value of a network like CANUnet, the proposal as a whole had weaknesses. The four expert reports were bluntly honest on many issues: it was clear that a lot of questions about how to build a network couldn't be answered for certain until a network had actually been built. Moreover, no one could predict with any accuracy what levels of traffic were likely to be encountered. And because each of the four research groups worked independently on what were interdependent elements of the network, each report had to consider the network's topology, protocols, usage and even its administrative structure from every conceivable angle. The proposal as a whole ended up giving a very mixed impression of what kind of network CANUnet might be.

The letters of support received also gave an uneven impression of the willingness and ability of the participants to contribute to the project. Reid was keenly aware that there was some ego at play in these reactions. McGill University, for example, went out of its way to emphasize that it really had no need of a network because of its existing wealth of computer resources. The Ontario universities didn't want to step on the toes of the Office of Computer Coordination, which was organizing the effort to construct an Ontario computer network. And for their part, the smaller universities were candid about their limited resources and expertise in computing – although they did also point out that a lack of resources was the main reason for their wanting to see a network developed in the first place.

THE NETWORK THAT WASN'T

Reid's report, along with the supporting letters and the four research reports, represented over a year's work and close to $75,000 invested. The final result demonstrated clearly that the participants were willing to put in the effort required for a project of this scope – especially given the limited amount of time available and the scope of the problems each group had to tackle. Despite the technical precedents set by ARPAnet, the possible approaches to computer networking were still many and varied.

Since some uncertainties were to be expected at this early stage of the project, Reid, Shepard and Parkhill remained optimistic. This was simply a proposal after all, and Reid's document laid out a solid plan for building the network. No one expected CANUnet to be built overnight. The SCC had anticipated years of planning before construction was to begin – though certainly not as many years as it eventually did take.

Reid submitted his report in March 1972 and the DoC rejected the proposal. Two months later the department issued its own exhaustive two-volume study of computer networking in Canada called "Branching Out." Amazingly enough, CANUnet and Katz's report weren't even mentioned, nor was the possibility of the government creating or funding a computer network like CANUnet at all discussed. Policy leanings had shifted, apparently prompted by the assumption that Canada's existing network providers, i.e. the telephone companies, would develop the required infrastructure and technology.

Reid was disappointed but felt that CANUnet was in any case more ambitious than he'd originally envisaged: "It was the DoC that conceived of it as a grand project, whereas I just wanted to get a couple of universities connected and find out how the technology would work." He turned back to the business of building the Université du Québec's computer infrastructure and, by the end of 1972, he had completed its inter-campus network with a series of terminal connections to the university's various computers. The network ended up being less sophisticated than Reid had hoped, but at least the connections were made.

The DoC's rejection spelled the end of the CANUnet project. No funding was forthcoming in any form for a network on any scale. It seemed no one at DoC except Shepard and Parkhill had the slightest inkling of how computer networks would grow in the coming years. Whatever the reason, the window of opportunity for computer networking on a national scale had closed and would not open again for more than a decade.

Although the CANUnet project may have been premature, it laid out the territory – and the issues – that determined debate about national computer networks for the next 25 years. These included the huge gravitational pull exerted by US technology, resources, and economies of scale that continually tempted Canadian universities, provinces and business interests to look south for solutions; the varying interests of computer users, from computer administrative services to computer science departments to mathematics and science departments with specialized research needs; the struggle to establish universal standards in the face of competing proprietary hardware and software; and the need for those Canadians who believed that a Canadian network would not only preserve Canadian sovereignty but also best serve the needs of all Canadians to act cooperatively and tenaciously in pursuit of their ultimate goals.

GRASS ROOTS PROGRESS IN ACADEME: NEW BRUNSWICK & PEI

During and after the abortive CANUnet project, the construction of remote terminal connections on university campuses continued unabated. Ironically, one of the first and largest of these networks was developed in a province Joseph Reid and other CANUnet proponents didn't even consider visiting in 1971. The main reason New Brunswick was left out of the initial CANUnet planning was also a big part of why New Brunswick's computer network was so successful: the schools simply didn't have all that many computers to go around.

A unique balance of power existed among New Brunswick's universities that created a stable and relatively cooperative environment. As David Macneil, currently director of the computer services department at the University of New Brunswick, explains:

> The province has a number of small universities and one big one. New Brunswick has an advantage over a lot of other provinces because the big one is about the same size as all the small ones put together. Things end up being reasonably balanced. We can make progress because we don't get stalemated and UNB isn't threatened by these other institutions because, individually, they are much smaller.

Initially the province's only major computer resources for its universities were located at UNB in Fredericton. Dana Wasson, UNB's head of computer science at the time, was concerned that the other universities might never be able to afford similar resources. Wasson felt that the province's educational computing resources could only expand by pooling their collective assets and sharing their existing hardware.

Wasson and members of his department approached other computer service departments at Mount Allison, Moncton and St. Thomas, and sketched out preliminary plans to build a provincial computer network. The New Brunswick Higher Education Commission, a new agency created to promote inter-university cooperation, agreed that the networking project fit their mandate perfectly. The agency committed the necessary money to back the newly-christened New Brunswick Educational Computer Network.

Wasson and his colleagues started building the network using dedicated lines, i.e. they leased copper lines owned by the telephone company for their own use. The lines were connected by a 2400-baud modem the size of a small kitchen stove. As a point of comparison, the 56 Kbps dialup modems inside most of today's run-of-the-mill PCs send data 20 times faster.

Although he was too junior a member of UNB's computer services department in 1970 to play a major role in the first iteration of the network, Macneil recalls how it changed the way computing was taught in the province:

> The "resources" being shared were batch cycles on a mainframe. They had a punch card reader in Sackville, for instance, and they were able to teach programming by sending a computer code (called a "job") up to Fredericton on a wire, where it was processed by the computer and sent back to Sackville, where the results were printed out. As far as the students were concerned, the computer could be in the next room – they couldn't tell the difference.

> And we did this for each of the small universities, even places like the University of Moncton's two branch campuses, Shippegan and Edmundston.

In 1976, the universities in Prince Edward Island joined the network, prompting the network's name change to the New Brunswick/Prince Edward Island Educational Computer Network (NB/PEI ECN), which it remains today. New applications and networking technologies were added over the years. In 1980 the ECN made email facilities available to its faculty and students, and the use of the network as a communication medium became highly valued.

When the ECN was first envisaged, it was assumed that the "resource" which the computer networks would share was the raw computational power of a mainframe computer. But this definition shifted even within the first few years as local support staffs at the various campuses were themselves supported by the computer experts in Fredericton. Human networks of expertise, every bit as important as the technologies themselves, expanded alongside the computer networks. As Macneil explains:

> We started sharing hardware and discovered in fact that what we were sharing was a lot of knowledge and information. Every six weeks or so the directors of computing in all the universities of New Brunswick and PEI would meet and we'd exchange a lot of stories, information, programs and documentation. We helped each other with user services and support.

Furthermore, the academic community was beginning to appreciate how much the ability of computer networks to share human and intellectual resources greatly reduced the importance of being physically close to major research and academic institutions. As Macneil points out, the network's introduction of email facilities, even on a small scale, "made New Brunswick and UNB, this remote institution in the north woods, intellectually a more central place."

Overall, the original applications, structure and funding patterns of this pioneering network established a number of important precedents for the networking story that would eventually unfold on the national stage.

2

NETWORKING TAKES ROOT

“I DIDN’T HAVE A LAB, BUT I HAD SOME MONEY AND A LIST OF SENSIBLE THINGS THAT HAD TO BE DONE. AND BUILDING A NETWORK WAS A SENSIBLE THING TO DO.”
– VINCENT TAYLOR

Canada’s first attempt at a national network failed to take root but the seeds had been sown. Parallel to the increasing use of computers and computer networks in academe and research, the technology had become an essential administrative tool in most large corporations by the mid-1970s. The challenge of meeting the growing need for a network that would allow these computers to talk to each other was taken up in the private sector mainly by computer vendors and telephone companies. And the technology that appeared to be the most promising was packet-switching, which was by now well established on the ARPAnet and other international military research networks.

In 1978, Bell Canada began offering its Datapac service using a packet-switched data network designed to take advantage of the benefits of packet-switching while preserving the traditional time- and distance-sensitive pricing scheme telephone companies used for their circuit-switched networks. Datapac was one of the earliest packet-based services used commercially by a telephone company anywhere in the world.

Datapac also became an important model for the X.25 technical standard, whose development was fostered by the Comité Consultatif International Téléphonique et Télégraphique (CCITT). For more than a century, this organization had served as a forum where national telecom players – typically large regulated monopolies, islands unto themselves on the domestic front – had come together to develop and settle on technical standards.

QUEBEC & EDUPAC

Universities were quick to consider the potential of Datapac and the X.25 standard. In the forefront again was the Université du Québec (UQ), where

Joseph Reid had established a network of remote terminals in 1972, all linked to a centralized Control Data Cyber 6000 mainframe located in Quebec City. A few years later the distant campuses upgraded their dumb terminals for Control Data or DEC computers of their own. However, because of their remote job entry (RJE) networking protocols, these distant computers still had to emulate remote terminals. In other words, they could only communicate with and through the central mainframe in a master-slave hierarchy, tapping into applications located at the central computer as opposed to sharing applications they ran themselves.

UQ continued to explore ways in which networking could help overcome their distributed campus structure. In 1975-76 it started an ambitious – and expensive – videoconferencing initiative using satellites to connect classrooms across the province. In 1977 it implemented Plato, a computer-assisted teaching program using highly specialized terminals from Control Data. The following year UQ investigated the advantages of networking with the new X.25 packet-switching technology. One possibility was subscribing to Bell's existing Datapac service. The advantage to this option was less up-front cost because Bell brought its own infrastructure along with the service, leaving only operating costs based on the volume and distance of network traffic.

On the other hand, UQ's existing network already covered a fair geographical distance, linking Montreal, Quebec City, Chicoutimi, Trois-Rivières, Hull and Rouyn. And since the network was used for internal administrative business as well as for teaching purposes, it carried a high volume of traffic. Moving to the X.25 protocol promised to increase the number of possible applications, and thus the flow of traffic. Taking all these factors into account, UQ's cost estimates showed that it would be cheaper to buy their own hardware (nodes and concentrators), lease their own copper lines and run their own network – just as the university had built and operated its own network of remote terminals.

Richard Lacroix, director of the computer centre at UQ's Montreal campus, his counterpart Guy Bertrand in Quebec City and several others jointly issued an RFI to design and implement the network. A number of companies responded, including Bell Canada and other large information technology firms. But the winning bid eventually came from a small Montreal company called Société de Télématique Rochon, Thibault, Chevalier. The X.25 network which the firm implemented replaced the university's more limited network of remote terminals and remained in use for more than a decade.

In 1981 Lacroix left UQ for a position with Quebec's ministry of education. The ministry wanted to improve the network of remote terminals which

Quebec's school boards had depended on since 1969. Lacroix's experience helping to implement the UQ network served him well in this new role. Since the main expense of the network was in leasing the hundreds of kilometers of copper lines that would connect the scores of school boards to the ministry's computer centre, the distributed peer-to-peer model of networking possible with X.25 would reduce costs significantly. The boards would only have to lease enough copper line to reach the nearest network node rather than all the way back to the ministry's headquarters. Costs could be further reduced by leasing access to the network to other public organizations.

This new network was dubbed Edupac, a contraction of "educational" and "packet-switched." By 1983 Edupac included all of Quebec's elementary and secondary school boards, the province's CÉGEPs and a host of other public organizations, altogether totaling more than 200. The network also established its own email service, putting all the institutions in electronic touch with one another.

Lacroix's only disappointment was that the network was used mostly for administrative purposes and never became popular with teachers. While some teachers used the network to teach students about computers and computer languages such as APL and LOGO, most were too daunted to take real advantage of what was easily one of Canada's most technologically advanced networks at the time.

THE WEST WALKS ALONE

Quebec's experience with Control Data hardware and the X.25 protocol was unusual for the time. No matter what their network configurations, almost all grass-roots efforts at universities such as UNB relied on IBM mainframes and proprietary IBM software. The computer services department at the University of British Columbia (UBC) in Vancouver also used an IBM mainframe, but the software route the university traveled led them to a different conclusion – and to Canada's first national networking initiative.

UBC: COMPUTER SCIENCE GETS INTO THE ACT

In 1969, UBC was the first university in Canada to invest in a new IBM mainframe called the System 360/67. It was designed to make use of features never before available in the commercial marketplace: dynamic address translation and virtual memory (two advanced technologies that effectively expanded the memory and computing power of the mainframe system).

UBC's IBM 360/67 used the Michigan Terminal System (MTS), an operating system specially designed at the University of Michigan for the 360/67. Nine universities – from western Canada (including Simon Fraser University in Burnaby, B.C. and the University of Alberta), the United States and overseas – formed the MTS group. It was a world unto itself, meeting every six months at working sessions attended by up to 100 people. These face-to-face working sessions naturally migrated to a networked approach as X.25-based services became available in the late 1970s. The Canadian universities and their international partners used the links to communicate with each other, and to share and develop software such as email services for the MTS systems.

The IBM mainframes running MTS at UBC might have served the needs of the computer services department for years, but they presented major drawbacks as a tool for teaching computer science – as Paul Gilmore discovered when he became the head of the department at UBC in 1977. Gilmore had been working in New York at the T.J. Watson Research Center, the headquarters of IBM's Research Division. Offered the opportunity to return to his home town and alma mater, he jumped at the chance. As he puts it, "Vancouverites, like salmon, all have to head home sooner or later."

Other than the homing instinct, there wasn't a great deal about the position to draw him away from New York. At the time, UBC's computer science department was part of the Faculty of Science, and not a big part:

> The department had 13 members and not a very high profile in the university. The physical accommodation was very poor, we were scattered over various places, so the first challenges were really physical ones – trying to get the department into one place so you could begin to really build some departmental spirit.

Another challenge was improving their computer equipment, which was a far cry from the high-end hardware Gilmore had grown used to in New York. As Gilmore explains, "Teaching computer science on MTS was not the ideal way to go! We were having to teach our kids using punch-card equipment, which was pretty unacceptable in 1977."

The UBC computer services department, recognizing that computer science was a distinct computing culture with special research and teaching needs, helped Gilmore get his terminals. But his department also favored ever-more personal computing, beginning with mini-computers such as the PDP-11, typically running the UNIX operating system.

The UNIX OS, developed by computer scientists working out of Bell Laboratories, was written in C, a high-level programming language that made it compatible with almost any kind of computer. At the time, Bell wasn't allowed to make a profit on activities outside of its core telephony business, so it couldn't sell UNIX commercially. Instead, the source code was licensed for a nominal fee to universities, and versions of UNIX began springing up all over North America (UNIX has achieved more popular notoriety today in the variant known as Linux).

As an OS designed by computer scientists for computer scientists, UNIX was an excellent teaching tool as well as an effective operating system. Its success meant the research action was happening on smaller and smaller computers. And it was the development and spread of increasingly powerful and accessible computing that was transforming the science of computing – all leading to a great demand in the private sector for graduates with the skills and training to program and operate computers.

Students flooded into UBC's tiny computer science department in the wake of Gilmore's arrival. He didn't have the resources to accommodate them all, nor was it possible to hire more faculty:

> There weren't the Ph.D. candidates available and those that were available were not terribly interested in coming to a tiny place out here on the West Coast. So it was suggested to me that Gerald Neufeld would make a good faculty member. He was exceptionally bright, very knowledgeable and had worked in the computing centre, in the communications group and also in our artificial intelligence group. Everyone praised him highly.

Unfortunately one thing Neufeld hadn't done at university yet was earn his Ph.D. Gilmore nevertheless convinced the dean to hire Neufeld as a lecturer.

CDNNET: A TESTING GROUND FOR OSI

Gilmore was also excited about bringing Neufeld into the department because Neufeld brought with him contacts he had made at the Comité Consultatif International Téléphonique et Télégraphique (CCITT). While at UBC Neufeld was planning to work with the CCITT to develop the standard for an email messaging protocol known as X.400.

The X.400 standard was one of several being developed by CCITT subcommittees under the auspices of a second international and venerable standards-setting body, the International Organization for Standardization (ISO). This body had also initiated work to define a common architecture for computer communications and in 1978 formally began its effort to design what it called the Open Systems Interconnect (OSI) standards. By 1979, drawing on work done by computer vendors, national standards organizations and the ARPAnet itself, the ISO set out its blueprint for the ambitious OSI standards project called the "seven-layer model."

This model specified the range of protocols underpinning a data network's architecture. Each layer defines a separate set of tasks that must be carried out to network computers successfully and allow them to share data. The first is the physical layer (e.g. copper wire, optical fibre) and the seventh or top layer is the application layer. All the layers in between are reserved for the protocols that establish connections, direct data, monitor the network and perform other behind-the-scenes tasks.

The point of defining seven discrete layers of network protocols was that each can be designed to work independently from the other. As long as the protocol is written in conformance with the standard, then it should be able to work perfectly well with any other existing layer or network function. OSI would allow data networks of all kinds to inter-communicate and was to be implemented on a global basis.

The OSI standard Neufeld was developing for the X.400 protocol would function as a top-layer application. The plan was that Gilmore would arrange funding from the Natural Science and Engineering Research Council (NSERC) to allow Neufeld to pursue his X.400 research and earn his Ph.D. concurrently. The work on the messaging standard would parallel the development of a test network. Gilmore explains:

> The network was partly a matter of necessity and partly political. The necessity was that you can't properly test the software unless you have it running on platforms in different places by different people. And the political motivation was that we had no money: the university didn't have money for us, the computing centre didn't have money for us. The only pot of money available was from NSERC, so we had to create something that would be useful to Canada's scientific community.

The network, having received federal funding, was given a suitably "national" moniker: CDNnet. But unlike CANUnet, CDNnet wasn't built to be a production network upon which applications would be built. Rather, CDNnet was more like an application for which a network had been built. CDNnet wasn't defined by its wires and routers, but by the email application that the network members would share and operate together.

The new CDNnet project received over $170,000 dollars from NSERC in 1981 when Neufeld was first hired, and the funding lasted until 1984 when he left temporarily for Waterloo to complete his Ph.D. The project organizers announced early on their intention to develop a messaging software. According to John Demco, a friend and co-worker on the project, Neufeld searched long and hard for an appropriate name. He looked up words for "messenger" in West Coast aboriginal languages and found a reference to an appropriate great spirit, but the person he talked to about it wouldn't actually utter the name since this might bring bad luck (eventually Neufeld settled on EAN, which stands for Ellen and Andrew Neufeld, his wife and first child). During 1982 he incorporated the OSI messaging standards into the EAN software and tested them in partnership with the CCITT sub-committee.

Meanwhile, Gilmore was arranging the CDNnet to implement and test the EAN software. He found that, in many cases, the building blocks were ready to hand and only had to be cemented into place. First of all, the EAN software had been designed to work equally well with different physical and transport protocols. As a result the network organizers didn't have to specify a certain kind of network infrastructure. Second, CDNnet opted to use existing facilities wherever possible in order to enable remote sites to join the network as well as reduce the complexity of managing the network.

Most of the participants had available to them the public data networks built by the telephone companies and operating on the CCITT-approved X.25 protocol. These networks opened up a wide audience to CDNnet. Moreover, there was very little up-front financial commitment for participants since Datapac, Canada's X.25 network, was a pay-as-you-go proposition. Network members only paid for packets of data actually sent and received. The down-side was that, as the network became more popular and useful, it got more expensive to use.

At the end of 1982, the CCITT announced it had completed the technical standards for the X.400 messaging protocol. A mere two months later, UBC announced the first successful implementation of those standards in the form of EAN software connecting a number of remote sites across Canada. Because it was the first working implementation of the standards, there was a great deal of international interest in the software. In Canada, computer science departments were also intrigued by CDNnet itself, which was growing rapidly.

A great deal of the coordination and management of the network was soon handed over to Demco, who had come to UBC to pursue his own Ph.D. in computer science. Demco became the manager of CDNnet in 1984 and oversaw its growth and development. Although the network was designed to operate as a distributed network without a central node, it in fact developed with UBC as its core. Messages were sent directly to Vancouver then relayed to their final destination.

Canada's first national network was up and running – and was soon followed by another with a direct connection to the ARPAnet itself.

ARPANET ON A BUDGET: DRENET

Use of the ARPAnet gained momentum throughout the 1970s. Starting with four operational nodes in December 1969, the network grew to almost 100 nodes by 1980. But the ARPAnet remained in effect a closed military facility. University, government and private industry research facilities without defence contracts were shut out. Even with a military contract, the cost of leasing permanent and redundant links to the ARPAnet priced network access out of the reach of would-be members.

By the early 1970s the ARPAnet was the largest, most advanced packet-switched network on the planet. But it was not yet an "internet," let alone the Internet. It was a single network without gateways to other networks: ARPAnet traffic stayed there not just because of security and cost restrictions but because there simply weren't any major domestic computer networks to interconnect with. Most academic, government, corporate and hobbyist networking initiatives wouldn't begin until the late 1970s and early 1980s.

INTER-NETWORKING

Internationally it was a different story. In 1972 the International Networking Working Group was created as a forum for research and discussion on various networking issues, especially inter-networking. The group's goal was to create an international network of networks they dubbed Catenet – from "concatenated network," i.e. a series of linked networks.

The Catenet was a lofty ambition but the fact was that, in the early 1970s, the protocols required to allow differently configured data networks to interact were not in place. Vint Cerf and Bob Kahn, the ARPAnet's resident experts on packet-switching, overcame that obstacle in 1974 with the publication of a milestone paper entitled "A Protocol for Packet Network Interconnection." Drawing on lessons learned from the ARPAnet and their international colleagues, especially Louis Pouzin's work on the Cyclades project then under construction in France, Cerf and Kahn outlined a Transfer Control Protocol that would allow different data networks to exchange traffic. It was the invention of this gateway protocol that allowed the Internet to flourish, and made Cerf and Kahn the "fathers" of the Internet they are considered today.

Meanwhile, national networking projects in other countries were beginning to bear fruit. In addition to Cyclades, a full-scale packet-switched network, Donald Davies' work in the UK was progressing at the National Physical Laboratory. And while the Catenet itself was still only a dream, ARPA and several members of the international military research community did establish a packet-switched network that operated over satellite links. This Atlantic Packet Satellite Network, or SATnet, connected NATO members in the US, UK, Norway, Italy and Germany. From this group was founded the International Collaboration Board, which continues to meet today.

Finally, the first successful test of inter-networking protocols and technology occurred in October 1977 when packetized messages were passed back and forth across the ARPAnet, SATnet, and a terrestrial wireless network in San Francisco. A year later the protocols used in this experiment were further refined into the Transmission Control Protocol and Internet Protocol – today's familiar TCP/IP, developed from the initial protocol suite designed by Cerf and Kahn.

Packet-switching had moved since 1969 from concept to reality and was starting to open up a host of potential applications for computer networking. The creation of the TCP/IP suite ushered in a new world of inter-networking possibilities.

THE DEFENCE RESEARCH ESTABLISHMENTS

In 1980 John Robinson had just graduated with his Ph.D. in mathematics and high-energy physics. Immersed for years in the world of sub-atomic particles, Robinson knew his leptons from his mesons but there wasn't much call for explaining the difference outside of academia. After weighing his options, Robinson accepted a post-doctoral fellowship at the Communications Research Centre (CRC) in Ottawa.

The CRC was under the jurisdiction of the Department of Communications (DoC) and formed part of a mammoth compound located in Canada's premiere high-technology neighbourhood. Although the compound once dominated the then rural landscape of Shirleys Bay, its modern list of neighbors reads like a roll-call of some of the largest high-technology companies in North America: Nortel, Newbridge, SPAR, Mitel, Digital Equipment. Shirleys Bay is now better known by its self-appointed title of Silicon Valley North.

In addition to the CRC, the 600-hectare compound housed the Canadian Space Agency and a military research laboratory belonging to the research and development branch (DRDB) of the Department of National Defence. DRDB has such laboratories throughout Canada, known collectively as the Defence Research Establishments (DREs). Nevertheless, DRDB's communications research has traditionally been carried out by the CRC's military research arm, the Directorate of Military Communications.

Robinson was given a great deal of freedom to explore the CRC's facilities in search of projects that caught his interest. Since part of the DoC's mandate was to manage Canada's over-the-air spectrum, the CRC was focused on developing new technology and applications for wireless and broadcast communications. What really intrigued Robinson, however, was work at the Directorate of Military Communications sponsored by DRDB and intended to create its own version of the ARPAnet. Since the US network had moved out of its experimental phase and was becoming an important tool within the international research community, DRDB and Canadian military researchers began looking into how they could get involved.

In 1980 DRDB had carried out two internal audits to determine if there was a need for an ARPA-like network and, if so, what the main options were. The resulting studies, written by Intellitech Canada in 1980 and 1981, concluded that a network allowing interactive access to remote computer systems, electronic mail and file transfer would greatly benefit the Department of National Defence as a whole. The reports also stressed that DRDB would benefit even more if its DREs had the same access to colleagues outside the organization itself. The studies pointed to the functionality of the ARPAnet as an ideal end-goal, and suggested two different technical options for reaching that goal.

The first option was to rebuild the ARPAnet in Canada. Technically this would be the most straightforward approach, although it would require a great deal of money to build and maintain the network. The second option was to take advantage of the X.25 public data network also being considered by the new CDNnet in British Columbia. This second model was technically more complex, especially since the network had to connect to the ARPAnet. The X.25 model might also limit the network's overall functionality.

DRDB found it easy to make arrangements with DARPA (ARPA had since had "Defense" prefixed to its name). DARPA was interested in promoting international collaboration on networking, as was evident from its involvement in the founding of the International Networking Working Group and its role in SATnet. In return for connecting to the ARPAnet, all that was expected of DRDB was that it participate in various networking collaboration conferences where Canada would share the knowledge acquired from developing its own network.

BUILDING A LINK THEY COULD LIVE WITH

In 1982 Keith Hooey, Robinson's boss at CRC, approached him to take charge of the ARPAnet project. Robinson was intrigued by packet-switching technology and jumped at the chance. Perhaps all those years studying particle physics had left Robinson predisposed to a network that moved bits around at high speeds like a particle accelerator – though in this case it was data packets not atoms that were accelerated.

Robinson's task was technically straightforward. The packet-switching technology was already developed, so it was simply a matter of implementing it on this side of the border. The location of the first connection had already been decided. The DRE laboratory in the Atlantic provinces (DREA) had wanted an ARPAnet connection for some time because researchers there had been working with colleagues from Stanford's Artificial Intelligence Lab. While Stanford had been on the ARPAnet for years, DREA was out of the loop and found it difficult to collaborate with their colleagues and communicate resources and findings.

In 1983 Robinson and Hooey flew to the Canadian embassy in Washington to meet Vint Cerf and Bob Kahn and work out the technical details of joining the ARPAnet. 1983 was a critical year for Cerf and Kahn and the ARPAnet as a whole since it marked the official switch-over for the network from its original Network Control Protocol to the TCP/IP suite developed by the two engineers. At the meeting in Washington they all discussed the proposed Canadian extension of the network and outlined the equipment and other technical elements the CRC would require. They decided that the initial DREnet link would use a router developed by BBN, one of the ARPAnet's original hardware manufacturers.

Robinson got to work setting up the IP link. Unfortunately the complicated logistics were made even worse by some extraordinarily bad timing. The plan was simple: lease a telephone line from DREA to the CRC and a second line from CRC to Rochester, New York where a connection to the ARPAnet would be made. Leasing a line across the border required the CRC to deal with Bell Canada for the Canadian leg of the journey and then arrange a carrier partner on the US side.

As luck would have it, the American telecommunications industry was in the throes of deregulation. The breakup of AT&T into several regional Bell operating companies left the CRC waiting at the border with only half of its southern connection for almost four months. Robinson remembers this as a very frustrating time, although the connection did get established before the year was out. The link represented Canada's first connection to what would eventually be known as the Internet. From the American perspective, it represented the ARPAnet's first permanent international connection over leased land lines.

The ARPAnet connection was successful, but Robinson and his colleagues at the CRC knew that the cost of the link was a problem. As a pilot project, DRDB had been willing to dedicate the required funds to establish an ARPAnet link, but the equipment and especially the leased-line costs made maintaining a real network in this fashion a prohibitively expensive proposition.

Although leased-line charges were certainly more expensive in 1983 than today, Robinson remembers the problem was also one of priorities. Few people, especially in government, thought of data communications as an essential resource. Today, organizations from multinational corporations to home office businesses treat their Internet link as a lifeline and an unavoidable cost of doing business. In 1983, however, the CRC simply couldn't justify spending hundreds of thousands of dollars building a research network to link the DRE labs.

There was no choice. Robinson and the team from Software Kinetics, a high-tech defence contractor from Ottawa, had to exercise the second option suggested by the Intellitech study and build an IP-over-X.25 internet.

PLAN B UNFOLDS

The choice of the X.25 packet-switching technology was ironic, since one of the reasons X.25 services had never really taken off beyond highly specialized corporate uses was price. The X.25 service charge is based on volume, distance and time. Not only does this approach make the service expensive, but it also makes it extremely difficult for users to estimate their bills.

For Mike Martineau, a Software Kinetics employee who worked on the DREnet project, this was his first exposure to the X.25-based Datapac service:

> The cost had distance, it had volume, some would swear it included the phase of the moon and the value of pi. I defy anybody to calculate what their bill would be. And then each country and telephone company ran their own version and there were gateways in between them and there were charges for going through the gateways. The telephone companies took something that was a good idea and tried to make a lot of money out of it. Instead they ended up with a very small audience who used it. But in the days when we were working on DREnet, which was low-volume, long distance work, Datapac actually made sense compared to leasing dedicated lines.

It took several months for Robinson to design the system, and months longer to implement and test. Software Kinetics and Robinson designed their network to run on smaller computers such as the micro-VAX, which would not be difficult or expensive for the DRE labs to get their hands on. In retrospect, Martineau marvels at the breadth of services Software Kinetics had to supply simply because there were so few established players at the time. Very few products were available off the shelf. “We actually built routers for DRDB,” said Martineau. “That’s how early this stuff was.”

Robinson started encapsulating IP packets on top of the X.25 packets and worked with Tim Symchych from Software Kinetics to build the specialized routers. The first prototype of the network connected a number of computers called PIN9101s to Bell Canada’s Datapac 3000 X.25 service ports. The prototype worked, but it was not a practical configuration. A second, more ambitious project was mounted in September 1985 that delivered data across several international packet-switched networks, including the ARPAnet, in order to run performance measurement tests.

But in 1985, just as the IP/X.25 network had been developed, Robinson left the CRC to take a three-year contract position at the NATO Technical Centre in The Hague. It was an opportunity he couldn’t pass up, but on his departure DREnet was left at an embryonic stage.

FACILITATING DRENET

Robinson had proven that IP over Datapac was a viable way to build a network. But the link to DREA really only proved that it could be done and done affordably – it was more of a pilot project than a network. The only access it provided was to two labs: DREA and the CRC. The rest of the DRE labs scattered across the country remained isolated from each other and their international counterparts.

The technical elements of a network had been created: the micro-VAX routers, the Datapac service, the 56 kilobit-per-second (Kbps) link to the US. Yet no one in a position to push things forward had recognized that these were the pieces of a promising puzzle waiting to be put together. Then in late 1985 Vincent Taylor arrived at DRDB headquarters in Ottawa as the staff officer responsible for Information Technology Research. As he puts it, “I didn’t have a lab, but I had some money and a list of sensible things that had to be done. And building a network was a sensible thing to do.”

Taylor had been a Scientific Liaison Officer at the Canadian Embassy in Washington during the late 1970s, acting as the interface between defence research communities in Canada and the US. He had watched the development of the ARPAnet with interest for some time, and his work led him to collaborate with DARPA on research initiatives. The focus was on computer security issues and especially on gaining a fundamental understanding of computer processes from start to finish. “We were really dealing with how to develop software that let’s you know what it’s doing, as opposed to relying on what you think it's doing," says Taylor. "It's a problem that's still not well solved today."

Taylor's job dealt with ensuring that software didn't breach certain performance specifications outside of visible test points:

> The whole idea in this area is to try and develop systems that do not violate policies. The policy might be a security policy or a safety policy, or something else altogether. But you need to be able to prove that your system won't violate these policies.

When Taylor arrived at DRDB HQ it was his expertise in building systems – human systems as much as technical – that would give DREnet its shape as a network.

Although the plan for DREnet did not include more than a dozen nodes, network management was a challenge simply because everything was being done for the first time. There were very few guides to follow, no peers to consult and no off-the-shelf software for anything, including network management. On the technical side, the prototypes developed under Robinson had to be turned into standardized equipment.

Looking back on DREnet's development, Taylor describes his activities as facilitating the creation of a network:

> We did two things. One was produce the network. No, I take that back. We produced the ability to have a network. And that's a little different. At no stage did we impose the network like a project. What we said was, here is what you need to have your network.

The second thing Taylor did was design the system of operations for managing and troubleshooting the network: connecting to the network meant defining two roles, the "responsible person" and the "responsive person." The responsible person dealt with issues of policy, while the responsive person handled technical and operational issues. It was a simple structure, which is probably why it is still used today.

It was also effective. Taylor remembers one time he was notified that the DREnet secret password file had been located at a site in Georgia:

> That was a little alarming, but it was very interesting to see how our measures went into place. What we did is basically cut off the network at the US border. The network of responsive people worked to ensure that each lab changed all of its passwords. As sites reported back that they were clean and had found no evidence of any tampering, they would be allowed selectively back onto the network. The process almost went into play automatically.

As a result, DREnet ended up more ad hoc and decentralized than the ARPAnet itself.

By the time DREnet became operational at the end of 1985, the ARPAnet had been divided into two networks: the ARPAnet and MILnet. The latter, which survives today, was reserved for US military research on a highly restricted basis. From the start DREnet was a distinct network in its own right (not just an extension of the ARPAnet) and was officially a peer network with MILnet. In the same year the network operations centre was relocated from the CRC to its Shirleys Bay neighbor, the DRE Ottawa lab. As Taylor points out, "We were really the first part of the Internet in Canada – although the term was hardly in common use at the time."

DREnet was once the most advanced network in Canada, but it was also (and remains) the most exclusive and invisible. Nevertheless, its construction served as an important training ground for Martineau and Software Kinetics, who would go on to play an important role in a much more public undertaking. And in the meantime, a new networking movement was already under way outside of both academic and military research circles – a grass-roots movement that gave a wider audience a taste of global networking years before the Internet arrived on the scene.

USENET IN CANADA

Canada's first official connection to the ARPAnet was made through DREnet in 1983. Unofficially, however, Canada had already been connected to the ARPAnet through the upstart international network known as Usenet, established in 1979.

Usenet is actually the software application that runs (mostly) on top of UUCPnet, but "Usenet" is commonly used to refer to the network as well. UUCP stands for Unix-to-Unix CoPy which, as the name suggests, copies data from one UNIX machine to another over regular phone lines with a computer modem. Although called "the poor man's ARPAnet," Usenet did not employ packet-switching or networking protocols, operating instead as a store-and-forward system with dial-up phone lines. The network also used manually updated lookup tables rather than fully automated computer updates to guide data to its destination.

Usenet, the software application that encouraged the growth of UUCPnet, worked like a distributed bulletin board. In addition to sending and receiving personal email messages, users could read and respond to messages publicly posted in newsgroups dedicated to various subjects. With the store-and-forward technology, messages they mailed or posted were first stored on the local computer. Then, at pre-determined times (often late at night, when phone rates were low), the computer would automatically dial out to another computer on the network to exchange news and mail. News went to all the computers on the network while email went only to its intended destination, but both traveled from computer to computer to computer, one hop at a time.

Usenet depended on the cooperation of many computers to serve as network nodes. A number of these nodes were loosely referred to as the network's "backbone" because they had the technical and financial resources to handle heavy traffic and send it long distances over phone lines.

Cost was an important factor because, when the network first started, it was a cooperative venture and charging for access to the network was considered poor etiquette. Nodes that, for whatever reason, could afford the phone bills had a great deal of control over which newsgroups were distributed and which weren't. The managers of the backbone computers became known as the "Backbone Cabal." They practiced censorship, not only to keep down costs but also because newsgroups with racy content might raise the eyebrows of executives who had agreed to foot the phone bills as a "company resource."

UUCPnet's low-tech quality actually helped it spread rapidly. Launched in 1979, Usenet quickly dwarfed the ARPAnet in terms of nodes, users and traffic. It was also unique in that it spread beyond academic and military circles into the residential and corporate sectors – at least within specialized groups including individuals such as Mike Martineau that knew about the network:

> It certainly wasn't business quality. People in the technical community were using it mainly for communication amongst themselves. The links weren't reliable, the delivery of mail could take hours or days... It was maintained very lovingly by volunteers whose companies may or may not have known they were donating time, equipment and/or communication links!

By 1983 there were over 600 nodes on Usenet serving thousands of users. Marginal though the network was, it was a powerful source of information (sometimes misinformation) and communication. Perhaps because of the heavy overlap between computer fanatics and science-fiction fanatics, the Usenet community developed a richly textured collection of myths and sagas about itself and its users. Network celebrities were immortalized as "Net Gods" and pivotal events in the network's development were described by grandiose titles such as "The Great Renaming." Usenet was also responsible for producing a great number of slang terms and acronyms that survive in email messages and Internet chat sessions today: IMHO (in my humble opinion), TIA (Thanks In Advance) and so on. But, as many old-time Usenetters explain, nothing marks a newcomer as clearly as excessive use of these acronyms!

Canada gave the Usenet world at least two Net Gods of note: one who played an important role in giving Usenet its form, another its content and character.

ZOOLOGY TAKES A BOW

In 1981 Henry Spencer was running the zoology department's computer facility at the University of Toronto (U of T). In May of that year he succeeded in having a computer in his department connected to the UUCPnet, thereby becoming the network's first international node:

> Zoology being the first Canadian site on UUCPnet was a little weird, but it was an accident of timing and personalities. The more obvious departments at U of T were all enthused about the ARPAnet and weren't interested in second-rate networking done intermittently at phone-line speeds. They changed their minds quickly when it became clear that they could get email now without waiting years or spending piles of money.

Mindful of his place in Usenet history, Spencer admits to a potential quibble over who can actually lay claim to the title of first international node (utzoo):

> If we insist on being picky, the first international Usenet site was the UNIX system that Ron Gomes ran at Bell Northern Research in Ottawa. I don't think anything was ever posted from it, it was a read-only system for all practical purposes. It faded and died after Ron left BNR. Utzoo was second by a couple of months.

Obtaining a connection at that time consisted primarily of getting listed on the logical maps of the network that were distributed to all the nodes and convincing someone to include a particular node in its file transfer process. In Spencer's case, U of T's zoology computer called a computer at Duke University in North Carolina. Exclamation marks, called "bangs," represented the UUCP network links or "bang paths." The bang paths were used to address email and other Usenet traffic by spelling out the route that data had to travel to reach its final destination. Thus, a Usenet address might look like this: duke!duke34!utzoo!henry. Spencer recalls:

> It quickly became common practice to sign your name with an indication of how you could be reached from a few of the best-known backbone sites, e.g. {allegra, ihnp4, decvax, pyramid}!utzoo!henry.

Becoming a major distribution point for Usenet was no small undertaking. While attending a Usenet backbone meeting at the Usenix conference in 1983, Spencer learned that DEC had spent over $250,000 in phone bills that year distributing Usenet traffic:

> [DEC] averaged 13 hours a day of long-distance calls. The main problem was characterized as "the system operator is being too soft-hearted about letting

> new sites connect." I couldn't complain too much since we were among said sites, initially with manual dialing at 300 baud. These soft youngsters have no concept of the joy with which we greeted our first 1200 baud modem – it had an **AUTODIALER**! By common consent getting it installed and working took priority over everything else.

In 1985, Spencer became well known in Usenet circles as the co-author, with Geoff Collyer, of the software used to operate the network. The program had gone through several evolutions when Collyer and Spencer wrote what became the standard version, known as C News. "Geoff did most of the hard parts," claims Spencer, "or at least that's the way it seemed to me."

C News was used throughout the UUCP universe until UUCPnet itself was incorporated into the Internet. But even after the Usenet backbone was replaced by Internet connections, the fringes of Usenet continued to communicate using UUCP links (and C News) that remained quicker, cheaper and easier to install and maintain. Local UUCP sites remained the most common means of connecting to Usenet, and what was referred to generally as the Net, well into the 1990s. In fact, UUCPnet was still the largest computer network in the world until the TCP/IP-based Internet finally overtook it in 1994.

GONE FISHING

A second Canadian nominee to the Usenet pantheon is Richard Sexton, whose legacy is as colourful as it is distinguished (in the Usenet universe the two often go hand in hand).

Sexton moved to Los Angeles in 1979 and eventually finagled his way onto the Usenet backbone. He had become friends with Brian Reid, who ran DEC's research lab in Palo Alto. DEC, as Spencer can attest, was an important supporter of Usenet and Reid was well respected in the community. Reid pulled a few strings to convince NASA's network node to share traffic with a friend's home computer, a node they called Gryphon. Sexton describes the ploy:

> Just because Reid thought I was funny, he wrote a letter to NASA JPL and said, "I'm concerned about network activity in Los Angeles and I believe you should give Gryphon a feed." The administrator at JPL said, "Oh my God, it's Brian Reid, whatever you say Brian!" Brian was a real Net God. So every 20 minutes JPL would call little Gryphon and dump off news and mail and pick up news and mail and transport it. The people at NASA thought Gryphon was a massive computer, contained in a giant frigid room cooled by synthetic blood and part of an enormous company called Trailing Edge Technologies which owned the patent on the trailing edge of aircraft wings and funded the entire operation. Actually, there was a PC on the sofa, possibly the first PC on the Net: a 386 running Xenix of all horrible things. But I had about 100 downstream feeds in the LA area. It pretty much was the backbone of Los Angeles.

A contentious issue at the time was newsgroup distribution, because of the added cost and time required to distribute every upstart group – but Sexton enjoyed contentious issues. He had long been a fish hobbyist, specializing in a tropical breed called Killifish. After being scolded for talking inappropriately about his fish in newsgroups that weren't fish-oriented, Sexton decided that Usenet was in desperate need of a "sci.aquaria" group. The rest of Usenet needed convincing.

The issue was debated in a special newsgroup that served as a quasi-democratic forum where people could vote for or against the creation of a particular group. However, the process wasn't exactly objective. Sexton again:

> Newsgroups used to be created automatically when you posted to them, but that didn't work for very long. So what the Backbone did was decide we'd have a vote and, at the time I was involved, the way it worked is you had to start a discussion as to why this group should exist, and that would go on for two weeks. Basically I would say we should have a particular group and everyone else would say, "No, that's a stupid idea." But what they failed to put in the rules was that the discussion was binding in any way. So everybody would agree it was a bad idea and then I would go to the next step, which is call for a vote. The person championing this group would then collect votes. Having someone that wants the group count the votes is a pretty stupid idea. It's so prone to corruption it's not funny.
>
> So people would send in their votes, yes or no. If you got 100 more yes-votes than no-votes your group was created. Again, Brian Reid egged me on. "You've got to do it, I'll support you all the way." And he did... I had the biggest voter turnout in history. Usually these things are 150 to 120. Sci.aquaria was 950 to 780. Everyone forged votes for and against this thing but it passed and actually it is still a decent group today.

Sci.aquaria earned Sexton a great deal of good-natured ire from the Usenet community. But his next stunt would make a truly profound, and questionable, impact on all of Usenet – and ultimately the Internet itself.

One day, while demonstrating the basics of Usenet software to a friend, Sexton jokingly posted an article suggesting the creation of the newsgroup called "rec.fucking." The article was never meant to leave his home PC, but a typo in the posting instructions sent it out to the LA Usenet community at large. Although the "group" was immediately quashed by the Cabal as inappropriate, the subsequent demand for a sex newsgroup was overwhelming.

Eventually the sex issue got swept up in the wider debate over control of Usenet by the Cabal. After repeated rejections of suggested groups by the Cabal, the so-called "alt" newsgroups (for "alternative") were created on an independent set of backbone sites that footed the transmission costs themselves.

The extremely prolific alt.sex newsgroup hierarchy became the home of a colourful array of sexual interests, peccadilloes and perversities. The creation of the alt hierarchy also signaled the beginning of the end for the authority of the Cabal.

While Usenet has since faded from view – integrated into browsers and overwhelmed by spammers – the network application has left an indelible mark on the culture of the Internet. Far more than any other network of its time, it demonstrated the power of online communities to serve as a resource, as recreation and as a way of drawing people from every walk of life into the vagaries of Net culture.

NETTING THE ARPANET (UNOFFICIALLY)

It was in 1981 that Usenet made first contact with ARPAnet. When the University of California at Berkeley joined the network in 1981, it created an overlap between the ARPAnet and Usenet. The Berkeley administrator, Mark Horton, began to funnel ARPAnet mail lists to the Usenet network. These ARPAnet links are represented by the familiar @ sign used in email addresses today.

While creating gateways between the networks was a clear violation of ARPAnet security regulations and acceptable use policies (AUPs), most administrators turned a blind eye to this unofficial expansion of the network. ARPAnet traffic began circulating through Usenet like any other news and mail. More gateways were built and began to allow for two-way traffic between the networks. The consistency of the connections varied, but these gateways allowed ARPAnet discussions and articles to flow beyond the borders of the ARPAnet, including into Canada via its link through U of T.

By the early 1980s, fears that Canada would fail to develop its own domestic networks and simply become an adjunct to ARPAnet had yet to come to pass. This was as much due to the US military's reluctance to provide access to their research network as it was to Canadian efforts to develop a domestic equivalent. In the latter category, Bell's pioneering Datapac service, inspired in part by the CANUnet project, had served as an enabling platform for two independent national networking initiatives, CDNnet and DREnet. These stood out from the many local and regional efforts being carried out at universities across Canada by virtue of their scale and innovative use of packet-based networking. They taught their operators important lessons about the construction and operation of advanced packet-switched networks, and they gave a sense of the possibilities that still lay untapped.

The UUCP network or Usenet, on the other hand, introduced the concept of an ad hoc approach to networking based on large-scale, distributed cooperation. UUCPnet didn't require any special services or cooperation from the telephone companies – in that sense it was more independent than CDNnet or DREnet. Neither centrally managed nor organized, UUCPnet was supranational.

Another new networking movement was now on the horizon, borrowing heavily from the UUCP network's independent and ad hoc (and lower-tech) approach. Centred in Canada's academic community, it would come closest to realizing the CANUnet vision.

3

A NATIONAL NETWORK (...AT LAST)

"SUDDENLY I KNEW WHAT COMPUTERS WERE FOR. THEY WERE TO TALK TO EACH OTHER." – *ROGER WATT*

In 1983, the ARPAnet moved from its original Network Control Protocol (NCP) to the TCP/IP protocol suite, still the technology at the base of today's Internet. But the early years of the decade also saw the onrush of a whole new wave of networks, and network protocols, running in parallel to the ARPAnet.

"BECAUSE IT'S TIME": THE BIRTH OF BITNET

The most common of these were proprietary networks developed by integrated hardware and software manufacturers such as IBM and DEC. DEC had designed its own protocol and accompanying network called DECnet, although the real giant in the field was still IBM, and university network initiatives such as the NB/PEI project depended almost exclusively on centralized, time-sharing IBM mainframes connecting remote terminals.

The IBM proprietary software included the VM operating system and a set of communication tools known as the Network Job Entry (NJE) protocols, which were designed to connect two mainframes and transmit files between them. IBM's position in the computer industry – especially within the academic and research communities – meant that almost every sizeable lab had at least one VM machine that could be used as a network gateway.

Universities used NJE to interconnect computers to share information on and between their own campuses. For example, U of T had been interconnecting computers between its suburban and downtown campuses using this capability since the late 1970s. Other universities such as Waterloo, Guelph and Queen's University in Kingston had been experimenting with using NJE to set up links between universities, but none of these experiments had led to any large-scale ventures. The time was ripe for the next step forward.

In 1981, Ira Fuchs at the City University of New York and Gordon Freeman at Yale found an innovative way to use their IBM time-sharing software. Both City and Yale were using IBM mainframes, the VM operating system and the NJE protocols. What Fuchs and Freeman did was lease a permanent telephone line to connect their computing departments, interconnecting the computers at each of their facilities as though they shared the same room.

This link established the first two nodes in what became known as BITNET (the "Because It's Time" Network). BITNET grew quickly in an ad hoc fashion. Because IBM was in its heyday, IBM equipment and software were prevalent enough to act as a de facto technical standard under which BITNET flourished. Computer networking began to reach the mainstream research community, making it accessible on a broader scale. BITNET was hot – in the US, that is.

On the other side of the border, Canadian universities that used the same computer systems as their US counterparts, moved in the same academic circles and were exposed to the same scientific discoveries watched BITNET take off with admiration and a touch of envy. David Macneil from UNB heard Fuchs talk about the network at a computer services conference in 1981 and was immediately interested in joining the effort.

Roger Watt, assistant director of computing services at the University of Waterloo, was caught up in the excitement surrounding the growth of BITNET at a SHARE conference (a biannual gathering held of the IBM user group):

> Suddenly I knew what computers were for. They were to talk to each other. And about that time I lost all interest, comparatively speaking, in what one could use a computer for other than to have a network. The whole field of networking became far more interesting to me than any other way of using computers.

It was an enthusiasm Watt shared with Canadian colleagues he ran into at these and similar conferences, including Alan Greenberg from McGill, Paul Molyski from the University of Victoria and Darwin Fedoruk from the University of Alberta:

> We'd all meet at SHARE twice a year and we'd all say to ourselves, "Geez, wouldn't it be nice if we could get our universities connected to this BITNET thing. This is really neat stuff." Except that the people who were responsible for managing BITNET... they had this policy that you had to connect at 9600 bps. Well, 9600 bps lines in Canada cost one hell of a pile of money and getting them across the border cost two hells of a pile of money.

The reality at the time was that prospective BITNET members had to pay for a connection to the nearest node in the network. In the US, where schools are typically larger and closer together, connection costs – even at the top-end 9600 bps – weren't usually all that expensive. In Canada, on the other hand, computer service directors tried for two years to get their individual universities connected to BITNET – unsuccessfully. The costs and regulatory headaches involved in acquiring the necessary lines were just too great.

A FLYING START FOR OUNET

The setback with BITNET soon turned to Canada's advantage. Canadian institutions stopped casting envious eyes south of the border and started focusing energy and resources on the grass-roots efforts that had developed around them. The universities of Guelph and Waterloo were ready to take the next step, having already experimented with logging on to each other's mainframes through data PBXs (telephone exchanges such as the one used by the university switchboard to direct calls).

DISTRIBUTED COMPUTING AT GUELPH

The ECN network developed by New Brunswick and PEI was one of the more elaborate and ambitious local networking projects in Canada but it was by no means the only one. In the early 1970s Kent Percival at Guelph had his hands full developing and maintaining a fairly sophisticated local computer network. Guelph had pursued a model of "distributed computing" rather than the centralized approach typical of most campuses at the time. Instead of a mainframe at the computer centre servicing terminals throughout the campus, different departments had their own computer systems to carry out specialized activities. These systems were then networked to the centre's IBM mainframe. Percival, the network operator at the time, explains:

> Our library operated their own computer system. Our registrar's office and several academic departments had put in small systems. One of the things I worked on was to link up all these computers on campus using various types of communication lines. All of these different kinds of computers were linked together to our IBM system, and the IBM mainframe became the central point for communication between them.

The university had developed an email system using the APL programming language for its own interactive computing environment as early as 1975. This reflected a computer-friendly climate, as Percival put it:

> I think there was a climate at U of G with our distributed computing models that put more information technology people out in the departments. They encouraged greater interest in using the network because there would be a few information technology people and a few keen professors in each department that were getting onto the email system.

Guelph's kindred networking spirit for some time was the University of Waterloo. Percival had worked on several projects with Waterloo, and on some with Queen's that hadn't come to fruition. These involved experimenting with remote logins by hooking terminals into data PBXs and logging into the other school's mainframe. The technique worked, but it wasn't the most efficient or easiest way to network: the universities would have had to duplicate all their administrative systems and information in order to support remote users logging into their mainframes. Percival comments:

> We realized that what we needed was something providing support to remote users other than our main computer systems. Something that would allow a person at the university to log on to the U of G computer and be supported from there to access other universities.

When the solution eventually appeared, these pioneering efforts laid the foundations for the first connection in a national network.

LINKING THE PROVINCE ... AND BEYOND

What the universities of Waterloo and Guelph really required was a computer switching facility, not a telephone PBX. In 1983 an opportunity arose that Percival and his colleagues couldn't turn down:

> Waterloo and Guelph ran joint graduate programs in chemistry and physics. I believe it was the physics group that came to us and said, "We have this big computer over at Waterloo purchased to do physics computing and we'd like to be able to move our data back and forth between the Guelph mainframe and Waterloo." We also wanted to be able to log into that computer from Guelph and I think what we eventually put in was a 9600 bps link and dedicated half of it to establishing an NJE link. We multiplexed it and half of it was an NJE link between the computers and the other half was connecting our PBXs together. I think we allowed four 1200 bps dial-up connections from terminals at Guelph to the computer in Waterloo.

The two 9600 bps modems at either end of the connection had to be readjusted by the operator every time they tried to exchange signals. The familiar screech of two modems connecting to each other is the automated version of this manual tuning process, known appropriately enough as "handshaking."

But these tedious details did little to slow progress. The spirit of collaboration behind BITNET had found its way north. As Watt puts it, attitudes were becoming more positive:

> I'll lease a line from me to you, you lease a line from him to him, and we configure our systems to move unit record spool files across these damn things just like we did internally, and poof, we have a network!

This atmosphere of cooperation was possible since Percival, Watt and other university computer managers in Ontario were colleagues who knew each other and had already worked together to organize inter-university networking of a more human and personal nature: conferences, associations, user groups, papers. These well-established personal connections planted the seeds for the new network to flourish. The members of one of these user groups, the Association of Computer Service Directors of Ontario (ACSD), would play an especially important role.

At an ACSD conference held in Toronto at Ryerson Polytechnic Institute in November 1983, the group focused on how they would respond to the ambitious step taken by Waterloo and Guelph. Percival had spoken with Dick Mason, Guelph's computer centre director, prior to the meeting about the feasibility of building a BITNET network in Ontario:

> I recall sitting down with some data and costing numbers I got from Bell. I had convinced our Bell rep to give me the whole price sheet on all of their connections so I was able to estimate costs between the various sites. I started thinking these numbers are getting close to what organizations might be willing to pay, especially if we did a cost averaging or equalization so that Windsor way down at the end didn't have to pay a whole lot more than Guelph linking to Toronto. So I raised this with Dick Mason, essentially my boss's boss, and very quickly found I had an ally. He said, "Go figure it out and convince us."
>
> He was keen enough to let me take some time to work out the numbers and arguments and try to sell it to the computer centre directors directly at the next ACSD meeting. So he got me on the agenda and I lobbied ahead of time... I did the presentation on how the network would work and how we would connect it up and showed them some dollar figures and got agreement to proceed – but subject to each computer centre director making the budget decision themselves. We started with a core group saying "yes."

From that meeting came commitments from six more universities and colleges looking to connect their mainframes to the two-node network of Waterloo and Guelph: U of T, York University in Toronto, the University of Western Ontario in London, Queen's, Humber College in Toronto and Ryerson. The group initiative was tentatively named the Ontario University Network – OUnet for short. As these links were being established over the next few months, the network also received requests for interconnection from Lakehead University in Thunder Bay and three out-of-province universities: the University of Manitoba in Winnipeg, McGill University in Montreal and UNB.

OUnet suddenly had a vision of stretching across Ontario and beyond – but it was still short "two hells of a pile of money" for anything to be done about it.

The challenge of overcoming Canada's geography through networks, whether transportation networks or communication networks, is hardly a new theme in Canadian history. Yet it still has to be faced by each new generation of technology. Percival explains:

> One of the first bumps in costing was when Thunder Bay [Lakehead University] asked to join. That is a very long line. It was a struggle for some of the other members to say, yes, we'll accept the fairly significant increase in our shared costing for a single member. I think there was some question when Alan Greenberg [of McGill] wanted to participate about whether we treated him as part of our community or as external. We agreed we would link Ontario and Quebec and he became part of the network community and paid his equal share. Then when Gerry Miller from Manitoba and Dave Macneil from UNB got involved, we were looking at fairly significant connections out there. Dave was very keen on operating a Maritime hub and taking responsibility for managing a separate community down there.

At the same time the OUnet members again cast their eyes southward. A connection to BITNET didn't simply mean making a connection with their American colleagues – it would also connect the Canadian universities to the international NJE research network. IBM had already subsidized a link connecting BITNET and the European Academic Research Network (EARN) in 1982. From EARN, connections would eventually extend to networks in Asia, Australasia, the Gulf region and Africa. The accessibility of the NJE protocols was quietly creating a global resource of news, data, communication and collaboration which the Canadian academics were anxious to be part of.

But expansion in all directions cost money, and the universities simply did not have it. Cost sharing had allowed the network to progress as far as it had but no further – not without help from a familiar figure on campus.

FOCKLER'S TALE

Ken Fockler was well known to most Canadian universities – especially around their computer departments. He was IBM's man on campus, not initially in a sales or marketing position exactly but as more of a roving ambassador. His position at IBM had developed in 1963 out of a project at Western: working as a systems engineer, he helped install one of its first large scientific computers. Afterwards, he says,

> IBM, in its wisdom or whatever, let me roam the campus once I got the mainframe installed. The computer science people were already using it, and my job was to encourage other departments to use it as well. It was a great learning experience.

IBM decided Fockler's campus walkabouts were productive enough to make such efforts his full-time occupation. He continued to visit campuses, learning about current projects, listening to complaints, making suggestions. He relayed the latest news from IBM and from the computer industry in general. He also helped non-computer departments take advantage of the computer resources already on campus, showing them applications they might find useful. "Like here's what the anthropology department in Princeton is doing with their computer," as Fockler put it. "They would get a hold of other people's programs or get ideas that way."

Fockler could be described as a one-man inter-campus network with his own messaging, news and file-transfer applications. But in the main, his job was to be helpful, to listen, to be interested. And he was – Fockler liked academia. And, for the most part, academia liked Fockler, though he noted some reservations:

> They didn't always like IBM that much. In a couple of instances some of the people who decided they didn't like IBM said, well, Ken Fockler can come around, because he's not selling, he's bringing some interesting ideas. But we'd rather not see any IBM salespeople on campus!

SOME FRIENDLY EXTORTION

Fockler eventually joined IBM Corporate Programs – a gift-giving arm of the company established to provide grants and enter into cooperative projects with its clients. In that capacity he ended up having a drink with some OUnet computer service directors early in 1984. Very quickly Fockler realized that his friendly chat over a beer with Paul Dirksen from Waterloo, Dick Mason from Guelph and Dave McNaughton from Queen's was becoming an exercise in good-natured extortion. OUnet was looking to expand. IBM had helped EARN connect to BITNET. Would IBM Canada be interested in helping out the Canadian leg of the NJE network?

The message was that OUnet was going to take off whether IBM was along for the ride or not. There were always other computer companies to approach. IBM arch-rival Digital Equipment might be interested for instance... Fockler decided IBM Canada would indeed be along for the ride – all he had to do was convince IBM Canada.

The ultimate goal, as OUnet and Fockler saw it, was to build not just an expanded provincial network but a national network connecting the major Canadian universities, at least one in each province, and connect that network to BITNET in the US. The network would remain facilities-based: rather than plug into the public telephone networks and rely on Bell's routers, the universities would lease access to dedicated phone lines and control their own traffic using the IBM hardware and network protocols. The real challenge was funding the long telecom lines necessary to link the various Canadian regions.

Given the potential scale and expense of the project, Fockler decided to convince IBM by building support incrementally. He broke the project down into two phases. The first phase involved defining and establishing the network operations centre, and providing the equipment and software required to run it. Guelph had in fact already volunteered to act as OUnet's NOC and maintain the connection to BITNET, so the only other support he needed for the first phase was from Lorne Lodge, IBM Canada's CEO and head of its Operations Review Committee (ORC). The second phase would involve expanding the network inside Canada, and that would require not just support from IBM but cooperation and support from many universities as well.

Fockler prepared his proposal for the ORC very carefully. The BITNET connection had many elements that made it an easy sell. It was relatively simple logistically, it wouldn't cost a tremendous amount of money and, rather than a straight cash donation, the funding included the donation of a medium-range mainframe known as the 4341. On the other hand, anyone going before the ORC felt trepidation: it was like "going before the gods." Lodge himself was considered something of a tough nut, and the ORC's collective scrutiny was known to cause a high rate of mid-presentation panic. Fockler felt pretty nervous himself.

Contrary to expectations, Lodge was open to the OUnet concept and genuinely supportive when Fockler presented it in May 1984. The ORC followed suit and gave Fockler approval for phase one of the project. With this partial blessing Fockler began to prepare for a whirlwind tour of the country to see if the rest of Canada's universities were ready to support OUnet. Fockler needed a critical mass of universities to sign on to the project and the right geographic mix to make a national network fly.

Meanwhile, at the annual Ontario University Computer Conference, the OUnet members had grappled with the fact that their name, catchy as it was, had already been rendered obsolete. Roger Watt explains:

> We realized we had to rename it, that OUnet wouldn't fly, because we had colleges in it and we had Alan Greenberg in Quebec saying I want to get McGill connected to it, and in Manitoba, Gerry Miller saying I want to get U of M connected to it and Macneil wanted the network to run out East. So we had to get rid of the "O" and we had to get rid of the "U" and we were left with Net, and hell, we were the counterculture movement in the north determined to create our own network, so we looked at "Northnet." Then we went and had some beers at the reception and we came back to the room where we had Northnet up on the blackboard, and I said "No, let's make it NetNorth." That's a different ring. Nobody's done that yet. And we got up the next morning and discovered we had all agreed to do it, and that's where NetNorth came from.

HEADING EAST

In the fall of 1984, the first leg of Fockler's journey took him to the Atlantic provinces. At his first two stops he received a warm reception. Dana Wasson, the director of computer science for UNB, was certainly no stranger to the benefits of computer networks. The NB/PEI ECN had been running for well over a decade and inter-campus email had been used for years on the regional network. Similarly, in Charlottetown, UPEI president Peter Menke whole-heartedly embraced Fockler's proposal because such a network would allow faculty to connect with their colleagues in the US and Ontario. As he pointed out, at UPEI a specialized discipline such as anthropology is wedged in between experts in theoretical physics and animal husbandry with no one to talk to.

Fockler had an equally gratifying meeting with computer centre director Peter Jones at Dalhousie University in Halifax. But despite his receptive audience Fockler cleared out quickly. He was about to be upstaged by the Pope, due to arrive in a few days, and all of Nova Scotia was caught up in the preparations.

At Memorial University in Newfoundland Fockler found more support – tinged with incredulity. The reaction seemed to be, "Why are you even here asking us? It's obvious we want such a thing." But the outcome still wasn't obvious to Fockler, who couldn't yet give any firm promises or details: "It's one thing to make a connection to Fredericton. It's another to make the hop to St. John's..."

Fockler left the east coast with a sense he'd found potential partners who were motivated – as well as curious as to how the whole operation would work. Whose network was it? Who would manage it? How?

HEADING WEST

A few weeks later Fockler took off from Toronto in the opposite direction. The western leg of his tour would turn out to be a bumpy ride. Still elated by the results of his first trip, Fockler was a little taken aback by his first two visits to the University of Manitoba in Winnipeg and the University of Saskatchewan in Saskatoon:

> These were fairly large universities, islands unto themselves, pretty capable. There was some interest in the project – they bought the concept, they wanted to communicate. But they weren't great IBM users, so there was a degree of "What's going on here? What's IBM up to?" And you can't blame them. Many of them were large Amdahl shops! Amdahl was guarding their territory well and starting up user groups. There was a reserved kind of response from Manitoba. It was positive, but kind of reserved.
> The University of Saskatchewan had never been an IBM shop in my time involved in universities, so it was a difficult call. I was a real outsider there,

> whereas in Manitoba I'd met some of the people and they'd had some idea of who I was and I felt comfortable. And maybe my experience in the academic community helped me talk to them and come across like I wasn't out to rape and pillage! But in Saskatchewan it wasn't enough. I left that call thinking, if something happens they would probably participate. But they sure weren't willing to give it a lot of enthusiasm.

Fockler may have been disappointed, but he hadn't heard any outright refusals. Still, it would help his case with IBM Canada if he had a partner in the West who would act as the network's regional gateway rather than have the network proceed in a straight line like a railroad track from Guelph to Vancouver. Although a straight line network would cost less, the western leg would be more reliable if built in a star configuration, with connections to the universities radiating out from a single western gateway or node that would manage those links. Moreover, Fockler knew that passive cooperation might not be enough to convince Lodge and the ORC that the project was worth pursuing.

To Fockler's surprise, he found what he was looking for at his next stop:

> The University of Alberta in Edmonton was where I struck gold… a big Amdahl shop, but they also had some IBM. The computer centre director was Dale Bent. Dale either had the vision or he saw it as a good political move because he was quick to say, "It's a great concept, I know about these things going on, and I'll undertake to organize the West and be your contact." I found that amazing.

"BUT WHERE DOES THE MONEY GO, REALLY?"

Fockler returned from his trip West armed with the ammunition he needed to meet again at IBM with the ORC. Phase one was up and Guelph was moving ahead with its connection to BITNET. NetNorth had grown to include 13 institutions, having added McGill and the National Research Council's Ottawa laboratories. Now Fockler had commitments from universities in every province in the country. He also had three regional volunteers to act as network representatives: UNB in the East (already the hub of its inter-provincial network), U of A (Edmonton) in the West and Guelph as the central hub.

Fockler packed his presentation with sound bites, promises and goodwill. Despite having secured support for phase one of the project, support for phase two was far from a sure thing. The ORC members were particularly concerned with the large cash grant involved – over $400,000. IBM Canada was more accustomed to playing a supporting role and donating equipment. Where was all this money going exactly? Fockler was on the spot:

> So I tried to say, well the money is going to go to the University of Guelph, which will disburse it out to Alberta and New Brunswick and make these line connec-

tions. Oh no, that's not what we mean. Who ends up actually getting the money? The answer was CN/CP [Telecommunications] and Bell. At the time IBM was still trying to figure out whether these guys were customers or competitors – or both. So why should IBM be doing this, or at least doing it alone?

A number of the ORC members were dubious, to say the least. The proposal got caught on this snag and threatened to unravel. Lodge then suggested that IBM investigate the possibility of opening the NetNorth project to other companies, especially the telephone companies, and the federal government. Fockler left the meeting without any definitive answers, while the marketing department went off to find external support to help defray costs.

Fockler didn't hold out much hope they would pursue the task vigorously – he wasn't even sure they really understood the project. Yet Fockler was also getting the sense Lodge had been manoeuvring on NetNorth's behalf. Marketing might not get anywhere with Bell, Ottawa, or CN/CP, the telecommunications arm of the Canadian National and Canadian Pacific railroad companies, "but at least they were given a vehicle to vent their frustration," notes Fockler.

Four weeks later the ORC reconvened. Fockler pitched NetNorth for the third time. Marketing had come up empty – IBM would have to bite off the entire project or let it drop. To Fockler's delight, Lodge gave his proposal the green light: "This is the right thing to do," he announced.

"I got what I asked for," Fockler said, "and went out quite happy." He received the $400,000 in seed money to be spent over the next three years. After that, NetNorth would have to be self-supporting. The cash was for connecting Guelph to UNB in Fredericton, with another line connecting Guelph and U of A in Edmonton. Lines connecting institutions to the eastern and western nodes were also partially funded where necessary. IBM was to provide technical support for the 4341 and would cover some of the administrative costs. All told, the NetNorth project would cost IBM a cool million dollars.

CREATING THE NETNORTH CONSORTIUM

While IBM was preparing to part with its grant, the NetNorth founders were preparing an organization to accept it. Although the network did exist, it ran without the support of an administrative entity. NetNorth was a joint project between an unfederated group of university colleagues – on paper there wasn't really anything or anybody to give money to.

Despite their lack of formalized ties, the participants could not have taken NetNorth more seriously. NetNorth had grown out of existing budgets, using existing staff. It wasn't a networking experiment; it was a network meant to increase effective communication and collaboration among faculty, staff and students at Canada's universities. The sense that NetNorth was both useful and feasible was accepted as self-evident. Meanwhile, the founders had been too busy building the network to build NetNorth as an organization. Management of the network had been left to a relatively local group of participants who gathered primarily to work out technical and logistical issues. Almost incidentally, they had also dealt effectively enough with managerial ones.

But these were pragmatic men, and IBM's arrival signaled a need to formalize NetNorth for a number of reasons. IBM needed names to put down on its contract with the organization: NetNorth had identified U of A, UNB and Guelph as the ones responsible for administering the funds on behalf of their regional colleagues. A fair cost-sharing formula had to be put in place that would make the network accessible to potential members, affordable for existing members and financially independent after three years.

Given the scale of IBM's involvement, there had been expectations that IBM would want to play a role in managing the network and making decisions on how to run it. In fact IBM's intention all along was to provide corporate support only, and to let NetNorth handle its own affairs. Nevertheless, NetNorth was mindful of IBM's significant investment and realized that to continue operating NetNorth on an ad hoc basis as it grew in size would ultimately threaten the dependability of the network.

In 1985, representatives from the member universities flew from across the country to Queen's in Kingston to address these issues. As Watt describes it, one of their major tasks was governance:

> It was an organizational meeting where we laid out the plan to develop an administrative structure for the thing we had decided to call the NetNorth consortium... What we ended up with was a structure consisting of an executive committee headed by the secretary of the executive committee and an administrative committee headed by the secretary of the administrative committee. We believed strongly in not picking titles that anybody with an ego would want to have.

The reps also formalized the details of the network's physical topology. In Ontario, the major node was Guelph, which housed the administrative centre and the link to BITNET, both managed by Percival. York would act as the major point of interconnection in Toronto. McGill supported a direct connection from Guelph and was the centre of a straight star network connecting institutions in Quebec. UNB already connected institutions in its province, as well as maintaining links to Charlottetown (UPEI) and St. John's (Memorial). Dalhousie in Halifax would be Nova Scotia's hub.

The western portion of the network started as a direct link to U of A in Edmonton, which served as the centre of a star network connecting the western universities. Eventually, however, once the various Prairie partners were fully up and running, the Prairie part of the network did end up very similar to the train track that Fockler had imagined. Simon Fraser University in Burnaby became the central node for the BC universities since it was one of the few western universities that used IBM equipment and software in its computer services department.

COMING FULL CIRCLE

The consortium had leased most of the network's inter-city lines from CN/CP, and comparisons between NetNorth and the CPR's transcontinental line had become common over the course of NetNorth's construction. But the number of historical parallels at work here would have surprised even the network's builders.

Towards the end of the 1870s, the federal government initiated two national networks considered critical to cementing Canada's claim to its territory, as well as its status as a sovereign political entity: the telephone system and the railway.

In 1880 the Bell Telephone Company of Canada was granted a charter to encourage its western expansion and create a coast-to-coast telephone system. The company was given sweeping powers to construct lines along any and all public rights of way (without municipal or provincial interference); interconnect with all existing telephone and telegraph networks; build its own equipment; and buy out other telephone companies.

The Bell charter complemented recently concluded plans to build Canada's national railway line. The CPR already owned and operated Canada's telegraph system and as the two companies moved westward they were expected to work together. The CPR would provide the long lines between local networks, while Bell furnished the telephone equipment at railway facilities and provided the railway with free long distance service. These intertwined networks, along with some of AT&T's long-line facilities in the US, acted as Canada's first national communications backbone. The cooperative arrangement lasted for nearly 50 years, before being replaced by a series of interconnection agreements. In 1931 it became the Trans-Canada Telephone System, which was renamed Telecom-Canada in the late 1970s.

Over the years, the two national rail systems had gradually pooled their telegraph facilities as CN/CP Telecommunications and began to take an increasingly competitive posture towards the telephone companies, their one-time allies. It was not until 1979, however, that the CRTC finally allowed CN/CP to compete directly with the incumbent phone companies in a number of lucrative commercial areas, including private-line facilities and long-haul data transmission.

When NetNorth needed to buy bandwidth to become a fully operational supplier it went to CN/CP. By doing so, NetNorth brought national networking full circle. It was building Canada's third-generation network based on computer communications using corporate facilities that had been used to create the first-generation network, the railway. CN/CP was later reorganized and now operates as AT&T Canada, partly owned by the US parent company that acted as the third partner in the original creation of Canada's national telephone system.

ALL POINTS EAST MEET ALL POINTS WEST

Towards the end of 1985, NetNorth was close to becoming truly national in scope. The meeting point would be in Alberta, where all points West would connect with all points East. The compelling analogy between NetNorth and the railway was strangely prescient. Coming up was November 7, 1985, the 100th anniversary of the railway's completion when the last spike was driven at Craigellachie, B.C. It was as good a date as any to mark their own national network's opening ceremonies. Fockler recalls the occasion:

> The University of Alberta did a great job of promoting what was going on. We had some press, although the press didn't understand the concept that much, but U of A had made a great backdrop. It was a bigger than life-size version of the famous picture of Lord Strathcona driving the last spike. And they had a terminal connection that displayed the transmission of a message that had been sent from Memorial University to the University of Victoria.

U of A had also had a large cake made for the occasion decorated with a map of Canada. Fockler, perhaps the most intimately aware of the geography that the network covered, realized to his dismay that one province was missing:

> Maybe it was the wrong thing to do, but I pointed out that Newfoundland had been left off the map of the cake! So somebody took a spoon, scooped out a hunk of the Arctic and plopped it down in the Atlantic Ocean somewhere a little off Nova Scotia, around where Sable Island is, then said "There it is!"

Over the next three years NetNorth continued to grow: 21 members in 1985, 40 in 1986, 57 in 1988, 65 in 1989. The network was a success for at least two reasons. First, the business model was working: the network had become financially self-sustaining well before the three-year funding ran out. Second, people were using it. In fact, it was quite a hit. By 1988 the network's 4341 minicomputer hub was heavily involved in file transfer and processing more than 65,000 pieces of email every day.

The NJE network that NetNorth became part of was popular and global but it wasn't the Internet as we know it today. In terms of presentation, monochrome and text dominated computer screens. In terms of bandwidth, the fastest sections of the main Canadian trunk line ran at 9600 bps. As a contemporary point of comparison, a subscriber in 2001 using Sympatico's DSL service (the High Speed Edition) for home Internet access can download data at a nominal 960 Kbps (960,000 bps), a hundred times faster than NetNorth.

These future possibilities mattered little to the NetNorth population, who were thrilled with what they had. The network made a deep impression on Roger Taylor, who was then working out of a lab in Ottawa in the physics division of the National Research Council:

> I was conducting research with someone in England over NetNorth and BITNET and it was really working very well. Such an enormous improvement over trying to do things by conventional mail and phone calls. Just being able to ship messages and get an answer back within minutes... it blew me away. To be able to carry on a conversation with someone... despite all the problems of going through a gateway at the Rushford Lab in the US and switching to the protocol that [the British network] was running in those days. Yes, going from NetNorth to what we have now is nice. But it hasn't made the same difference to my life as that step from no net to NetNorth.

4

CANADA CATCHES UP

"WE SIMPLY FELT IT WAS IMPORTANT TO ESTABLISH THE INTERNET IN CANADA ON A CANADIAN NETWORK." – *GERRY MILLER*

NetNorth and CDNnet had succeeded in networking more than computers. They had brought into existence a well-defined community of networking specialists, one that identified itself as Canadian in scope. The network of communications they had erected to serve academic research now also served to manage and develop their computer network on a national scale.

The NetNorth consortium was to play an especially critical role in this regard because its members had shared the experience of building a network from the ground up. The fact that members leased their own lines for connecting to the network and managed it in a decentralized, ad hoc manner meant they had a strong proprietary interest in the status and management of the network. At the same time members were also committed to helping each other: the NetNorth consortium was developed on a cost-sharing basis in order to extend the network's infrastructure as broadly as possible. This had been a critical means to overcoming Canada's all-too-familiar barrier – too few people spread out over too much geography. Overall, therefore, the network wasn't just a facility; it was a cooperative effort that required members to participate actively in order to make decisions.

Against the background of these positive achievements, however, was a growing awareness of the technological limitations imposed on NetNorth and other networks linked into the global NJE network by its store-and-forward design, which was still based on circuit-switching. Because a file was sent out as one lump, its size affected not only its own travel time but the travel time of other files that got caught behind it in a queue. The first-come-first-served system could create serious bottlenecks if a few sizeable files got stacked on top of each other.

Chat applications had a particularly bad reputation for clogging up the system. The chat application simulated real-time communication by sending small amounts of data, typically one line of text at a time, to the head of the queue for quick dispatch. Hundreds of simultaneous conversations over the NJE network would seriously clog the progress of larger files, and its users had to be depended on to exercise self-control.

Not only was interoperability between base protocols a problem, but the NJE email applications had to learn to read other addressing systems. David Macneil experienced this when he tried to email back to UNB while on sabbatical at Stanford: "I remember exchanging messages and finding it was kind of catch-as-catch-can as to whether they would work." Another time a message with an especially large subject heading sent from UNB caused a computer crash at CERN, the European Laboratory for Particle Physics in Geneva (the birthplace of the World Wide Web in 1989-90).

The highly decentralized nature of NJE network management meant that every node calculated its own distribution path, theoretically in conjunction with other computers using the same map of coordinates. Sometimes, however, the network computers didn't or couldn't cooperate. Messages could die mid-route for a number of reasons: the network map was out of date, one of the computers was offline or simply turned off. As the network expanded globally, it was also the victim of cultural gaps. For example, computer centres were shut off every Friday by most of the members of GulfNet in the Middle East. On a global network, it was always a holiday somewhere. Ultimately the patchwork nature of the NJE networks reduced their accessibility, reach and overall quality of service. Data transmission was limited by the state of communication between networks as much as within these networks.

NetNorth represented a giant step forward in networking computer services and facilitating research for many academic departments. But its drawbacks left computer science departments and other large technology-oriented facilities in government and private industry – especially those not blessed with military contracts – still hankering for the ARPAnet. Having an ARPAnet connection was a mark of prestige and an important differentiator between first – and second-class facilities. As a result, many constituents in both Canada and the US felt they were being left behind. In response, a number of "alternative" networks sprang up in various high-tech communities.

The first of these was CSnet, the fruit of efforts by the US National Science Foundation (NSF). In 1980 NSF joined hands with DARPA and representatives from the US university community to organize and implement a computer science network using packet-switching technology, which allowed networks such as ARPAnet and CDNnet to avoid the congestion of the NJE networks. Led by Larry Landweber at the University of Wisconsin, the initiative resulted in CSnet (the American equivalent of CDNnet), which used components of the ARPAnet, the X.25 protocol and regular dial-up phone lines to link virtually the entire computer science community.

In 1985, the NSF went on to fund NSFnet, which linked up five supercomputer centres in the US. Because it was explicitly designed for non-military, academic use by universities, the NSFnet was accessible to a much wider community of users than ARPAnet. Most importantly, in a decision that was anything but a foregone conclusion, the NSF chose the TCP/IP protocol suite as the platform on which NSFnet, and eventually the entire Internet, would operate. With NSFnet in place and functioning, it replaced ARPAnet as the focus of Canadian efforts to extend their networks south of the border and internationally.

BCNET: THE FIRST PROVINCIAL IP NETWORK

British Columbia has been a locus of computer networking for almost as long as there were computers around to network. Starting from the initial connection to the Michigan Terminal System (MTS) network, whose mail relay network also provided a gateway to BITNET, both UBC and Simon Fraser University (SFU) had expanded their networking capacities considerably in the early 1980s. At UBC, the cyclotron facility (TRIUMF) established an early connection with the Energy Sciences network in the US, providing the facility with connections to CSnet, BITNET and high-energy physics labs throughout the US and Europe.

CONNECTING TO NSFNET

UBC had enjoyed a direct link to CSnet since 1983 using an X.25 connection. When CDNnet came online in 1984, UBC provided a gateway for users of the new network and CSnet, which in turn provided a gateway to a growing number of international facilities. UBC's southern link was thus serving as the principal gateway for Canadian computer science departments wanting to reach colleagues in the US and overseas.

While the CSnet connection was a useful resource, it had its limitations. For one, it had the now familiar drawback of being an unpredictably expensive service. Second, the US-Canada public data network was even more expensive and slower than the Canadian Datapac service, because Bell and Sprint's Tymnet, its American counterpart, used different versions of the X.25 standard. This required a gateway based on yet another standard (X.75) to connect them. "The result," as John Demco puts it, "was a lowest-common-denominator connection."

In 1986 Demco was attending a networking workshop in Dublin. He was chatting with Landweber when the CSnet founder made an intriguing suggestion: get a leased-line connection to NSFnet. Demco hadn't realized that the NSF would even consider accepting an international link, but he was riveted by the idea of getting UBC connected. Steve Wolff, then director of NSF's division for networking and communications research and infrastructure, was also at the conference. When Demco brought up the subject, to his surprise Wolff seemed entirely receptive.

But when Demco returned home, what seemed straightforward became increasingly complex at every stage. Although the acceptable use policy for the NSFnet had indeed been loosened, it wasn't like signing up to an ISP. Arranging a cross-border connection was onerous. The technology was a mystery to most people. Routers weren't necessarily interoperable. In the end, CDNnet made arrangements with Northwestnet, the regional IP network run out of the University of Washington. Demco had a great deal to learn and during the entire process he was still running CDNnet, which continued to grow dramatically.

Finally, the inevitable issue of funding had to be tackled, and in fact wasn't dealt with until quite late in the process. Demco:

> There wasn't any ground swell of demand to connect to the Internet yet. There might have been more from non-computer science departments if they had known about it. And the people who really needed the network, well, they already had it in some form or another.

The connection would have to wait for a bigger boost from a new network with a broader funding base – BCnet.

MEANWHILE, BACK AT COMPUTER SERVICES

By 1986, UBC was running or participating in at least four computer networks: the MTS mail relay network, CDNnet, NetNorth and CSnet. In addition the university had close working ties with the University of Michigan, which was now part of the consortium running the NSFnet.

Through his Michigan contacts, Ed Froese, director of the UBC computing centre at the time, had been observing the growth of regional IP-based networks in the US. Froese was among the first to urge UBC to champion a network, patterned on the US regional IP networks, which would be linked to other regional networks across the country. In 1986 he and his colleagues from computer service departments in other BC universities and colleges secured start-up funding from the BC government to establish what eventually became BCnet.

Initially the networking partners planned to use high-speed wireless connections to link their campus networks. They had received an estimate that was far lower than the cost of leasing equivalent bandwidth from the telephone company. Unfortunately the process of securing a licence for the wireless technology got held up by red tape, as well as by jurisdictional clashes between the provincial and federal governments.

Finally the BCnet consortium ran out of patience with the Department of Communications. Jack Leigh, then associate director of the UBC computing centre, remembers his frustration:

> We lost eight or nine months trying to use the wireless links waiting for our prospective supplier, Pacific Digital, to get federal approval. Somewhere in my files I've got a letter from the then Minister of Communications that says: we're not against what you want to do, but it doesn't say yes, doesn't say no and indicates that they'll probably think about it for another couple of years.

At that point the organizers decided to exercise other options. Cost was still a serious consideration, so only a single line was leased to reach the University of Victoria (UVic) from the mainland. But to connect SFU and UBC, both located in the Vancouver area, the universities were able to use existing cable facilities built to deliver video services such as distance learning but largely unused. This cable network ran throughout as well as between the campuses.

The final network turned out to be a combination of telephone and cable facilities using Ethernet, DECnet, X.25 and TCP/IP protocols. "We didn't have a whole bunch of money, so we had to use what we had," as Leigh puts it. So by October 1987, one thing they did have was Canada's first provincial IP network. The initial participants included UBC, UBC's cyclotron facility TRIUMF, SFU and, on Vancouver Island, UVic, the Dominion Astrophysical Observatory in Saanich and the Oceanographic Institute at Patricia Bay. In December of 1987 BCnet came online as an inter-organizational wide-area network and was still working out the details of its connection to NSFnet.

As timing would have it, while the newly established BCnet was looking for other Canadian regions to connect to, Demco and CDNnet in the computer science department down the hall were still looking for ways to establish a direct link across the border to NSFnet. So BCnet and Demco reached an agreement. CDNnet gave up its X.25 connection south and instead leased a 19.2 Kbps line to Washington state. As a result, CDNnet members who previously had to pay separately for their use of the CSnet gateway now had the cost covered as part of their regular CDNnet dues.

Although CDNnet was growing quickly, the permanent connection to the TCP/IP Internet through BCnet ironically foreshadowed its eventual demise.

As organizations like BCnet made direct Internet access increasingly feasible, the use of intermediary networks such as CDNnet became less necessary. Moreover, applications like EAN, designed to work with multiple underlying protocols, became increasingly irrelevant as TCP/IP networks became more pervasive.

ONET: ONTARIO'S HARDWARE BLUES

The University of Toronto (U of T) and the rest of Ontario's universities (much like Toronto and the rest of Ontario's cities) share something of a love/hate relationship. Because U of T is significantly larger than any other Ontario university, it inevitably controls what appears to be a disproportionate amount of educational resources – including computer resources.

Once NetNorth was established, the academic community in Ontario hoped that the network would help distribute these resources for all to share. NetNorth did address this problem, but with mixed results. For one thing, the NJE store-and-forward protocol placed significant limitations on the kinds of computer resources it could make available remotely. Second, the NJE protocol had been designed to work with IBM hardware and, although software work-arounds were available, universities with hardware from other manufacturers were at a disadvantage.

THE NUMBER-CRUNCHING CRUNCH

These problems were a source of great frustration to Andy Bjerring, who was appointed Director of Computing and Communications Services at the University of Western Ontario in 1985.

The department was brand new – the first in Canada with computers and communication as its prime mandate – and Bjerring himself was beginning a new career (his third). Leaving a position as a systems-control engineer in the pulp and paper industry, he had enrolled in Western's philosophy of science PhD program in 1972. After graduating he taught philosophy of science at Western before moving to a position in the President's office helping to plan and budget university resources.

Bjerring was appointed director of the new computer and communications department just after NetNorth and CDNnet were established and growing coast to coast. Although Western was part of NetNorth and found the network an invaluable communications and file-sharing tool, the university's computer technology was actually better suited to a network like CDNnet. Instead of sticking with the usual IBM equipment for cost and efficiency reasons, Western had pursued a multi-vendor strategy because, as Bjerring explains, its emphasis was on high-performance computing:

> The distinction was between a time-shared system which could handle multiple users all on a single mainframe environment, where you were providing support to students whether they were in computer science or engineering or whatever, versus a number-crunching shop, where you would have a smaller number of users but the jobs that you were running were very, very large jobs, required a lot of disk space for swapping big files in and out and very fast processors.

The result was a lot of computers made by different manufacturers that could not talk to one another:

> We had everything from Control Data, to Digital Equipment, to VAXes and old PDP-10s, to IBM gear. We had just about everything you could name and none of them talked to each other. And so we had very little opportunity to create a network, even among the mainframes. Let alone make it really easy and convenient for users that wanted to switch between the mainframes. Using a switched network environment, they basically had to log off one mainframe and then log onto the next mainframe. It really reduced the opportunity for effective utilization of the strengths of each system. It went to zero.

Although especially acute at Western, this problem was hardly unique. By 1985 solutions for heterogeneous local computer networks were being developed, and Western itself had begun to implement Ethernet, developed at Xerox PARC 12 years earlier. Ethernet is a fast transmission standard – typically 10 Mbps in an office environment, but newer versions can deliver data as fast as 1,000 Mbps. It is also robust, i.e. computer hosts on the network can malfunction without disrupting network operations, and it can be used by almost any computer. Ethernet very quickly overshadowed other local area network (LAN) protocols such as the Token-Ring networks being promoted by IBM.

The growth of Ethernet-based LANs at universities was quietly creating thousands of small computer networks across the continent, and dozens across the province. The development gave new urgency and opportunity to the problem of interconnecting these networks and the increasing number of resources available through them. But Ethernet solved only part of the problem: even if all of Western's computers overcame their incompatibilities, Western and other universities with similar environments and requirements still needed high-level Control Data or Cray number-crunching machines – hardware on a scale they just couldn't afford. In Ontario, only U of T had that kind of computing power.

ITRC TO THE RESCUE

U of T had no desire to hoard its Cray supercomputing facilities but NetNorth couldn't support the interactive use of remote facilities connected to them. And the fact was, as Bjerring says, that even with its considerable means, U of T could barely afford to maintain the Cray itself, let alone facilitate a network to help distribute access to it:

> They agreed that network access was critical but it wasn't their job, and even so they ran out of money and had to go back to the [provincial] government for more.

In that same year, however, the government of Ontario formed a number of Centres of Excellence with the mandate to fund, promote and facilitate scientific research and development between Ontario universities, industry and government. Bjerring recalls:

> The government, among other things, made it clear they needed to make the computer facilities more Ontario wide. So they changed the name of the U of T facility to the Ontario Centre for Large-Scale Computation and expanded the board of directors and I became a member of the board. But there was still no money for networking.

Another of these centres, the Information Technology Research Centre (ITRC), had its facilities distributed across campuses at Waterloo, Toronto, Queen's and Western. The ITRC, and especially Zenon Puilyshin, one of its main organizers, saw a clear need to establish an IP-based network for the kind of resource sharing Bjerring was advocating. Puilyshin was located at Western and often employed the services of Peter Marshall, one of Western's senior networking staff, as an advisor. Marshall and Puilyshin discussed establishing a network, and Marshall in turn spread the word among his IT colleagues at Waterloo and Toronto.

Meanwhile, the ITRC was already developing a computer science network to be operated by the Centre. But at a meeting in Waterloo late in 1987 representatives from the university IT communities pitched the notion of making the network broader in scope, with the ITRC playing the role of "sustaining member" (that is, it would pay a disproportionate amount of the costs).

John Wong, head of the ITRC, agreed and eventually a funding partnership was worked out among six universities (Western, Waterloo, Toronto, Guelph, McMaster and Queen's); the federal Ministry of Industry, Science and Technology; and another Centre of Excellence named the Institute for Space and Terrestrial Science. Even the cash-strapped Ontario Centre for Large-Scale Computation found funding to double the proposed network bandwidth from 9600 bps to a full 19.2 Kbps.

While these negotiations were taking place, U of T had been looking at establishing a connection between its campus network and the TCP/IP networks flourishing south of the border. In November 1987, Eugene Siciunas, assistant director of U of T's computer services, received the go-ahead from Cornell University to link U of T's computer centre to the NSFnet through Cornell. The cost to lease the bandwidth they wanted, a 56 Kbps line, would cost $80,000 a year. Siciunas thought that a bit pricey, even for U of T:

> So we were looking for ways to share the cost. We knew that NetNorth already had a 9600 baud link to Cornell and our computer science department had a link to the ARPAnet through Rochester at 9600 baud too. We proposed buying the 56 Kbps link, multiplexing it into two channels and giving a 19.2 Kbps channel to NetNorth for the same price they were paying for the 9600 baud line – double the bang for their buck. The remaining bandwidth would be an IP link shared by U of T and the computer science department.

Since U of T was a key link in both NetNorth and the proposed provincial IP network, it seemed reasonable that the cost of the desirable NSFnet connection should be shared by both networks. Finally, in August 1988, the first links of the new IP-based network christened ONet were established between Western and Waterloo. In October U of T finally got its connection to the NSFnet, becoming NetNorth's new link to the US – and ONet became Canada's second "live" regional IP network.

RISQ: AND QUEBEC MAKES THREE

Faced with the same concerns as their colleagues in BC and Ontario, Quebec's academic and R&D communities had the basic resources required to build an IP network among member institutions, as well as a link south to the NSFnet. But Quebec also had the unique advantage of precedent and momentum, created by a networking predecessor based on a technology that was originally intended to compete with IP, not encourage it.

AHEAD OF THE CURVE WITH CRIM

A prime mover in the founding of Quebec's regional network was Bernie Turcotte. Graduating from McGill in 1980 with a B.A. in mathematics and a minor in computer science, Turcotte held various positions in networking, including a short stint programming network interfaces, before becoming network administrator at the Business Development Bank of Canada.

From its headquarters in Montreal, the Bank supported a national network of branches. Turcotte was involved in establishing a computer network that would provide remote processing through these branches, as well as a variety of applications in what became a relatively sophisticated networking environment for its day. In doing so he oversaw the bank's transition from one singlevendor networking solution to another – from an all-DEC environment to an all-IBM one, “wall-to-wall Big Blue” as Turcotte puts it.

Turcotte enjoyed the opportunities provided by the private sector and recalls that “in the early 80s a programmer could change jobs every second week if he wanted to.” But he preferred the more creative, advanced projects in the academic and R&D communities. In 1985 Turcotte was offered a position with a newly formed non-profit computer research institute associated with five Montreal universities and supported by private industry. The Centre de recherche informatique de Montréal (CRIM) had a mandate to create and fund post-graduate research opportunities for the increasing number of computer science students graduating from Quebec’s universities. Turcotte was CRIM’s fourth employee, and its first computer expert. In less than a year, the institute successfully launched one of the most sophisticated inter-university networks in Canada.

The opportunity to build the network was made possible because of a large donation from DEC, which provided most of the funding and equipment necessary to build a 56 Kbps network connecting CRIM, its associated universities (Université du Québec à Montréal, Université de Montréal, Concordia University, École Polytechnique de Montréal and McGill) and SIRICON, a separate research centre at Concordia.

The network operated using DECnet protocols and routers. Unlike NetNorth’s store-and-forward approach, DECnet was an interactive, packet-switched protocol. However, unlike TCP/IP, it was a proprietary technology developed and owned by DEC, and primarily compatible with DEC-manufactured equipment.

The CRIM network was greeted with great enthusiasm. Not only did it feature an advanced form of networking, but through DEC’s largesse, CRIM could afford to lease its own 56 kilobit lines. As Turcotte puts it, “back then, a 56K communications channel was like gold, and it was incredible that someone would ever offer this to you.”

There was only one hitch: DECnet was not compatible with IP, which was a source of great frustration to users such as McGill electrical engineer Michael Parker, better known by his computer handle, der Mouse. Parker was, and remains, well-known in computer circles as a UNIX guru, especially on his home turf of Montreal. And UNIX speaks IP, not DECnet, even though DEC had in fact developed its own flavour of UNIX known as Ultrix. Parker recalls that each CRIM-connected university had at least one Ultrix machine on hand:

> At the time the Montreal universities were islands: IP-connected within themselves but connected with one another by DECnet-only links. It was frustrating to have perfectly good lines connecting the universities but be unable to use them because the routers were too stupid to speak IP as well as DECnet. The obvious answer was to encapsulate IP in DECnet, which is more or less what I did.

The process of encapsulation allows computers to communicate by embedding one kind of protocol within the packets of another: DREnet used this technique to send IP packets over an X.25 network. In CRIM's case, Parker sent IP packets over a DECnet network, which took a fair bit of doing. Because Ultrix could speak both DECnet and IP it could, in principle, do the encapsulation. However, without any documentation or source code to work with, Parker's efforts to write the encapsulator required a great deal of creative, educated guess-work and persistence.

His success ultimately threw open the floodgates, allowing many more computers to connect to the network. In September 1987, CRIM established a 9600 bps connection with the CSnet, giving it access to a vast array of organizations that were part of or connected to that network.

THE BIG RISQ

In late 1986, the National Research Council (NRC) was ready to announce plans to fund the construction of a national IP-backbone. Turcotte watched with interest the meetings hosted by the NRC in Ottawa and by U of T in Toronto in 1987. But later in the year the general feeling was that negotiations were at a standstill and the project had somehow become derailed. This only intensified the desire of Quebec's universities to develop IP networking on a broad scale themselves. Turcotte again: "The writing was on the wall that if a university was a real research university it needed to have networking."

The instigators of the new initiative were the university computer-centre directors acting through their working group, the Regroupment des directions informatiques des universités du Québec. Taking their lead from the networks established in the early 1980s by UQ and the Ministry of Education, the group had discussed as early as 1983 the possibility of building a private inter-university network using the X.25 protocol. This "UNIPAC" network got as far as sketches on paper and principles in outline but by the late 1980s the popularity and utility of TCP/IP were undeniable. So the working group and CRIM decided to abandon their existing network technologies (X.25 and DECnet respectively) and committed to using the Internet protocol suite for the new network to be called the Réseau inter-ordinateur scientifique québécois (RISQ).

The initial funding came from CRIM. Turcotte says that "CRIM had a lot of money in those days, and it was supposed to do "interesting research things." They were able to convince CRIM that this new network qualified as a "very interesting research thing." RISQ was launched in September 1989 using 56 Kbps lines linking its founding members: the CRIM network, UQ head office, École des hautes études commerciales à Montréal, Université Laval, Université de Sherbrooke, the Atmospheric Environment Service in Dorval and Hydro-Québec. The network operations centre was located at McGill. That same month RISQ established a link to the NSFnet through the New York State Education and Research Network.

The fall of 1987 through the summer of 1988 saw the establishment of two provincial IP networks (BCnet and ONet) and the first steps towards founding a third (RISQ). Depending on one's point of view, this was either a sign of progress (the first implementation of an important networking technology) or a further unraveling of the vision of a unified and autonomous Canadian networking infrastructure. Either way the development of these networks in Canada's three most populous provinces hinted at the disparities between them and provinces that could not be connected without cooperative action at a national level.

THE NATIONAL RESEARCH COUNCIL

In 1986, the NRC celebrated its 70th birthday as one of Canada's most distinguished scientific organizations. Frederick Banting had worked for the NRC. Its staff had cooperated with English scientists to develop radar during the Second World War. Gerhard Herzberg, the NRC's director of pure physics, had been awarded the Nobel Prize for chemistry in 1971. The organization currently employs over 3,000 highly trained personnel, maintains research labs in every region of Canada and receives nearly half a billion dollars in government funding annually.

That same year, Roger Taylor was appointed director of a newly formed division known as "Informatics," a coinage derived from the French referring to computer communications. But rather than being able to work within the traditions of this venerable scientific institution, Taylor had to adopt an ad hoc approach that had more in common with Usenet's seat-of-the-pants consensus. The surprise in store was just how compatible these two legacies could become.

Since the NetNorth and UUCP links joining the NRC's laboratories (most of which were located in Ottawa) were no longer suitable, one of Taylor's major tasks in his new position was to upgrade and improve the networking facilities.

Taylor's first step was to sit down with John Curley, whom Taylor had appointed networks manager, and plan the proposed network:

> We were supposed to link all the NRC labs. I have a feeling when they told us to link all the NRC labs, they really meant just those in Ottawa. We interpreted the assignment literally as all the labs. That means Victoria to St. John's. As we discussed this, neither of us can remember who said it first, but we decided that if we're going to do a coast-to-coast network, why not do it properly – involve all of Canada's research establishments.

Some months after Curley and Taylor decided to build their NRC network on an unexpectedly large scale, they received some interesting news from a mutual colleague. Andy Woodsworth had just learned that his lab, the Dominion Research Observatory of Astrophysics on Vancouver Island, had been chosen to become the official repository of the digitally stored data captured by the Hubble Telescope. Peter Shames, the NASA representative attached to the project, was also well known within networking circles because of his involvement with NSFnet. In April 1987, Curley and Taylor arranged through Woodsworth to meet Shames and Steve Wolff of NSFnet in Baltimore.

NSFnet was an obvious model for Curley and Taylor's project, and the two of them did a thorough job of picking the brains of their American counterparts. They were also testing the waters and watching how Shames and Wolff reacted to the proposed network, which would depend heavily on a connection to the NSFnet. Their response was more than positive – it was downright enthusiastic. To have Canada build its own TCP/IP backbone was good news for a number of reasons.

For one, it meant that traffic moving from one location in Canada to another would most likely stay in Canada, advantageous from the perspective of an already badly overloaded NSFnet. It also meant the US would have excellent compatibility with their networked colleagues in Canada on such projects as the Hubble Telescope initiative. Finally, it would represent a significant addition to the growing number of interconnected IP networks, further validating the NSF's choice of networking protocol. The enthusiasm and cooperation of the NSF staffers would prove useful to the NRC network on a number of occasions.

Taylor and Curley had a month to plot out the case they would have to make to the NRC management committee, which would meet in May to set the spending agenda for the rest of 1987. The scope of the project would no doubt come as a surprise. It was a far cry from the internal networking project the committee members were expecting, since the proposed network was intended to serve Canada's research community as a whole. It added an element of publicly visible risk to the project.

On the other hand, the two researchers knew they wouldn't have to restate the case for the value of computer networking. The NRC had been part of NetNorth for several years and had found the service invaluable. As to Canada's need for an IP network backbone, the advantages of having a real-time, interactive, packet-switched network weren't hard to demonstrate. TCP/IP could run circles around the older store-and-forward networks, allowing for remote log-in, database access and computer resource sharing on a much larger scale than was possible with circuit-switched networks.

But was building such a network the job of the NRC? Since the mid-1980s the organization had been shifting its mandate towards research collaboration. The government agency wasn't supposed to carry the mantle of scientific progress on its shoulders alone anymore. Its job was to stimulate the larger scientific community in industrial and academic circles. Curley and Taylor's networking project filled the bill nicely. It would require broad-based support from the research community and the benefits from it would be distributed on an equally broad basis.

YELLOW LIGHT FROM MANAGEMENT

NRC's management board was guarded in its reaction. While they didn't veto the project, they wanted to see more details. "They weren't quite sure what they were letting themselves in for, but no one tried to rein me in," said Taylor. He and Curley then formed a network committee and began to build support. Although Taylor had come directly from a pure science background to his position as director of informatics, he and Curley navigated the obstacle course of the federal bureaucracy with the instincts of seasoned pros.

They needed three groups to get on board. The first was the existing networking community, essentially drawn from the computer service and computer science departments on university campuses across Canada. This wasn't a tough crowd to win over. The gossip that NRC was cooking up a national backbone had been circulating through the community networks for months.

On the other hand, Taylor knew his window of opportunity with this crowd would be limited – not because they would lose interest, but because regional efforts might pass the NRC by. Provinces with the means to arrange their own connection to the NSFnet were already doing so. Once these provincial efforts were established, would they agree to join a speculative east-west networking initiative based on a cost-sharing model designed to support their less conveniently located colleagues? Would BC want to share the cost of connecting Newfoundland? If provincial governments got involved, how would jurisdictional issues get played out?

The second group to be convinced was the NRC management board. They needed to approve the project in principle as well as its budget. The NRC wouldn't support the project indefinitely or, as Taylor puts it, "Nobody would have bought into something that would appear to be a black hole you had to keep pouring money into." Whatever money the NRC gave would have to serve as a catalyst to get the network up and running long enough for it to become self-sufficient, after which it would be supported by user fees. How much money that required was still a big question: Taylor was looking at a ballpark figure of one to five million dollars. Furthermore, the issue of how hands-on the NRC should be was hanging in the air. Would the network be managed centrally or on an ad hoc basis like NetNorth? Would the NRC operate the network or would it only fund it? The answers would go a long way to determining the management board's final decision.

Third and last, the NRC had to convince the federal government at large, and Treasury Board in particular, that this was both a project worth pursuing and an appropriate one for the NRC. Clashes with other departments that felt their toes stepped on wouldn't be constructive.

BACK ON THE PLANE

Timing was important on all these fronts. After receiving tentative approval – or at least no direct disapproval – from the management board in the spring of 1987, Taylor and Curley took the first steps. The network committee, chaired by Curley and including Morven Gentleman, Art Hunter and Andy Woodsworth, organized a meeting of Canada's "networking gurus" in late October 1987.

The committee flew them in from all over Canada for the meeting. Henry Spencer represented Canada's loosely affiliated Usenet community (another attendee admitted to Spencer that UUCP links were still the easiest way to get on the Net). John Demco was there as CDNnet's representative. A Software Kinetics representative attended on behalf of DREnet. Paul Dirksen, head of Waterloo's computing services, spoke for NetNorth. David Macneil wore his NB/PEI networking hat for the occasion. Also in attendance were observers from NSFnet, including Peter Shames, and representatives of Canada's nascent regional networks in BC, Ontario and Quebec.

Along with this group of networking celebrities were representatives from the NRC's management board and the federal Treasury Board. Taylor hoped the enthusiasm of the assembled group would be infectious. He asked Ross Pottie, the NRC's vice president, to open the meeting with a few words and stay on for the meeting. He did and was charmed by the good will and energy in the room. The group quickly moved on from wanting to go forward in principle to debating details of network management. The NSFnet crew strongly recommended a centralized management approach: they had tried to manage the network in an ad hoc fashion and the experience had been a disaster. They then handed the management of the network over to a three-member consortium made up of IBM, MCI and the MERIT group from the University of Michigan.

Several days after a successful meeting, Taylor spotted a newspaper article about a proposal to create centres of excellence around the country that would promote the development of computer technology:

> I took that thing out of the newspaper and walked to Ross Pottie's office, banged it onto his desk and said, "Ross, look at that. The network we're proposing will make that happen." He looked at it and said, "Yes, you're right." Pottie was sold on the idea.

In the meantime, Taylor had also hosted an informal meeting with the Government Telecommunications Agency to make sure the project wasn't going to ruffle anyone's feathers in that department. The agency agreed that the NRC was the most appropriate agency to lead this kind of project. The network committee then submitted its initial report in November, and got immediate approval from the management board. But this was official approval for the concept. The devil was in the details, and the NRC had still heard very few. How much would it cost? Who would manage it? When would it be ready?

The network committee couldn't give them the answers, and they clearly needed help putting together a business plan for the project, by now dubbed the National Research Computer Network, or NRCnet (the name had been suggested, half in jest, as a way to work the NRC's initials into the name but it stuck and nobody seemed to mind). In December, John Curley flew out to UBC to hold another round table on the issue. The BCnet and CDNnet crews, about to finalize their own connections to NSFnet, were in attendance as were other members of the local networking community. The meeting was productive. Jack Leigh was especially enthusiastic about the effort and had strong ideas based on his experience in BC about what an ideal network would look like.

Two months later the NRC would meet with the U of T computer services group. There the plans to build a national network would take on an even more concrete tone – U of T would propose that it take the initiative in building and running Canada's version of an NSFnet.

HOG WILD IN HOGTOWN

Neither Warren Jackson, U of T's computer service director at the time, nor Eugene Siciunas, his assistant director, can recall what exactly prompted them to submit their unsolicited proposal to run the NRC's still hypothetical network. Looking back on everything happening around 1987-1988, Siciunas suggests: "It just seemed like a good idea at the time."

By the start of 1988, a number of realities had become apparent to the U of T computer services staff. Everyone knew the NRC wanted to build a national IP backbone. NetNorth was also committed to remaking itself in the image of IP, and U of T was in the process of taking over from Guelph as the NOC for NetNorth. Not only that, the university was playing a leading role in establishing ONet. But so far, no one had stepped up to make a concrete proposal about building the IP network or to put their organization on the line for it. That person would eventually come from within the academic networking community or, failing that, government or industry.

U of T knew it had the facilities, staff and experience to run the proposed network. As Jackson puts it:

> You've got to realize, a bunch of academics like us and our guys are very passionate people when it comes to this kind of thing. This is something that needs doing, we think it's feasible, we're going to do it somehow. That tended to be our attitude. We can tell you exactly how we're going to spend the money – all we need is the money.

Jackson, Siciunas, systems manager Bill Lorison, IT planner Al Heyworth and other members of U of T's systems group worked to develop a paper which outlined the scope and technical details of the project. Having worked with the people at NSFnet to set up U of T's own connection, they modeled the network closely after the NSFnet, both technically and organizationally.

The university had already approached IBM Canada to enter into a joint project whereby IBM would supply the same technology that IBM in the US had supplied NSFnet – the computers that served as routers, directing data around the network. IBM Canada agreed in principle, but didn't want to be seen as favoring one university above any others. It preferred to keep the focus on national collaboration and so insisted that Jackson and Siciunas present their proposal to not only the NRC, but also representatives from McGill and UBC, before submitting officially.

At the February meeting with NRC, McGill and UBC, Siciunas remembers their audience was distinctly "underwhelmed." If the U of T staff had seen the behind-the-scenes activity, they would have known their guests from NRC hadn't intended to be discouraging. In fact, the U of T's presentation gave them a great deal of what they needed to put together their business plan. But too much was still up in the air for them to leap up and declare U of T the operations centre for a network that still didn't exist. Things were certainly getting close, but the last few hurdles – getting final approval and money for the project – would turn out to be the toughest yet.

BEDEVILED BY DETAILS

The network needed a sound business plan to get up and running, and eventually stand on its own financial feet. The funding was intended to serve as seed money to get the backbone built and operational, but the lead time it provided would have to be relatively short. For one thing, the NRC wasn't receiving any new money to build this network; it would simply be dedicating funds from its existing budget. For another, broad public demand didn't yet exist for a national data network, so any kind of long-term subsidy just wasn't in the cards.

Revenues would be generated in a straightforward manner, with the provincial networks paying fees to gain access to the backbone running east-west and south into the US. BCnet's hands-on experience and U of T's unsolicited proposal had given the NRC committee the facts and numbers it would need to make its proposal to the management board.

Constructing the business plan was a speculative endeavor to say the least. Not only did the plan have to take into account the time and cost to build the major network. In addition, its revenues depended on the success of ten other startup networks looking to deliver access to a network most people still hadn't heard of, using technology they didn't understand. Although in retrospect becoming Canada's first ISPs might not sound like a risky proposition, at the time it must have sounded like investing in Florida swampland. Taylor recalls:

> It was an enormous challenge. The regional networks were very enthusiastic to see the backbone project go forward, because for most of them it was their best chance to get connected to the Internet. But at the same time they're saying to themselves: here's the budget, here's the cost to run the network every year and we've got to generate enough revenue to pay for it. They had no idea whether they could or not, and the NRC's seed money would only give them about a year to work up the business, get things going and start charging people.

> It was a gamble. You think to yourself now, how could that be a gamble, it's a cinch! But it was a gamble and I admire their courage in agreeing to take the chance they did.

Drawing on the information at hand and estimating that the cost to run the network for five years would run in the neighborhood of $5.5 million, the network committee proposed that the NRC authorize $2 million in seed money for the project. That would leave the regionals to pick up the $3.5 million deficit by the end of the five-year period. By that time the network was to become financially self-sufficient. The plan was submitted to the management board in early spring, 1988 and was well received. Approval to go ahead was granted.

NETNORTH AT THE CROSSROADS

The fall and winter of 1987-88 was a crucial time for networking in Canada. While the provincial networks were getting established and the NRC grappled with its national project, NetNorth was nearing the end of its funding agreement with IBM Canada.

IBM had served as the patron saint of NetNorth since the network established its first link to Cornell University in 1984. It spent upwards of a million dollars supporting the network, and although IBM was motivated by a desire to be a good corporate citizen, its charitable activities were also an extension of its public relations strategy. Not that the universities weren't themselves interested in the potential of cooperative self-promotion.

NET 87

Ken Fockler recalls these as the good old days, when he had a staff of seven or eight bright people with public relations and marketing backgrounds to help him get things done. For example, in 1986, a year after the opening ceremonies for NetNorth, Guelph and IBM organized an event to showcase what had been accomplished. The press came to hear about how professors could now enter into interactive collaboration with partners in New South Wales, or stay on top of projects when they were on sabbatical in Israel. "I don't know if they had a lot of coverage," recalled Fockler, "but I was really struck by the content of the demonstrations. I suppose the examples may seem rather trivial now but at the time it was very exciting."

By 1987 NetNorth had moved from its start-up days into a period of steady growth and regular maintenance. That year the consortium was planning its annual board meeting when IBM saw another opportunity, as Percival describes, “to create a forum in which IBM’s name would be spoken.” Rather than have a closed meeting, why not open up the meeting and make it a more public event? The first conference was to be held at Guelph and, while the university donated the use of its facilities, IBM underwrote the $15,000 cost of the event (and did so for several years after). The initial idea was for each of NetNorth’s regional groups to present an update of recent activities and their current status. Even at this first conference, however, the outcome was a far broader exchange of information.

A major topic at this Net 87 conference in June was the possible merger of CDNnet and NetNorth. NetNorth was built to provide reliable network services to member institutions and was considered a “production network”: downtime was minimized, upgrades were planned, events were scheduled to keep service disruptions to a minimum. CDNnet, on the other hand, was an example of a “test network” built in order to determine the viability of the X.400 email standards in the form of the EAN application software. Once the EAN software proved to be a success, however, the network became a critical connectivity resource to computer science labs having no other means to link to each other and the rest of the world.

Macneil remembers that he and the other NetNorth members had a sense that if networking was to progress in Canada, it would have to draw on the community as a whole. “I wanted CDNnet to come because in Canada I firmly believe that we are not big enough to have two or three national networks” said Macneil. “Here was a little competition, some fresh ideas, a new group to collaborate with.” John Demco from UBC and Bob Cavanaugh from Queen’s discussed the possibilities of cooperation between Canada’s two national networks.

But if hints at future convergences were part of the conference spirit, so were future divergences. While Ira Fuchs gave a presentation on the ARPAnet’s networking technology, a representative from IBM gave the attendees a preview of its own proprietary network protocols. Talk drifted to the growing stress on NetNorth’s volunteer management structure and the possible need for commercial involvement, perhaps even commercial use of the network – a fundamental issue that revealed the differences between NetNorth and other regional and international IP networks.

WHAT KIND OF NETWORK?

These issues arose again once the NRCnet proposal became unofficial public knowledge in the summer and fall following Net 87. This proposal for a national IP network was hugely attractive to NetNorth, since provincial initiatives such as CRIM, BCnet and ONet were establishing themselves – and establishing strong links to CSnet and NSFnet south of the border.

The NetNorth community viewed NRCnet as an advanced production network and agreed with NRC on the starting point: the network would take root in the academic community because they were the only group with the right combination of need, experience and resources to tackle the project. Linking the diverse members of the public sector research community afforded an excellent means for promoting collaboration and the cross-pollination of ideas – the lifeblood of academic research.

But an advanced production network that would also be open to the private sector with the hope of increasing revenues was another matter entirely. This was particularly an issue for NetNorth, which did not share the NRC's combined public and private sector mandate. While networking colleagues and resources located at two different universities had obvious advantages, linking two rival private sector companies did not. In fact, the idea of a permanent data connection was more likely in those days to raise serious concerns with regards to corporate security and industrial espionage.

From a funding perspective, moreover, little or no economic incentive existed for private enterprise to underwrite a network that would serve the public interest – to say nothing of the highly speculative nature of such a project given the uncertain long-term value of large-scale, packet-switched networking.

These concerns became more acute towards the end of 1987 as NetNorth neared the end of its three-year funding cycle with IBM Canada. Member dues subsidized the cost of connecting each organization to the network regionally and to that extent the network could be considered self-sustaining. However, IBM's money was paying for the long-haul line between the three-region network and into the US: the reality was that, given the low volume of traffic on the network, these lines could not be maintained without further subsidy.

GUELPH PULLS BACK

With these concerns in mind, Kent Percival at network HQ in Guelph began to investigate alternative sources of funding. He arranged meetings with representatives from potential organizations such as Bell Canada, the Communications Research Centre and the National Library. The results were mixed. Bell was unenthusiastic. The federal organizations were interested in data networking but only in projects that used OSI standards – which were now being promoted by the ISO and CCITT at the international level. Percival remembers being somewhat irked at the CRC meeting:

> They started the meeting by telling me that NJE was the wrong way to do networking. I reacted to that a little bit. I said "Look, we've built a network that's running across Canada now. Tell me about your network."

Percival next found himself sitting down in January 1988 with Roger Taylor from NRC and Guelph's vice-president, Jack McDonald. Taylor was there to sell the notion of NRCnet. NetNorth was the most likely vehicle for this kind of broad-based network initiative and Taylor was testing the waters for a network manager. Although Guelph was known primarily as an agricultural university before taking the lead with NetNorth, it had proven itself as more than capable in managing the network's operations centre.

Percival was eager to accept the challenge. He knew the NetNorth board was all but officially committed to migrating the network to the IP protocol. Guelph vice-president McDonald was less enthusiastic. His university was undergoing a round of belt-tightening and McDonald had already approved the budget for rebuilding its campus network into a high-speed LAN. Percival made the case that the network had been and would continue to be an excellent way of raising Guelph's profile. As the interface for NetNorth and the rest of the world, Guelph's name was well known internationally. But McDonald was unmoved. He asked if Percival could demonstrate how NetNorth had delivered any concrete return on its investment. Percival saw he wasn't getting anywhere and sadly, Guelph was going to focus on its internal networking needs.

As Guelph faltered, U of T had shown itself willing and able to replace Guelph as NetNorth network manager. But while U of T was finalizing the details of its new role in NetNorth, and its founding and participating role in the new provincial ONet, it was also establishing a permanent NSFnet link to Rochester.

These developments convinced NetNorth's managers that they had to act quickly to help establish a national IP network. On the one hand, NetNorth was still functioning well and the network itself was popular and continued to grow. But the writing was on the wall: with BCnet in place and the largest institutions in Ontario and Quebec committed to developing newer IP-based networks, it was clear the major players wouldn't be interested in or capable of maintaining two separate networking activities indefinitely. And without the larger institutions carrying their technical and financial load, the dozens of smaller institutions wouldn't be able to keep NetNorth going on their own. The question for all of the members thus became: how was NetNorth going to reorganize itself in light of IP's looming presence on the horizon?

As the NRCnet initiative gained momentum, a new and powerful player appeared on the scene, highlighting the issues NetNorth was grappling with.

TREASURY BOARD POLITICS

In the spring of 1988 Taylor spoke at a meeting sponsored by the management board of a networking initiative called SUPERNET. "Super" actually referred to the computers on the network, not the network itself. The Atmospheric Environment Service at Laval University was in the process of upgrading from its current Cray 1-S supercomputer (used primarily for weather forecasting) to a more powerful model. But instead of disposing of the old computer, Laval and interested parties from the federal government were looking into the possibility of distributing the computer's number-crunching capacity by providing networked access to the Cray. Taylor was attending the conference to discuss NRCnet and its potential use as a vehicle for the Laval project.

ISTC AND DIGGERNET

Peter MacKinnon was also attending the SUPERNET conference on behalf of Industry, Science and Technology Canada (ISTC), the federal department known today as Industry Canada. MacKinnon was intrigued by the NRC's proposal, and especially by its potential as an advanced production network that would also serve private industry, not just the R&D community. He saw a role for ISTC in bringing this technology to Canada's private sector. MacKinnon had roots in that community, and had in fact been seconded to ISTC on assignment from Cognos, a private high-tech company.

Mackinnon's enthusiasm for the network proved infectious around ISTC. Part of the department's mandate was to ensure that research and development initiatives such as NRCnet didn't ignore the development half of the R&D equation. And the department had certainly become attuned to the prospect of data networking as an important vehicle for technological and economic development. The government's increasing interest in this potential was good news for the future of data networking. Whether or not it was good news for NRCnet remained to be seen.

Questions arose, especially within the ISTC, as to whether or not the NRC, and by extension the academic community, was indeed the best choice to build Canada's national data network. Adding to all this doubt was a nascent project taking shape within ISTC itself headed by NRC physicist Digby Williams, who was himself on secondment to the ISTC. Word of his initiative began to travel through the academic grapevine and was quickly dubbed "Diggernet."

Williams worked hard to create support for his initiative within ISTC and larger industry circles, including computer and telephone companies across Canada. From the networking community's perspective, it suddenly seemed Diggernet was turning up everywhere and purporting to best NRCnet in every way: it would be larger, faster and cost the government less money. It would be OSI-compliant. It would involve more industrial cooperation.

And Williams was blunt about what he saw as the shortcomings of the NRC-backed initiative: the economic potential and importance of data networking simply couldn't be left in the hands of a group pursuing what Williams considered to be a lightweight network intended primarily for the benefit of academic research.

ISTC's views on NRCnet and its own high-speed initiative started to play an important role on the next stage in the project that counted most – funding approval from the Treasury Board.

HALF-STEAM AHEAD

Although the NRC had committed a portion of its own budget for NRCnet, an expenditure of that size and an initiative of that scale would be carefully scrutinized by Treasury Board, which gives (or withholds) final approval for all government expenditures. As project manager of NRCnet, Andy Woodsworth took charge of the effort to develop a plan that would meet with the Board's approval.

The NRC initiative had two persuasive arguments in its corner. The first was simply that the NRCnet was ready to be built. The technology was proven and the provincial networks were being put in place. By contrast, the ISTC-led initiative to build a high-speed OSI network was still in its formative stages.

The second argument was founded on equal parts of truth and rhetoric. Canada had fallen far behind the US, and a number of other western countries, with regard to its data networking infrastructure. The NSFnet had been operational for almost three years and had already gone through a series of bandwidth upgrades. Meanwhile, a comparable Canadian IP effort had yet to get off the ground. Without an east-west backbone in place, regional network connections would continue to connect by taking the shorter route south to NSFnet.

Using the NSFnet as a long-term solution, however, was no foregone conclusion. Any attempts to use the NSFnet as Canada's de facto network backbone would certainly not be welcomed by the NSF, if only for reasons of cost. Moreover, Canada's telecommunications regulator, the CRTC, had its own reasons to prevent any such development: it had put regulatory barriers in place to prevent transmissions beginning and terminating inside Canada from being routed through American facilities.

But even more importantly, if Canada missed this opportunity to build the next generation of computer networks, Canadian networking would return to the regionalism that characterized it prior to the NetNorth and CDNnet initiatives. Canada's ability to take the next step forward in networking technology would be hobbled by an inability to collaborate and share resources between regions. This threat to Canada's technological future was certainly not lost on the networking community, which played the nationalist card to win support from Ottawa whenever they had the chance. While they knew this ploy was viewed by some as rhetoric or mere flag-waving, the network proponents at both NRC and NetNorth were quite sincere about their nationalist sentiments. As Gerry Miller recalls:

> We simply felt it was important to establish the Internet in Canada on a Canadian network. I don't think Canadian traffic should have to go through the US. Sure there are network efficiencies to keeping traffic in Canada, but that's not really a big deal. It's a nationalist argument. It was important to maintain a Canadian identity because if you have regional networks stitched to the US, then effectively we just become an extension of the US.

Meanwhile, the project itself continued to move forward as it awaited final approval. A separate implementation committee, chaired by Leigh and composed of government representatives and members of the existing networking community in almost equal proportion, was formed to oversee the next steps in actually building the network. Included were the familiar faces of David Macneil (UNB) and Paul Dirksen (Waterloo), Bob Cavanaugh from Queen's, Bernie Turcotte representing CRIM, and Peter Jones from Dalhousie. From the government side were Woodsworth and Curley (NRC), John Gilbert from the Government Telecommunications Agency, Chris Hughes from Energy, Mines and Resources, and Ron Watkins from ISTC.

The committee's task was to oversee the development of an official Request for Information (RFI) and use the resulting submissions as the basis for a Request for Proposals (RFP). This approach is a well-established method for executing large, complex projects using the best possible resources and the most capable contractor, within budgetary constraints. But the committee had to tackle more fundamental issues than merely selecting contractors – such as deciding if the project would continue to exist at all in light of dwindling support from government and the private sector as the Treasury Board began examining the proposal.

Although NRC had supported the project from the start and worked hard to consolidate support on other fronts, consensus was not forthcoming from within the federal bureaucracy as a whole. For one thing, it was public knowledge that Canada officially supported the OSI standards being developed for network interconnection by organizations such as the telecom-backed CCITT and the ISO. A number of federal departments including the Treasury Board were concerned that the NRC initiative to fund construction of a national IP backbone would run afoul of the government's support for the OSI protocol suite – the basis of the rival option promoted by Digby Williams and ISTC.

Bruised egos aside, the real problem the NRC had with the Diggernet proposal was that they felt Williams was selling his project as an alternative to the NRCnet. The way Williams and ISTC saw it, the OSI high-speed network would eventually become a full-scale production network. But such a network could not be built overnight; in fact, it would take many years to mature from a preliminary test network into a full production network. In the meantime, if NRCnet was not funded, NetNorth would be left on its own to establish a national IP-based network. If it did not succeed, the country's only national production network would devolve into individual regional networks dependent for Internet access on their own connections to NSFnet. To the NRC, therefore, to fund the ISTC project and not fund NRCnet was to raise the possibility, for the immediate future, of having only a limited access network, in this case CDNnet, as the country's national data network.

As a result of the competition between the two projects, potential supporters of both networks – the telephone companies, computer companies, provincial networks – were all getting mixed messages about the government's intentions and a technology that was poorly understood outside of specialized circles. The unhappy result was that most of the potential supporters either backed off completely or wouldn't support one network at the expense of the other. The process of getting Treasury Board approval was delayed by months.

DOING IT THE OSI WAY

The Canadian government's reluctance to give its blessing to an IP-based network was neither surprising nor irrational. Like the major players in telecommunications, it was committed to centralized control and to implementing change from the centre to the periphery. It was equally committed to the OSI standards and the process they represented: top-down planning and implementation initiated by the CCITT and ISO on behalf of corporate telecommunications interests.

These organizations both carried the weight of historical precedent that reached back to the last century. Both were large bureaucratic agencies where decisions filtered through numerous technical committees and sub-committees whose recommendations were coordinated and approved at numerous levels by multiple boards.

Each wave of technological change in the field of telecommunications, including the introduction of digital transmission, has been planned, developed and implemented through the work of highly regulated national telephone companies, typically operating as monopolies, and the international standards-setting organizations. In telecommunications, standards are everything. Telephone networks can only communicate between jurisdictions if their equipment can inter-operate – which means the equipment must be standardized, traditionally by organizations such as the CCITT; for them, technological change was a planned and orderly affair.

For some time, packet switching went unnoticed by the telecommunications standards organizations simply because it didn't appear to be relevant to their main concern – voice communications. Only later, as real-time data communications did indeed become relevant to large industries such as airlines and banks did the CCITT begin to develop its own packet-switching standards, beginning with the X.25 protocols.

In 1983, after years of working in parallel and in collaboration with the CCITT, the ISO agreed to join with the CCITT to work together on OSI using the seven-layer model established by the ISO. While the end goal of the ISO-CCITT alliance was similar to that of the academic and military research organizations adopting the TCP/IP protocols, the two camps couldn't have been further apart in terms of culture and procedure.

No technology is neutral in its effect, and the telecommunications industry initially viewed packet-switching as disruptive for two principal reasons. First, packet switching encourages the migration of intelligence from the centre of a network to its end nodes – a notion contrary to the electronic circuit-switching approach that had evolved out of manual switchboards and the more advanced automated, electro-mechanical switches. IP was a technological innovation that began at the edge of the network, where the nodes served as the switching mechanism and the existing switches were bypassed altogether.

Second, this new technology was developed not by the telecommunications industry but by the computer science community. The TCP/IP protocol suite created by Cerf and Kahn was more than just a brilliant engineering solution to many of the problems inherent in computer communications. It was also a threat to the established order – to national telephone monopolies and the international standards organizations that kept them working together. Furthermore, the TCP/IP community was characterized by a pragmatic approach using what worked rather than what was supposed to work – a commitment to operating by "rough consensus and running code."

IP was a far cry from OSI – or at least from what OSI promised. Nobody could guarantee service quality with IP. It was hard to adapt to existing pricing models. It hadn't even been voted on! Not surprisingly, the still loosely defined TCP/IP community and the official OSI groups mixed like oil and water. The divide grew rapidly as the decade wore on: by the time NSFnet implemented IP in 1985, almost every major western country had signed on instead to support the OSI standards.

Meanwhile, OSI continued to be "in development." Bits and pieces of its various layers were employed as prototypes all over the world: CDNnet's X.400 project was a prime example. But the truth was that OSI was still an experiment while IP was developed by and for a community that had used it at production levels for over six years. IP worked. So far, the same could not be said of an OSI-based network.

NETNORTH COMMITS, U OF T SUBMITS

During the spring of 1988, as Treasury Board pondered the stalled NRCnet project, NetNorth decided to go public with its intentions to back a national IP network. The NetNorth directors assembled a planning group during their summer meeting at the Net 88 conference in Fredericton. Gerry Miller was to chair the group and act as the executive committee's representative while Roger Watt would represent the administrative half of NetNorth.

In October 1988, Miller, Watt and NetNorth's directors and representatives from across the country gathered in Winnipeg. A strategic plan had been laid out by Miller and Bjerring, a proposal that essentially committed the members of the NetNorth consortium to migrate from a trans-Canada NJE network to a national-backbone IP network. The group overwhelmingly approved the plan, which not only committed the network to migrating to an IP, packet-switched technology, but to reorganizing the network based on individual provincial boundaries rather than using its existing three-region structure. The feeling was that if NRCnet failed, they would build an IP-based network themselves.

In the same month, the NRC implementation committee finally issued its RFI. The request described the broad goals of the network, pointed to the NSFnet as a model and outlined the funds available. One of the questions it posed was how the network could incorporate, and eventually migrate to, OSI standards where possible.

The staff at U of T Computer Services (UTCS) received the RFI with some satisfaction, seeing much of their previous proposal to NRC reflected in it. The systems team outlined a UTCS-led consortium that included IBM Canada, with CN/CP to provide most of the long-distance facilities currently used by NetNorth. It also outlined a proposed topology and made explicit recommendations with regard to equipment, network management, costs, staffing and even potential research opportunities that could be hosted by the network.

IBM's role was to supply the network's Monitoring and Management Processors (MMPs) as part of a joint study program. The management of a real-time IP network was a very different job than UTCS was currently handling on behalf of NetNorth. An IP network can be (and generally is) centrally managed. The network operations centre (NOC) is able to monitor the status of the entire network at all times and can detect malfunctions, locate, diagnose and often fix the problem, all from the NOC. The MMPs would report information back to the NOC and allow staff to manipulate the routers and network traffic remotely.

For its part, CN/CP offered to supply bandwidth to the network. Unfortunately, being a regulated carrier, the company was not allowed to discount or donate the lines to UTCS. Instead, CN/CP promised to make a cash contribution to offset the costs of the facilities to be used by the NRCnet.

The highly detailed response was very well received by the committee, which had received nine responses from interested parties in total. UTCS had met all the required criteria as set out in the RFI, except for an obvious resistance to incorporating OSI standards into the proposed network or describing in detail how or when a migration to those protocols would occur. As UTCS stated in the RFI, the problem was that OSI covered a comprehensive spectrum of network functions but very little commercially available equipment conformed to OSI specifications:

> Committing the backbone network at this time to an OSI-only approach is likely to lead to a network with little or no traffic. Commitment to a pragmatic evolution of the national backbone network into the OSI realm is more realistic, practical and desirable.

Although buoyed by the RFI responses and the solid support of the NetNorth consortium, the implementation committee was left somewhat directionless while waiting for the Treasury Board approval that would allow them to issue an RFP and subsequently award the contract to build the network. The government participants on the committee were a house divided, Leigh recalls:

> It was interesting to hear them discuss things, because they had their differences of opinion. The government wanted to be seen as having one voice in this, but it was pretty clear they had their differences, and in thinking back I assume that Digger was part of the differences behind them.

So the debate in the committee raged on. Although everyone remained dedicated to the project as months passed, the group felt it had to renew its commitment to going forward. “Two or three times during the process,” said Leigh, “which went on for over a year, we took a vote in our little committee as to whether we should continue, which we did. And I would inform the NRC and they would be reassured a while longer.”

In the middle of all this, Williams and the ISTC undertook a feasibility study in March 1989 as to whether the government should help build “an intelligent, multimedia network for collaborative research in strategic technologies.” The network was described as using OSI standards operating at T1 capacity (about 1500 Kbps, almost 80 times greater than ONet’s original bandwidth) “as a means of promoting a more internationally competitive information technology sector.” The aim of the study was to help decide whether the government should step into the open-systems arena at the test network stage in order to steer industry onto the right path. In effect, the study was trying to determine if there was sufficient economic payback for government facilitation of this process.

Woodsworth, in the meantime, had drafted and redrafted the Treasury Board proposal but to no avail. The government’s fiscal year ended in March 1989 without the Board’s approval for the project budget.

5

THE BIRTH OF CA*NET

WE COULD NOT EXIST AS A COUNTRY WITHOUT THE WHOLE NETWORKING PRINCIPLE. WE MUST BE PARTNERS, OR WE WILL BE NOTHING.
– DR. WILLIAM WINEGARD

As the various players waited for Treasury Board's decision, NetNorth held its annual networking conference in May 1989 at Concordia University in Montreal. The NetNorth conferences had grown larger and more diverse every year and were becoming the central forum for the discussion of next-generation networking activities in the country. Net 89 drew a record-breaking crowd that included its first bona fide political celebrity – Dr. William Winegard, Minister of Industry, Science and Technology.

THE MINISTER AT NET 89

Winegard opened the conference with a few words about the importance of networking to Canada. Ottawa's goal was to connect Canada's "three solitudes": government, business and academia. His speech almost echoed the NetNorth group's own words:

> This country has been built by networks of one kind or another, from the railroads and the Seaway to our modern telecommunications systems. We could not exist as a country without the whole networking principle...
> We must be partners, or we will be nothing.

Winegard wasn't there to sing the praises of IP over OSI, yet the mere fact that he was speaking on NetNorth's turf was a positive sign. It was also a good opportunity for the academic networking community to bend the minister's ear about their IP network. Digby Williams, who was overseeing ISTC's feasibility study of government involvement in the test phase of a high-speed, OSI-compliant network, had the same idea.

The rival groups met with Winegard in a Concordia boardroom. Andy Bjerring, Ontario rep on the NetNorth National Executive, decided that somebody had to do the talking – and better him than Williams:

> I must have done a lot of the talking. I told Winegard, "I think there's history in the making here. And we should be recording this because it's at this point that you're being given the opportunity to change the direction of history. We're ready to take this next step. We're building the infrastructure for tomorrow." And all that kind of stuff. And he asked, "Are the universities behind all of this?"... "The universities are eager, we're ready to go."

But the truth was the academic community did not quite speak with one voice on the subject of the NRnet (the "National Research Network" or NRN), as it was now called. The conflicting views of the NRC and ISTC over IP and OSI were reflected in the differences of opinion at Canada's universities over NetNorth and CDNnet. The possible merger of the two existing networks discussed at Net 87 hadn't happened. As set out in the annual executive reports from NetNorth and CDNnet presented at the conference, the burning question was still – to IP or not to IP?

NetNorth, with its background in production networks and ad hoc organization, was firmly in favour of migrating its existing network to IP – and effectively merging the NetNorth initiative with NRnet and the provincial IP networks. This merger had actually taken place in Ontario with Andy Bjerring as the chair of the merged entity. To say they were in favour of a national IP backbone is something of an understatement. From their perspective, with UTCS the front-runner on the basis of their response to the NRC's RFI, NRnet would become in effect a NetNorth initiative that had the NRC's backing.

NetNorth's executive committee was in fact in the midst of working out the final details of their plan to migrate NetNorth to the IP protocol. Their 1989 strategic plan called for the encapsulation of NJE traffic within the TCP/IP protocol suite. This joint solution meant a member institution could initially support both NetNorth and the proposed IP network over a single leased line. Separate facilities for NetNorth were to be phased out as institutions made the switch to IP.

CDNnet, on the other hand, whose raison d'Ítre was the development of the OSI email standards, remained committed to its own ideals. As Demco commented at the conference:

> The Canadian government and other governments are headed toward OSI. Our OSI work is important and should be supported. After all, we're not operating in a vacuum. It's an international community out there and we in Canada have to be part of it.

The CDNnet position was not inherently anti-IP. For one thing, CDNnet was based on an application designed to work with multiple network protocols, including the TCP/IP suite. Nevertheless, as the OSI proposal began to be seen as a competitor for government favour and funding, tension developed between the NetNorth and CDNnet communities.

CLAIMING THEIR NETWORK

In June 1989, rather than treating the NRC and ISTC initiatives as incompatible, the Treasury Board accepted the NRC's proposal – on condition that the NRC would only use TCP/IP until the OSI suite was ready. This solution cleared the way for the NRnet to proceed. The ISTC-led high-speed test network was now viewed as a project that would build on NRnet and eventually transform the NRC's "woolly-headed" academic project, as Leigh quipped. With the money guaranteed, the NRnet implementation committee was able to issue its official Request for Proposal.

The RFP was structured along the same lines as the RFI but went into greater detail, laying out technical guidelines for bandwidth, latency, redundancy, security and protocols. It asked for concrete plans to incorporate and implement OSI-based technology as it became available. The document noted the ISTC initiative, pointing out that "there is no commitment that the operators of the NRN phase of this joint initiative would receive preferential treatment in the selection of the eventual operators of the proposed ISTC network."

As the committee waited for responses, the NRC decided to make a break with generic-sounding names and anything with the letters N and R in it. Instead they wanted something that wore its patriotic heart a little more on its sleeve. Taylor suggested Maple Net, which got a laugh but no approval. Curley finally settled the matter, embellishing his suggestion with some nifty design work. He unveiled the new name cum logo at one of the network committee meetings: CA*net. The asterisk was the keyboard equivalent of the maple leaf that would appear whenever the name was graphically displayed as a logo.

The nine responses to the RFP were submitted by mid-August. U of T's submission was similar to its initial reply to the RFI, with a few alterations. Instead of CN/CP, U of T was working with a new player on the Canadian telecommunications scene: a bandwidth reseller called Integrated Networks Services Inc (INSINC). Resellers were a recently created category of telecommunications provider with regulatory approval to buy bandwidth from facilities-based licensees such as Bell Canada and CN/CP and resell the bandwidth in smaller parcels at discounted retail prices. The cost savings allowed U of T to keep their proposal within NRC's budget and still create an ambitious network topology.

The other significant alteration was the switch from advanced Cisco routers to IBM RTs, the routers used by the NSFnet. Siciunas thought of it as a "hand-me-down" way to equip the network – not exactly leading-edge, but at least proven technology. This choice meant that the NRnet and NSFnet would have no compatibility problems. IBM was also donating the Monitoring and Management Processors (MMPs) that would monitor the network. The MMP was in fact a second RT that monitored the first and sent the data back to the NOC. One advantage of this router configuration was that, if the RT router broke, the MMP could be used as a backup.

In October, less than a month after the RFP was issued, the implementation committee sent notice that it favoured U of T's proposal. Essentially they had won the contract. But there were a number of issues to be tackled before the second- and third-place proposals would be dismissed entirely.

An exchange of letters between UTCS and Earl Dudgeon, the NRC VP overseeing the NRnet project, reveals the last-minute negotiations to meet Treasury Board priorities: the topology was to be modified to provide PEI with as high a bandwidth connection as the other provinces; more private sector participation should be included; and the network management board (NMB) should make room for more government observers.

UTCS accommodated these requests. It reduced other network priorities such as redundancy and monitoring equipment to provide for a 56 Kbps trunk line to all the provinces, including PEI. It also adjusted the NMB on condition that, "as in our proposal, the NMB becomes the 'owner' of the network and bears full responsibility for all financing of the network."

The final request from Dudgeon was for the U of T group to specify just how they would incorporate OSI standards into the network. The question at first received an informal reply from Siciunas. In return, Dudgeon asked for a report detailing how the implementation of OSI would be handled. U of T submitted a 2,000-word report outlining the various aspects of the network – routing, addressing, monitoring, management and top-level applications such as email – and discussed where and how the OSI-compliant technology could be introduced. But underneath the detail, the message was clear: we won't use OSI if it disturbs the operation of the network and, besides, no element of the network currently qualifies for immediate OSI implementation.

The NRC responded with still more questions about the details of the implementation. But Siciunas, who was only a week away from the first meeting of the NMB, cut the conversation short in his last letter to Dudgeon:

> I sincerely regret that we seem to be incapable of providing you with the "...words you want to hear..." about OSI migration. It is not that we are being obstinate or devious or secretive. I think the problem stems from a lack of mutual understanding as to what "migration to OSI protocols" really means. We are responding to what we think your question asks, but you are obviously expecting a different answer.
>
> Our primary emphasis will be on establishing a functioning, stable, production network. Strictly speaking, we could install the network nodes complete with OSI code and watch what happens. However, I would consider this to be an unprofessional and cavalier way to run a "production" network. Instead, the OSI code should be tested in an off-line test environment of several routers and, once satisfied with the functionality and stability, the code could be propagated to the production network.
>
> I expect my comments have raised more questions than provided answers, but hopefully the meeting with Dr. Woodsworth will bring about mutual understanding of what we all mean and desire.

Days later, the NMB met to decide how they would run the network they had now claimed as theirs.

HOTEL MEECH

UTCS wanted the first meeting of the NMB, hosted by U of T, to work out fundamental issues such as the network's costing structure, topology, governance, etc. The meeting was set for mid-November. Once again a call was sent out to Canada's networking community to come together to plan their network's next incarnation. The core of the NMB would be drawn from NetNorth's existing regional committees: at least one representative from each of the 10 existing or planned provincial networks attended the meeting in Toronto. For a group that intended to run 10 separate networks as one seamless whole, they had their work cut out for them.

The two-day meeting, chaired by Warren Jackson, came to order in November 1989 at the Skyline Hotel in Toronto. By a strange and somewhat ominous coincidence, the NMB's meeting was taking place in the shadow of the Meech Lake constitutional negotiations. Just days before, Canada's premiers had met with federal ministers behind closed doors in Ottawa in an attempt to reform Canada's constitution. Now the NMB gathered in an airport hotel to tackle some nation-building of their own.

The networks represented at the meeting ranged from fully operational, inter-city networks with links to the NSFnet such as BCnet, ONet and RISQ, to situations in Saskatchewan and Manitoba where university LANs had yet to make an interconnection. The variation wasn't surprising given the differences that generally characterize Canada's regions. Densely populated southern Ontario, Quebec and BC, centres of industrial and political activity and close to the US border, could be expected to have a head start in their networking initiatives. Smaller centres did not have the same opportunities.

Between these extremes were provinces with significant network infrastructures in place but incompatible technology. The NB/PEI ECN, for example, had continued to grow and diversify the number of services its low-speed 19.2 Kbps of bandwidth could deliver: the multiplexed lines now carried remote-terminal access, NetNorth and email service.

At the Skyline meeting most of the provincial representatives were from university computer service departments. They included Jack Leigh from UBC, Glen Peardon from Regina, Dean Jones from Saskatchewan, Gerry Miller from Manitoba, Andy Bjerring from Western, David Macneil from UNB, Peter Jones from Dalhousie, Jim Hancock from UPEI and Wilf Bussey from Memorial. There on behalf of the existing provincial networks were Mike Patterson from BCnet, Walter Neilson from the Alberta Research Network, Bernie Turcotte from RISQ and Mike Martineau, who was now leading the Nova Scotia Technology Network (NSTN) initiative.

Present as observers were Alan Greenberg on behalf of NetNorth, John Demco on behalf of CDNnet, and John Curley and Roger Taylor for the NRC. Jackson was officially there as the observer for UTCS. Although these members were active participants they couldn't vote on final decisions.

The NMB's two most important tasks in this brief encounter were determining the formula for calculating how much of the network's cost each province would have to bear and finalizing the network's topology. The work on the costing formula had two components – the first was straightforward, the second a little more contentious.

The NRC seed money was intended to ease the immediate financial pressure on the provincial networks during the start-up phase. Thus the first job was to establish how much of the network's cost would be paid by the NRC in the first year, and how fast the provincial networks would increase their dues. By the end of the five-year funding cycle, the network members would be carrying the total cost of the backbone themselves.

The second component to the costing formula presented more of a challenge: determining what each of the provinces owed as their "fair share" of the backbone's expense. Dividing the costs evenly among the provincial networks would not only ignore the existing disparities between the provinces, but also the prospect that some provinces would generate a great deal more traffic than others. On the other hand, the group didn't want to adopt a metered system obliging the provincials to pay for their traffic by the bit, since they felt that would discourage use of the network. One of the popular characteristics of NetNorth had been its flat-rate pricing scheme. Computer service departments could estimate in advance what their network costs would be for the year ahead, making it much easier to budget and control expenses.

The dilemma was how to formulate a cost of access that was sensitive to usage but that didn't charge by the bit. The group needed to estimate the proportion of traffic that each provincial network would probably generate over the coming year – without any records of network traffic to use as a guide. So in lieu of actual usage data it was proposed that a province's population be used as a proxy. A province such as PEI would benefit a great deal from this approach but Ontario and Quebec suddenly faced shouldering a much larger share of the network's costs. This prospect was especially troubling to Turcotte and Greenberg because, even though Quebec is Canada's second-most populous province, the new provincial network (RISQ) was just getting off the ground. To use Quebec's population as an index for traffic was therefore questionable.

There was considerable argument and infighting on the particulars of the costing model but, underneath it all, everyone agreed on the big picture. Still, on that day in Toronto, the particulars dominated and the costing issue had to be resolved another day. Looking back on it now, most of the networking group agree that their eventual success in resolving this problem and many others stemmed from the fact that they got along and trusted each other well enough to speak their minds plainly.

The board also tackled the issue of the network's topology, which led to another hot and heavy discussion: the two major factors at issue were bandwidth and redundancy. Bandwidth, the line's capacity for sending data, depended on the circuits leased from INSINC. More bandwidth meant a faster network. Redundancy, on the other hand, refers to extra pathways joining any particular node to the rest of the network. Investment in redundancy meant a more dependable connection because if one line went down another could take its place.

The NRC's RFP had originally called for every provincial network to be connected to the backbone at 56 Kbps and for each one to be redundantly linked to a second node on the backbone. When UTCS mapped out and priced a network to meet those requirements, it far exceeded the project's budget. Instead, U of T recommended a slower network that maintained its redundant links. But this network architecture left Manitoba and Saskatchewan, and the entire east coast, connected at 19.2 Kbps while Toronto, Montreal and Vancouver would enjoy a 56 Kbps connection. Fewer users, less demand, was the reasoning. In response, the NRC had insisted there would be no "poor cousins" on the network and asked UTCS to adjust its design, drop its redundant links and ensure that every node had at least one 56 Kbps connection to the network.

At the Toronto meeting, however, the NMB discovered that a network's topology is a bit like a car seat: once you start fiddling with it you can't seem to get it right again. The members worked on finding the design that provided the greatest amount of redundancy with the fewest number of lines. But instead of making progress it felt as if the proposals were getting increasingly convoluted. The first day of meetings would leave the second issue seemingly less resolved than before they had arrived. The group retired for the day to take refuge in the comforts the hotel had to offer and make a fresh start in the morning.

They quickly discovered that the hotel didn't offer much in the way of comfort so they set out for the nearby Holiday Inn to get some dinner. It had been snowing all day and it was a cold night. In the middle of what had turned into an unexpectedly long, cold walk the group narrowly avoided another argument ("Who chose this restaurant anyway?") by discussing the network. Perhaps they were bonding in the face of adversity but, for whatever reason, they began to make progress: by the time they returned from the restaurant, they had chosen the board members of what was now called CA*net. The next day turned out to be a memorable one. After meeting in the morning, the group had lunch together and shared a collective epiphany regarding the network's topology. The vision manifested itself on Taylor's cocktail napkin as a figure-eight:

> I realized as we probed we were all talking about something that was topologically equivalent. I pulled out a napkin and drew two interlocking rings on it and said, they interlock at Montreal and Toronto and everything to the right is Atlantic Canada and everything to the left is western Canada and logically you can put cities anywhere you want on these rings, so it should be pretty easy to optimize.

A simple way to provide every node with redundant links using the minimum number of lines would have been to make the backbone a single narrow ellipse. Every node would then have two connections to the network. Unfortunately, it would only take two downed lines for the entire network to be split in half.

By twisting the circle to make two ellipses, and then redundantly linking the NOC in Toronto to both halves of the network, the individual nodes would be redundantly connected and large sections of the network couldn't be easily isolated.

The breakthrough seemed to restore the group's momentum. After lunch the meeting resumed and the members continued to work on the costing formula. What they finally arrived at is what Jackson called a "classic Canadian compromise": half of the cost to join the network would be based on an equal division of the expenses between the ten provinces, while the other half would be based on usage. However, until they had reliable usage statistics, they would pro-rate the other 50% according to a province's population. Looking back on his role as chair, Jackson said with a smile that this was "one of the only times I can say I actually helped with nation-building." Siciunas remarked that "it really was a classic Canadian compromise, which is strange because Warren is an American."

Before the group went their separate ways, they also worked out a number of issues regarding ownership and governance of the network. The NMB would essentially become the network's board of directors and it would be incorporated as a not-for-profit organization. Peter Jones was elected as the board's first chair. Andy Bjerring, its first treasurer and Gerry Miller, its first secretary, arranged to have CA*net Networking Inc. federally incorporated with its head office in Toronto on June 27, 1990.

GETTING UP AND RUNNING

The group proposed to build the network in three basic stages: preparing the provincial networks to join the backbone; establishing the backbone's connections to the NSFnet; and establishing the lines between the provincial nodes and getting the backbone operational.

By the time the Net 90 conference was held at the University of Victoria that summer, the project had finally cleared the paperwork and red tape and was under way. Having started the actual construction process in May, UTCS was aiming to finish by September. The network's first phase, completed June 1, saw U of T's NOC connected to nodes in Fredericton, Montreal and Vancouver. The next phase ending July 1 would connect Montreal to Halifax and Winnipeg, and Vancouver would link with Edmonton. The final phase would see the University of Saskatchewan in Saskatoon join through Winnipeg, St. John's through Halifax and Charlottetown through Fredericton. These connections wouldn't immediately create the interlocking rings the NMB had planned on, but it would be enough to get all the provinces joined and communicating over the backbone.

Meanwhile, NetNorth's planned phase-out was also under way. During the transition, UTCS was serving as the NOC for ONet, CA*net and NetNorth. Siciunas recalls:

> As people saw the logic of not paying for a parallel pipe when the same packets can be carried over a single, much faster line, the old NetNorth lines just fizzled away. The actual management of the old distribution table stayed around for a couple of years longer... The table management authority was taken over by some European outfit that decided to start charging for it, so everybody that wanted to stay in the routing tables had to pay. Our levy for Canada was $20,000 or $30,000 at the time. We said, "Oh boy, we'll pay this once, but encourage people to get off it. If you don't get off it, you pay yourself." Everyone eventually did get off it and whatever money was left – there were some NetNorth funds left over – we just divided back among the regions.

After NetNorth was phased out, UTCS still operated both CA*net and ONet out of the same building, a situation that had its own particular problems. Fortunately UTCS was working quite closely with the MERIT group in the US, who had experience with a similar situation since they were responsible for managing both their regional network and the NSFnet backbone. MERIT was also using similar IBM equipment, so their expertise was invaluable as they worked with UTCS to install and configure the IBM RTs being used as routers and monitoring equipment throughout the network.

Implementing the technology being used for NSFnet wasn't always a straightforward process. The systems crew, including Bob Chambers, Dennis Ferguson and Rayan Zachariassen (a graduate student at the time), often had their work cut out for them. For example, the RTs being used at the MERIT NOC were connected to a single T1 line, which provided roughly 1.5 Mbps (1,500 Kbps) of bandwidth. At U of T, however, the RT used on the southbound connection had five separate lines plugged into it: three delivered a combined 168 Kbps of bandwidth south, while the other two led east and west. Because the RTs had trouble keeping up with that many lines, Ferguson and Zachariassen had to make constant adjustments to the system to get it to work – what Siciunas describes as "shoe-horning code."

The task of setting up the lines between each point-of-presence (POP) where a provincial network accessed the CA*net backbone presented special challenges that called on Siciunas' unique powers of persuasion as much as his staff's technical abilities:

> In New Brunswick, at NB Tel, their VP didn't like resellers and he refused to put in the link from Fredericton to Charlottetown. The telephone company already had their own lines to the university for the existing educational network and he didn't like this notion of the reseller and this national network stealing business from the phone companies. I whipped off a really scathing memo that I got our

VP to send off to this guy because he knew him personally. It was something along the lines of "I'll be more than happy to inform Canada's research community why some small-minded telephone bureaucrat doesn't think a particular province should enjoy connectivity with the rest of Canada." Our VP just took the whole memo and forwarded it to the VP in New Brunswick and added, "As you can see, there are some strong feelings about your actions... "

I remember having discussions with our rep at INSINC. I used to give him such a hard time. I'd say, "Gee, you must be one of CN/CP's biggest customers. You guys should be able to get better service out of them. Beat up on them! Give 'em hell!" We had a grand time...

THE US LINKS

Compared with going east-west, setting up CA*net's north-south links to the NSFnet was a breeze. The first contact was made in December 1989 when NRC's John Curley met Steve Goldstein from NSFnet. The mandate of Goldstein's department was to encourage the development of IP networking, a separate function from the technical job of co-ordinating networks. He and Curley met at one of the earliest INET conferences, held that year in Australia. Instead of the thousands that typically attend INET conferences now, Goldstein recalls that the meeting down under involved an informal group of only 30:

John Curley and I went out for lunch one afternoon and we were talking about connecting our two networks. He had some degree of trepidation about just how robust a single 56 Kbps link across Canada might be. So he asked me, if for some reason CA*net were to break in the middle somewhere, could they heal CA*net by routing traffic through the United States since CA*net would be connected at more than one place to the U.S.? I said, "Sure, John, if you give us the same privilege of healing the NSFnet through CA*net." Which on the face of it was rather ridiculous, because of the vast disparities in bandwidth. Really it was said part in jest, but also in part to maintain parity in the agreement.

In principle Curley and Goldstein had established a peering agreement, i.e. neither network had to pay for access to the other. Instead, the two networks acknowledged the mutual benefits of sharing traffic and the connection costs between them. Roger Taylor was later delegated to conduct the actual negotiations, which were quite straightforward:

That was a nice exercise actually because I was given a limit that I could spend up to $100,000 a year or something like that on establishing the links. I went down to Washington and sat down with Steve Goldstein and in pretty short order we worked out a formula whereby basically he said, "I'll bring my line to the border and you bring your line to the border," and it turned out it cost us about $57,000 instead of $100,000.

By this time access to the NSFnet provided access to the Internet more or less as we know it today, beginning with a high level of use in international academic circles. As Goldstein points out, the US was slowly moving from essentially being the whole Internet to becoming the hub of the global Internet:

We had connections to NORDUNET, the backbone of the Nordic countries. We had a connection to the U.K. and, through some arrangements with our sister agencies in the United States, I believe sometime in the early eighties, we began to have connections with Germany and other places. And then in the early nineties we worked with a consortium called PACCOM and shared connections to Japan and Australia. Canada was able to share in all those international connections. In fact Canada's external traffic, external to North America, was pretty much routed through the NSFnet.

Using the NSFnet as CA*net's global hub was an important advantage for the CA*net group and Canadian networking in general. However, the government might have been taken aback if it had realized that technically this link infringed on Teleglobe's international telecommunications monopoly. It was yet another example of the Internet turning established communications regimes on their heads.

CUTTING THE RIBBON

In August 1990, Siciunas announced to the CA*net board of directors that the network was up and running with all the provincial nodes connected. For once things were running a little ahead of schedule. It wouldn't be until a few months later that the network would hold its official opening ceremony at U of T in October.

Jackson played master of ceremonies and William Winegard was "invited to apply his ministerial digit," as Jackson said in his opening remarks, to the button that would officially "cut the ribbon" and open the network for business. The network had been operational for some time but Zachariassen, who was running the show from behind the scenes, lent some of his technical ingenuity to create a little dramatic flair for the occasion.

Instead of turning the entire network off and on, which UTCS thought would be irresponsible, Zachariassen had the idea of simply turning off the monitoring equipment. When Winegard hit the button, the monitoring equipment was switched back on and simulated the powering up of the network. The assembled crowd was able to watch the provincial nodes come online while data flowed from one end of the country to the other.

After the ribbon cutting, Jackson carried out a few demonstrations of the network's reach and capabilities. The MERIT group was contacted and, as the audience watched on a projected screen, a message was typed out to one of the MERIT people attending the ceremony. It turned out to be an urgent matter that sent the MERIT guest running off to call her office. In another demonstration, a person contacted in Japan asked about his favorite literary character, Anne of Green Gables, a question which turned into a public relations windfall.

TRAINSPOTTING

After the network was technically up and running the next task was to get CA*net's NOC and network information centre fully functioning.

The primary job of the network operations centre was to monitor network activity, perform maintenance, implement upgrades and troubleshoot problems as they arose. One of the early bugs they encountered was that the network was redirecting traffic around damaged lines so quickly they couldn't detect a line was down at all. Siciunas:

> After we built the links to close the network, it screwed up our monitoring because it was the type where the SNMP suite would send a query to a router and give it a window within which you'd expect the response to come back. If it didn't come back in time, you'd say "Oh shit, something's wrong. That one's broken." The odd part was that when we closed the loops, if a link was broken the damn router sent the packet the other way and we'd still get a response back in time and think nothing was wrong. We had to shrink the window so we could force it to come back along a particular path rather than the other one. It took us a week or two to figure this out...

Siciunas recalls another unexpected form of "network monitoring" stemming from the fact that most of CA*net's lines were leased from CN/CP whose physical lines ran beside CN/CP railway tracks:

> That's how we knew about train derailments before anybody else did, other than the people on the train. When a train ran off the track, it would cut the line and there goes our network. Then we would hear on the news that some train ran off the track and say, "Oh, that's why it's down."

Ironically, the redundancy the network management board had worked so hard to ensure in theory was sometimes rendered moot in practice. What appeared on the network map to be a fat loop representing both the primary and the fail-safe secondary telecommunication lines was in some places two parallel lines laid just a few feet or even inches apart. So when a train jumped the track, it ended up cutting both lines and eliminating the prized redundancy.

The network information centre served as the interface between NOC and the rest of the network members, producing copious reports, charts and matrixes that gave a complete picture of network operations. For example, usage data was needed for the costing formula, which stipulated that population would determine 50% of each member's total costs until usage data was available. So UTCS and the information centre carefully monitored traffic and tracked it in a number of ways – origin and destination, application protocol, by the bit and by the packet – for use in the costing calculations.

Warren Jackson recalls that this service was a real point of pride at UTCS:

> Bob Chambers and Eugene ran and still run as professional a networking operation as you will find anywhere, period. The U of T staff are really, really good guys. Just amazing.... ONet's won a number of awards for customer service and professionalism... When they were running CA*net, they did things just as well.

6

FROM CA*NET TO INTERNET

BASICALLY WE DID THREE THINGS RIGHT. WE KNEW WHEN TO GET IN. WE KNEW HOW TO NURTURE IT IN SPITE OF THE GOVERNMENT. AND WE KNEW WHEN TO GET OUT.
*– GERRY MILLER, CA*NET INSTITUTE*

Between 1990 and 1993 the CA*net's managers and operators at U of T and in the provinces focused on establishing the basics of a stable, production-class, self-sustaining network. At the U of T network operations centre (NOC), this meant overseeing the installation of the remaining communication lines, debugging the network as necessary, and experimenting with the best ways to measure and report network activity. And at the headquarters of the provincial networks there was more of the same, though in a slightly more entrepreneurial vein.

CHALLENGES... AND TRIUMPHS

The raw power of CA*net as a communication and information resource attracted new users at a rapid rate. In June 1991 ONet provided Internet access through CA*net to over 40 organizations; a year later the number had almost doubled. Between 1991 and 1993 the amount of data traffic carried over the CA*net backbone more than doubled each year from around 20 gigabytes (20 billion bytes) per week to over 180 gigabytes per week.

Nevertheless, the ultimate success of the CA*net initiative rested on whether enough end-users could be secured to bear the cost of building and maintaining the national backbone and provincial networks. Building a customer base for Internet services in 1991 wasn't as easy as simply declaring the network open for business and providing competitive pricing. For starters, there wasn't anything comparable to compete against! What's more, the mandates of CA*net and the provincial networks weren't written to promote the widest possible use of the networks but rather to serve the research and educational communities.

Keeping technology high and costs low was a constant concern. And looming over the horizon was the March 1993 deadline, by which time the NRC grant would expire and the network had to be financially self-sustaining.

THE HIGH COST OF ACCESS

The provincial networks acted as early Tier 2 ISPs, offering access to organizations that in turn provided access to end-users. Drawing customers to them was less a matter of competition than of education. Even those groups that should have jumped at the chance for an Internet connection (those already on CDNnet or NetNorth) often had to be convinced to pay for the more expensive service. And those organizations that already had an Internet connection complained they had far too much trouble finding and taking advantage of the resources the network was supposed to make available.

Institutions in Canada that did connect to CA*net through their provincial networks paid between $15,000 and $50,000 for Internet access. The provincial networks in turn paid between $70,000 (PEInet) and $250,000 (ONet) for access to the "high-speed" backbone (predominantly 56 Kbps) – plus it cost roughly $1.2 million a year to operate the backbone. The US regional networks, on the other hand, had enjoyed at least T1 connectivity (1.5 Mbps) through the NSFnet for years, as well as other kinds of public support. The higher cost of bandwidth and lack of direct funding meant Canadian users had to put up with much slower connections to the Internet.

As CA*net treasurer Andy Bjerring wrote to the Board in 1992:

> The so-called "Field of Dreams" argument (build the network and the applications will come to fill it) has not yet had much effect in Canada, largely because our research networks are coping with tariffs for communications links that are 7 to 10 times those in the United States. As a result, the network capacity we can afford is still only that enjoyed by NSFnet users in 1986.

NEW TOOLS FOR NEW NETWORKS

A number of innovative applications were created in the early 90s that made the already-existing resources on the Internet more accessible and useful – especially to Internet newcomers. The Internet in 1990 was like a fantastically large library with thousands upon thousands of documents and no card catalogue. For the most part, Internet veterans from the 1980s knew what resources they were looking for and where to find them. Staying up to date with what was available on the network had been, for them, a manageable task: if you didn't know where a resource was, you simply asked a colleague or posted a public query on a newsgroup or some other network community. On the other hand, despite great strides by CA*net in making the physical network more accessible, "newbies" online for the first time had no idea where to begin.

By 1990, however, even experienced users were finding it impossible to keep track of network resources, manually or informally. That year, two McGill graduate students, Peter Deutsch and Alan Emtage, developed a software application they called Archie as part of their graduate work. Archie (a play on "archiver") automatically searched File Transfer Protocol (FTP) sites on the Internet and collated available file names in an archive of its own. Users could log onto an Archie server hosted at McGill and search the database for keywords or phrases. Successful hits turned up file names and FTP site locations from which the user could download the file. In short, Archie was a rudimentary version of an Internet search engine.

Archie helped fill the immense void of resources for Internet search and navigation, and was an unprecedented success. Under normal circumstances, Canada is a net importer of data: we download almost nine times as much data from computers located outside of Canada than computers located inside our borders. Thanks to Archie at McGill, however, Quebec's IP network (RISQ) was the first – and possibly the only – provincial network to become a net exporter of data, if only for a short period. Not long after its introduction, Archie became available on dozens of servers around the world, reducing some of the pressure on McGill's server.

That same year, Peter Scott and Earl Fogel at the University of Saskatchewan invented another groundbreaking and popular application known as Hytelnet. The name refers to its hyperlinked file structure and archive of Telnet sites (Telnet, one of the first network applications, allows users to log on remotely to another computer). Hytelnet integrated hypertext to let users jump from a link supplied by Hytelnet's searchable archive directly to a network resource. Employing a user-friendly format, the application compiled and made available hundreds of Telnet sites in Canada and around the world, including library catalogues, community computer networks and campus-wide information systems.

Then an Internet storage and retrieval tool appeared that, while technically just another top-level application, would eventually become synonymous with the Internet – the World Wide Web (WWW). Developed from an original proposal published in 1989 by Tim Berners-Lee at CERN, the European Laboratory for Particle Physics in Geneva, the WWW protocol (HTTP) used hypertext and hypermedia links embedded in Web-formatted documents to retrieve other Web-formatted documents, regardless of their physical location.

The Web gave content creators and seekers the most flexible tool yet created to make and access content even though, unlike today's browsers, the program was originally a simple text-based application. By 1991 Web-formatted files were circulating on the Internet, though for the time being it counted as just one of many online search and navigation tools, which included Gopher, Veronica, Jughead, WAIS, Listserv and a host of others.

Despite these advances, the Internet remained a text-based and often cumbersome tool – until 1993, when a piece of software saw the light of day that transformed the Web and, ultimately, the entire Internet. Mosaic was a sophisticated browser that made Web-based information much more accessible. Many Internet veterans at the time described using Mosaic as a jaw-dropping experience. In a sense, Mosaic did for the Internet what the graphical user interface (GUI) would eventually do for personal computing – replacing text-based commands with icons, pull-down menus and similar features.

Written by Marc Andreessen, then a young graduate student at the National Center for Supercomputing Applications, Mosaic's power derived from the graphical and relatively user-friendly window it opened on the still fledgling Web. Browsers for the Web had been mostly alphanumeric up to that point and couldn't take advantage of the Web's multimedia potential. Mosaic could and did, opening up new dimensions in Internet content.

Gerry Miller, CA*net's chair at the time, remembers his first encounter with Mosaic, which took place under rather special circumstances:

> We were visiting the National Science Foundation's headquarters ... and they asked us if we wanted to see something neat, called Mosaic. I said sure. I had heard of it but had never seen it... We went to [some European site] and looked at the pictures. That was shocking stuff back then – all these pictures slowly popping up.
>
> The funniest part was that there was this kid under the table pulling cables. It was Marc Andreessen – the guy who developed Mosaic. I only found out later who he was. At the time, he was just a fat little blond kid under the table showing us how it all worked.

OPEN FOR BUSINESS?

An especially difficult challenge for CA*net in the early years was defending its mandate. In addition to the old-guard R&D community, the CA*net backbone was now introducing tens of thousands of students and staff at educational institutions to the global IP network. There was, however, growing resentment from individuals outside these communities, many of whom wanted Internet access – but for purposes that didn't meet the acceptable use policies (AUPs) of CA*net and its member networks.

ACCEPTABLE USE POLICIES (AUPS)

The founders of the ARPAnet, NSFnet and CA*net initially viewed these networks as communications tools for academic and government research whose value would be compromised if access was thrown open to the private sector and recreational users. As a network funded in part by government for a highly delimited purpose, the NSFnet had a strict AUP that defined who could access the network and who could not. Most importantly, its AUP forbade commercial activity of any kind.

Since the NSFnet and CA*net were now interconnected at various points, they were committed to respecting each other's AUPs as much as possible. In practice, however, as the unofficial Usenet connection to the ARPAnet showed a decade before, policing AUPs was extremely difficult and network managers had to turn a blind eye to many infractions. In Canada, since users did not connect directly to CA*net itself, AUPs were largely enforced by the regional and provincial networks, whose policies varied a great deal. A case in point was the provincial network in Nova Scotia (NSTN), which was steering something of a unique path away from its CA*net colleagues – as a for-profit, commercial operation.

In 1989, Software Kinetics had won the contract to establish and run NSTN largely based on its past experience doing the same for DREnet. Mike Martineau, already a veteran in what was a brand-new industry, was charged with overseeing the Software Kinetics operation. In NSTN's first two years, Martineau slowly expanded its user base from university and government users to commercial R&D companies:

> There was no competition at that point. We were building the Internet in Nova Scotia. We were the only people to offer it, so it wasn't a question of "buy from us and not them." We were actively selling the value of the Internet itself. It was a very evangelistic activity.

Martineau made no bones about NSTN's open-for-business access policy:

> The joke going around the CA*net community was: "What's the NSTN acceptable use policy? Well, if you can fog a mirror and sign a cheque you can be a customer!" And I used to say, "If the cheque doesn't bounce, I don't care if you can fog the mirror!"

Joking aside, discrepancies between provincial networks and their AUPs caused tension at CA*net board meetings on a regular basis because of conflicts with the AUPs of networks outside Canada. Warren Jackson, the person responsible for CA*net's NOC, was keenly aware of the fact that a great deal of NSTN's commercial traffic in its early years broke the NSFnet AUP at that time. The problem was made more sensitive by the fact that NSFnet shared the cost of their north-south connections.

Martineau recalls that he and Jackson "ruffled feathers" many times over these issues, which were largely insoluble given the nature of the technology and the varying policies of the different jurisdictions. With the hindsight of history, however, it's clear that Martineau's business-like approach was a harbinger of what was about to become a tidal wave of commercialization sweeping over the Internet beginning in 1992-93.

UUNET AND ALTERNET

For individuals and organizations denied access by the provincial networks, the only option remained the comparatively slow and functionally limited networks such as Usenet that operated using temporary dialup connections between nodes over regular phone lines. These networks, which also included commercial online services CompuServe and America Online, offered an increasing number of gateway services such as email over the Internet.

But commercial ISPs looking to resell access to a provincial network and, by extension, the CA*net backbone, were turned away. Richard Sexton, coming from the far more liberal UUCP/Usenet community, remembers being particularly frustrated with the ONet crew because their not-for-profit philosophy meant that his first and best option for Internet access in Canada was denied him:

> CA*net was always an odd duck. They were a bit of a pariah because they would only sell to provincial networks, and the provincial networks would then sell to [R&D institutions]. There was a fair degree of almost hatred towards ONet and CA*net because they didn't allow downstream reselling.

CA*net and the provincial members didn't stay the only game in town for long. The first serious ISP serving Canadian businesses, UUNET Canada (UUNET stands for UNIX-to-UNIX Network), was founded on the heels of CA*net by Rayan Zachariassen in 1991. Zachariassen had started off as a U of T computer science student who found the challenges posed by the school's internal computer networks more interesting than his own graduate work. Along with Bob Chambers, Dennis Ferguson, Glen Macintosh and others, he had been one of the behind-the-scenes people at U of T responsible for much of the design and implementation of CA*net and its NOC.
But Zachariassen wasn't interested in simply maintaining an existing network. He knew the Internet was going to become an important part of the business world's infrastructure and that it wouldn't be served well, or at all, by his current employers. Seeing his opportunity, Zachariassen began offering Internet

feeds under the name UUNET Canada. The company was 20% owned by the US network provider UUNET Technologies, founded in 1987 as the first commercial provider of UUCP links to Usenet. Since that time it had become one of the first commercial ISPs, constructing its own network called AlterNet. UUNET Canada gained access to the US via AlterNet and soon constructed its own network facilities known as AlterNet Canada.

UUNET Canada was one of CA*net's first facilities-based competitors, though the two networks had quite different mandates, and it certainly wasn't the last. New competitors such as fONOROLA soon entered the scene and commercial pressures created new challenges for the young national network.

CANARIE ON THE HORIZON

On the government side, the official voice for commercial interests in the exponentially expanding Internet continued to be the federal department of Industry, Science and Technology Canada (ISTC). While CA*net used its NRC funding to establish an IP-based production network, the ISTC continued to work towards the much grander vision of the "Information Highway" – a phrase that was increasingly on the lips of policy-makers, journalists, telecommunications carriers and cable companies.

The results of the feasibility study conducted by ISTC regarding government backing for the project were encouraging. In the spring of 1990, the government gave Digby Williams approval to start planning the ISTC test network, with one catch – the project required getting Canada's academic networking community on board. In September, after a conference at Deerhurst Lodge in Ontario that brought all the stakeholders to the table, Williams got grudging consent to proceed.

Soon afterwards, Williams left ISTC and the test network project sat until January 1991 when Jocelyn Mallett was appointed director general of ISTC's Information Technology Industry branch, where Williams had been working, and took the project to the next stage. After seeking advice from CA*net representatives, she held her own network organizing conference in Ottawa in April of that year to develop a business plan that would secure government funding for the proposed high-speed test network.

The working groups established at the April, 1991 conference submitted their proposals in January 1992, and the Business Plan for what was then being called "CANARIE", or "Canadian Network for the Advancement of Research, Industry and Education", began to take shape. The project was to have three phases and several program elements. As far as the network was concerned, Phase 1 was to start in 1993 with an upgrade to CA*net to T1 (1.5 Mbps) or higher capacity across Canada. A further upgrade to T3 capacity (45 Mbps) was to take place by December 1995. A separate national test network running at 155 Mbps was also to be developed using Asynchronous Transfer Mode (ATM) technology (in the mid 1990s ATM was seen as a fast, highly scalable transport protocol that could handle data, voice and video traffic using "cells" of fixed size; it was in part the telecom world's response to the vision that was driving the Internet). Finally, in Phase 3, or roughly by the end of the decade, once the much-vaunted Information Highway had taken shape in the form of a fully upgraded production network, the government would withdraw from the project.

MORE THAN AN UPGRADE

The proposed upgrade for CA*net was certainly attractive. Although the provincial networks had enough customers to stay afloat, the bandwidth on the national backbone had fallen far below standard. Half of the network's links remained at 56 Kbps and many were severely congested, but the cost of upgrading to T1 was not affordable with the NRC funding, which was to run out in 1993, or on the basis of the fees paid by the provincial networks.

The CANARIE plan, however, only addressed the upgrade to CA*net's infrastructure, and did not help with other challenges facing the CA*net board. For one thing, CA*net's coast-to-coast public-service mandate and its cost-sharing price structure were becoming problematic. As commercial ISPs slowly began to enter the field, especially in the cities with the highest demand for Internet access, CA*net was faced with the possibility of regional networks being offered better facilities at a lower price. If CA*net didn't get a subsidy to continue the east-west backbone beyond the end of the NRC support, many of the provincial networks would likely take advantage of north-south connections to the US backbone, NSFnet.

Cost-sharing and serving the public interest work well when the service provider operates as a regulated monopoly, as had the telephone companies for many decades. But CA*net was far from a monopoly: it was a government subsidized network with a narrow mandate whose funding was about to run out. With no other patrons on the horizon in the spring of 1992, a grant as part of the CANARIE project appeared to be the only ray of hope.

Unfortunately, funding still had to be obtained for the CANARIE plan, and ISTC still required provincial support before committing any funding of its own. Mallett worked hard over the summer shopping the proposal around the provinces, but the learning curve was apparently too steep. Mallett left half her meetings feeling that her hosts hadn't really understood what she was talking about. In the meantime, the project had caught the interest of

Canada's two national telephone networks: the Stentor consortium of incumbent telephone companies led by Bell, and new entrant Unitel (formerly CN/CP Telecommunications, later AT&T Canada).

TELCOS AT THE READY

ISTC valued CA*net as a network and appreciated the experience of the CA*net board and the regional networks. Nonetheless, it did not see CA*net Networking Inc., the not-for-profit, federally subsidized corporate entity, as the right group to take networking, specifically the Internet, into the commercial arena, especially when the private sector was beginning to show interest in running such networks.

To a large extent, the CA*net group agreed. CA*net hadn't been built as a networking business. It was simply the quickest and only made-in-Canada way to establish the network required to support research and communications activities. Nevertheless, they were also aware that they, and the academic and R&D communities they served, still had knowledge of and experience with IP networking that the private sector largely lacked. They were far from ready to leave their mission-critical operations in anyone else's hands just yet.

One of the main reasons the academics had started CA*net in the first place was that the telcos had no apparent desire or ability to do so themselves. But the telcos were now paying close attention to CA*net, for several reasons.

First, the government was inspiring a great deal of interest through its direct and indirect support of the development of high-speed, packet-switched networking. Experts in other countries pointed to this technology as a crucial means to stay competitive as a nation in what was becoming known as the "information economy". The government saw its role as encouraging industry, in this sector as elsewhere, to take the financial risks necessary to meet the challenges posed by the emerging economic order.

Second, a great deal of the talk about packet-switched communication was not simply about its role in the networking of networks, but also about its role in the development of what later came to be known as "convergence".

Traditionally, communications networks had been optimized for one kind of service – notably telephone networks for voice service and cable networks for television service. The Internet was radically different. It was a "dumb," all-purpose platform designed to deliver data packets of whatever kind to their final destination as quickly as possible. On a fast enough platform, the Internet could conceivably deliver telephone service as easily as TV service, or any other information that could be digitized and sent down a wire. This possibility had caught the attention of the phone companies – whether they appreciated or understood packet-based networking or not – because it meant both opportunity (new services to sell) and danger (new services, including traditional telco services, being sold by their competitors).

Third, selling Internet access, whether to individual end-users, businesses or other ISPs, looked like it would become a profitable business in its own right. If anyone was going to make money selling new services over the networks operated by the telcos, it should be the telcos themselves (or so they reasoned). In short, the telcos were finally thinking about a move into packet-switched data transmission.

Meanwhile, ISTC's funding bottleneck had not opened up. Having exhausted the provincial possibilities, Mallett took a new approach: if she could win the support of Canada's two major telephone networks, this backing from industry would be sufficient to replace the required provincial buy-in. She did in fact succeed in getting arch-rivals Unitel and Stentor to sign a joint endorsement (on one-of-a-kind joint letterhead), and their commitment finally broke the financial logjam.

In December 1992, it became clear that the federal government was preparing to commit $26 million to start up the CANARIE project. By early in 1993, CANARIE Inc. had been incorporated, and had as founding members both universities and many of Canada's major telecommunications and high-technology stakeholders. By June, the formal agreement between CANARIE Inc. and the federal government was signed, and the long-awaited bandwidth upgrades to CA*net could begin (the allotted budget was $5 million). CANARIE's mandate further included the development of an ultra high-speed experimental network for testing next-generation networking technologies and applications, which could then be integrated into the evolving CA*net infrastructure.

By its own Bylaws, half of the CANARIE Board had to represent institutions like universities and other not-for-profit organizations. At its inaugural Annual General Meeting that June, the membership of CANARIE elected several members of the CA*net Board to the new CANARIE Board, including Andy Bjerring, David Macneil, Jack Leigh and the recently appointed President of CA*net Networking Inc., Ken Fockler.

FOCKLER MEETS HIS FATES

Throughout the summer of 1992, it had become evident that managing CA*net was too tall an order for the CA*net board, all of whom had day jobs.

A cloud of uncertainty hung over not just the immediate issue of financial survival but the long-term goals and definition of the network. Should the board keep it on the straight and narrow path of government subsidy, while upholding those provincial AUPs that excluded commercial and consumer clients? Would some industry clients be acceptable (for example, companies engaged mostly in R&D), but not others? What would CA*net's role be in relationship to CANARIE?

Moreover, ever since the U of T had made clear to the NRC management board that OSI compliance was simply not a practical short-term possibility, the board had heard very little about the OSI requirement. In short, for those CA*net members who still held the torch for an autonomous, public-sector IP-based network, the light of hope was not entirely extinguished. On the other hand, the general consensus was that the network could not be self-sustaining in its current form, and would have to adjust to the changing commercial realities of the Internet environment.

What CA*net needed was a full-time organizer to navigate these difficult waters and oversee the necessary transitions – someone who could talk business with the best of them but also felt comfortable with the academic and research communities, and who could somehow bridge the growing gap between commercial and non-commercial interest in the Internet. One name came immediately to mind.

Although he had not been directly involved in the CA*net initiative, Ken Fockler had stayed close to the networking community, especially through the annual Net conferences which IBM continued to help organize. But by 1991 IBM found itself in a very different computing industry than the one it helped build and had dominated for so many years. Ironically, computer networks such as those IBM had helped support were one of the disruptive forces behind these changes.

Fockler's position at IBM had been created at a time when signing and maintaining large contracts for expensive mainframes had been the staple diet of the computer business. But now the mainframe market was being overtaken by flexible, just-in-time, service-oriented production that dispensed with all unnecessary overhead, and the company was undergoing a massive reorganization. Fockler recalls a difficult time:

> In late 1991 IBM offered an early retirement package which wasn't exactly an offer: it was the door or working in the plant somewhere!... They made us a blanket offer, a good offer. But it was scary because you'd been with IBM for so long, as somebody described it, wearing golden handcuffs. They treated you well with benefits and a good salary and other perks, and you simply didn't break away. Was there even another world outside? I was about to find out.

Fockler spent an uneasy Christmas with his family and returned in the New Year resolved to make a go of it as an independent consultant. Although he met with some early success in organizing a major project for the United Way, Fockler wouldn't remain "independent" for very long.

Embarking on the search for a full-time executive director, the existing executive team on the CA*net board – Gerry Miller, Andy Bjerring and David Macneil (the Three Fates, as Fockler nicknamed them) – made up the selection committee. Macneil advertised the position and, in September 1992, Fockler received a call from a former colleague at IBM's New Brunswick office. Macneil had a message for Fockler: take a look at the job CA*net was advertising in The Globe and Mail – it was a good fit.

Fockler was dubious at first. Did Macneil want him to move all the way out to New Brunswick? Even after reading the ad and realizing the job was actually located in Toronto, he was still a little doubtful:

> I read the ad and it read very technical. I thought, I'm just not that technical. It's all communications stuff. But I called David and he encouraged me to apply anyway.

The four men met in Toronto for an interview. It had come down to a short list of two applicants: one offered a strong technical background and the other was Fockler who, as Miller describes it, "could schmooze with the best of them and was not administratively challenged." The group decided they had enough technical expertise already; it was time to diversify their strengths. They hired Fockler. Looking back, Miller comments: "It was probably the smartest thing we ever did."

PEERING PROBLEMS

By the time the hiring committee brought Fockler on board, they were under few illusions about the ultimate fate of their network. Both government and the major private-sector players were taking a greater and greater interest in data networking as a tool not simply for research, but for economic and industrial development that would benefit a much wider constituency than had been served up to this point by CA*net. Eventually, it seemed to them, their pioneering work would have to be ceded to these more powerful interests. But they were also determined to see that the transition preserved their legacy and enhanced the benefits CA*net had been built to provide. Fockler's mission was to ensure that the whole process took place as smoothly and constructively as possible.

Among the most important issues Fockler and the board had to address at the start of 1993 was the question of peering. CA*net had received requests from various network providers to peer with them, i.e. grant mutual access to each other's infrastructure at no charge in order to exchange traffic. In the early days of the Internet, peering between networks played a crucial role in expanding and consolidating the global networking system. Peering results in more robust and efficient network operations because data traffic has more options when traveling to its final destination – a vital attribute in a packet-switched network, since data packets do not follow pre-determined routes.

But the equity of CA*net's cost-sharing formulas, intended to allow reasonable access across the country as part of its nation-building mandate, would be at risk if CA*net peered with other networks. And yet, the longer CA*net closed its doors to commercial or non-commercial peering connections, the more incentive there was for end-users – and even the provincial networks – to go around it, taking long, circuitous routes via US networks, thereby encouraging inefficient use of both Canadian and US network facilities.

As noted above (see section on "Acceptable use policies"), a microcosm of the future could be seen in the evolution of the Nova Scotia provincial network. In 1992 NSTN introduced dialup service using phone lines and desktop modems, the way the majority of homes still connect to the Internet today. Creating the dialup service took some creative technical work-arounds, but NSTN was able to slash the price of access from $400 a month to $75 a month. Over the next six months the price continued to drop – by about 50%. The result, not surprisingly, was a significant increase in Internet use, as Martineau recalls:

> Can I point to any one event? No... slowly, about '93 on the dialup side, it started to get public press, people started to hear about it and the tide shifted. People would call up and say, I want to buy Internet access, what's it cost? They still had questions, but now they were calling us.

NSTN's pioneering efforts put more pressure on CA*net to re-examine its AUP, at least as far as the other provincial networks were concerned, its peering policies and its network efficiencies.

Typical of CA*net's soul-searching on this issue was the debate on peering with representatives from UUNET Canada. Despite an expanding network of its own (Alternet Canada) and a connection to AlterNet in the US, UUNET Canada was interested in peering with CA*net and proposed it in early 1993. But CA*net had already refused a peer-to-peer networking arrangement with DREnet, and took the same position with UUNET.

A series of public email exchanges between Andy Bjerring on behalf of CA*net and Tom Glinos and Sean Doran from UUNET tells part of the story. At first Bjerring explains that CA*net's reluctance to peer had more to do with their own organizational philosophy than the qualifications of any prospective partner:

> I think the cost of creating and maintaining a "true" national network in Canada is such that we need all those who benefit from its existence to make a contribution to it. CA*net and the regionals that it connects constitute the only such network that I'm aware of, with points of presence in places that DREnet and UUNET have no intention of ever spending money.

At this point Doran joins the debate. Doran was UUNET Canada's chief networking engineer at the time, and later went on to manage Sprint's national IP backbone. A networking guru and flamboyant personality who sported a bright pink Mohawk during his time at Sprint, Doran took issue with Bjerring's claim that CA*net was willing to reach places a commercial entity wouldn't:

> Really? Name five. We are quite prepared to drop a Point of Presence anywhere in Canada we can get bits to. The only difference is in the pricing scheme – we simply believe that people in outlying areas should pay their own way, and not lean on subsidies gotten from sites in larger cities like Toronto and Ottawa.
>
> While the ONet/CA*net model fits with one university-controlled network in each province, it does not seem to appeal to industry, and especially not to small business. Each network does what works for its members/customers. Neither model is the One True Path.

The peering debate would not be resolved in the short term. The existing situation didn't prevent either party from operating and it was clear that both suffered from a failure to communicate. As Bjerring points out, there were also more practical financial reasons behind CA*net's reluctance to peer:

> We had to build the base of users in southern Ontario in order to share the cost of the long haul links up to the north. This was a promise that we essentially made when we amalgamated NetNorth central and ONet... There was some fear that if we didn't keep the south together, that if it started to splinter, then we would never be able to live up to that commitment to the north. And the market never would have done this... It's not as if you can argue the virtues of the marketplace when there isn't a market yet.

At the end of the exchange with Glinos and Doran, Bjerring remarked to Fockler, Miller and Macneil:

> As you can see I've been fighting the "good fight" on behalf of truth and justice, if not beauty. The "peering" model does seem to represent the antithesis to the "consortium" model that we're trying to push.

This point was lost in the conversation with UUNET. The fact was that CA*net saw itself as a consortium of ISPs sharing the cost and responsibility of linking their provincial networks. These links constituted a national backbone, but from CA*net's organizational perspective that was almost incidental – it didn't see itself as being a backbone first and a consortium of ISPs second. CA*net's success was still synonymous with the success of the consortium of provincial networks that it connected. Opening the network to peering was therefore incompatible with their consortium model of networking.

The issue would be revisited by the CA*net board, as was CA*net's own self-definition, in terms of whether or not to allow commercial access. In the meantime, the exchange with Doran troubled Bjerring:

> One thing that Sean says is probably quite true – that the lack of a peer connection hurts our users more than it does theirs. I don't know how to resolve that paradox.

WHO RUNS THE SHOW?

The commitment of $5 million in the CANARIE plan to the upgrade of CA*net appeared to resolve a major problem for the CA*net board, although nobody was certain what the relationship would be between CA*net and CANARIE. As an autonomous, incorporated entity that owned and operated a valuable piece of infrastructure, the CA*net board was concerned about the future ownership and management of the network. From this perspective, the board recognized that, despite the promise of much-needed funds, CA*net needed to explore its options.

Fockler tried to query the Chair of CANARIE 's new Board, Bill Hutchison from Ernst and Young, about the relationship between the two organizations. Although he received some answers, they didn't put to rest the concerns of the CA*net board.

Hutchison's letter began by acknowledging the contribution Canada's academic and research community had made in building CA*net. It went on to say that, for a number of reasons, Industry Canada (IC) (the re-named ISTC) felt it practical to focus on one network and one organization – CANARIE Inc. – rather than to try to maintain a separate and additional relationship with CA*net. The letter also pointed out that CANARIE remained a joint effort, and the CANARIE board didn't want to displace the existing knowledge and experience represented by the CA*net group. Hutchison invited CA*net's members to help lead CANARIE and help expand the networking community they had fostered.

The letter did not address Fockler's main concern: What were CANARIE's intentions regarding the future ownership and operation of CA*net? The fact that CANARIE was the designated source of government funding for the upgrade of CA*net said nothing about who was to control the network and its evolution. After all, as Fockler later pointed out to the CANARIE board, CANARIE was also funding the development of the test network in conjunction with the telephone companies, and yet IC wasn't expecting to own any telephone companies at the end of the project!

In August 1993, Fockler met with Carmelo Tillona, CANARIE's first president, and Rafiq Khan, who had worked with both Digby Williams and Jocelyn Mallett and had recently jointed CANARIE after retiring from Industry Canada. As Fockler reported to the CA*net board:

> The Good News is that Carmelo is very aggressive and wants to make things happen right away. The Bad News is that Carmelo is very aggressive and wants to make things happen right away!
>
> He is also saying from his conversations... that [the future version of CA*net] could be and should be run by business and specifically the carriers, not the resellers.
>
> No real problem with the above.
>
> New News: he is also saying that he has made real progress with Stentor and Unitel, and they agree and say they want to, and will run such a network. I'd like to see and hear this myself.

So did the rest of the CA*net group. Jim Hancock, PEInet's representative on the CA*net board, echoed the collective opinion on the telephone companies' appreciation of packet-switching:

> In all this we must remember that initially the telcos weren't interested and it was the universities that had the know-how and desire to make [a packet-switched network] work... In dealing with Island Tel it has taken most of a year to get them to understand how CA*net works. Even now I don't think the top management really understands it. They still think networking is point-to-point.

But the strongest warning came from Jack Leigh at BCnet, which served a region with a history of marching to the beat of its own drum:

> Here's my reaction to your message. My initial and biggest reaction is concern. I'm concerned that Carmelo and Rafiq really are not very knowledgeable with regard to this type of networking.
>
> The one message that must be made very clear to everyone in CANARIE is that the operations (research and education) of my institution are dependent in a moment-to-moment manner on reliable, cost-effective access to the Internet. We cannot do anything that will jeopardize that access, nor can I be involved in any-

> thing that might. If I get concerned about what's happening I will take my ball (BCnet) and play in another game (CIX, fONOROLA, Northwestnet or whatever).

Meanwhile Fockler was turning his attention to another option for ensuring CA*net's future – expanding the network's client base. Meeting with various government departments he invariably found that a great deal of explaining was required to clarify what CA*net was and did, and how it was different from CANARIE and the government's interest in the Information Highway. Fockler's challenges were compounded by the fact that he didn't have all the answers himself. He had to explain that support for the somewhat utopian "Information Highway" didn't necessarily translate into immediate support for CA*net, a functioning production network that people actually depended on right now. The meetings produced little by way of tangible results.

A NEW PRESIDENT FOR CANARIE

In September, 1993, after little more than two months on the job, Carmelo Tillona resigned the position of President and CEO of CANARIE Inc. While Pat Sampson, one of Jocelyn Mallett's directors from IC took over as acting President, the Board search committee that had hired Tillona was hastily reconvened, and one of its members, Andy Bjerring, who was still treasurer of the CA*net Board, was prevailed upon to throw his hat in the ring. His first term as CANARIE's second President and his leave from the University of Western Ontario began in October. From that point on his attendance at CA*net Board meetings would be as a CANARIE observer.

One of the first tasks of the new President was to meet with Fockler wearing his new hat, and to try to address the concerns that he knew only too well remained at the heart of the relationship between CANARIE and CA*net. The two concluded a Memorandum of Understanding (MOU) by December that was later ratified by the two boards. Although no formal document was going to instantly resolve the tensions between the two organizations, the agreement was a step in that direction, and was sufficient to allow the funds to actually flow from CANARIE to CA*net for the first time.

NET (REALLY FAR) NORTH

Early in 1994 Fockler was alerted to action being taken by Jim Tousignant in Whitehorse to build a Yukon regional network. Tousignant had been canvassing start-up support for an operation that would become a self-sustaining entity like its provincial counterparts once a user base was established. His request was passed on to Fockler, who was enthusiastic about the project. After all, as Bjerring had stated in his email to UUNET Canada, part of what distinguished CA*net from other ISPs was its willingness to deliver access to areas where UUNET and the rest had no intention of spending any money. The Yukon certainly qualified as one of those areas and, while they were at it, the CA*net board could see about establishing a POP for the Northwest Territories as well.

Communications facilities in the territories were minimal by computer networking standards, and included reliance in part on high-frequency wireless connections. While most communities enjoyed phone service, the phone lines could only support a maximum bandwidth of 9600 bps – and only then after expensive line conditioners were applied. In Whitehorse, connecting to the Internet was achieved by dialing south to the nearest Internet service provider – another expensive solution. In Yellowknife, conditions were a bit more advanced: a 19.2 Kbps line had been established to ARnet, Alberta's provincial network. The territorial government had also established a local X.25 network that connected 10 communities using leased lines from NorthwesTel, the local phone company owned by Bell Canada.

The NWT also had an informal group interested in establishing a regional network.With Fockler's encouragement, both northern groups formalized their activities. They also developed business plans in support of grant applications for hardware upgrades from a separate funding program of CANARIE. Tousignant led the effort in the Yukon, while Bob Johnson, president of a local computer company, took charge in the Northwest Territories.

Fockler flew North in the mid-summer of 1994 to get a first-hand look at what was going on and help participate in some consciousness-raising. The level of interest in CA*net's networking efforts surprised him. Arriving with six or so scheduled meetings, he attended four times that number by the time he left:

> Even when I took a few days off, I would find myself talking to a teacher about the educational uses of the Internet, an art gallery operator in Whitehorse about future commercial applications and the operator of a bed and breakfast in Dawson City about tourism applications. All wanted to know when they could get started.

Connectivity to southern Canada was a major issue for these communities. Among other benefits, access to the Internet meant valuable services in areas such as education and health. Fockler's visit seemed to have a catalytic effect, encouraging key organizations that hadn't been certain about the project to come forward with support. He himself was able to provide assurances that CA*net would undertake to install lines to a CA*net POP in both territorial capitals by October.

When Fockler left the North, the Yellowknife group had worked out internal differences about how to proceed and decided to act as a non-profit rather than for-profit entity. In Whitehorse, representatives of NorthwesTel surprised everyone by declaring financial support of their own for both new territorial networks. By late fall 1994, leased 56 Kbps lines from NorthwesTel via INSINC were successfully installed. The addition of YukonNet and NTnet allowed CA*net to boast of being a truly national network.

THE BIG SWITCH

During 1994, with the $5 million in funding from CANARIE secured through the MOU negotiated between Bjerring and Fockler, and with the T1 upgrade in place, the CA*net board began to focus on the next needed upgrade, eventually to T3 capacity. At the same time, the CANARIE board was working on the new business plan needed to secure the second round in funding from Industry Canada that would make the upgrade possible.

In December, 1994, with second round funding for CANARIE confirmed by the Minister of Industry, John Manley, the CA*net board issued an RFI requesting a concrete plan to address CA*net's future bandwidth needs. Beyond being supplier of lines, the Board wanted to hear from possible partners who would help explore all the options facing CA*net and work closely with the board to map out its future evolution.

A short list was drawn up from the responses received and four applicants were invited to an interview to discuss their proposals. In the end, Bell Canada won the contract – though not without a little help from some well-placed partners.

BAC: BELL ADVANCES CAUTIOUSLY

After seeing the RFI, Bell Canada had realized this was its best opportunity to jump into the IP marketplace on a national scale – an opportunity to build the most advanced network of its kind, and one that would serve a pre-existing customer base reaching from coast to coast. But Bell also knew that, on its own, it could not demonstrate enough familiarity with or knowledge of the existing network or its current users to guarantee it would win the contract. So Bell ended up working with two partners, both with close connections to CA*net.

The task of representing Bell's interests and winning the contract fell to Bell Advanced Communications (BAC). One partner, surprisingly, was UUNET Canada, which might have been expected to compete rather than cooperate with Bell in bidding for a contract it was capable of winning. The next step was to approach Eugene Siciunas at UTCS. This relationship would reassure the CA*net board that BAC would provide a smooth transition period during the switch-over of service providers from UTCS to BAC.

Siciunas remembers how Zachariassen of UUNET approached him to create this unlikely trio:

> He approached us, playing marriage broker, and asked if we'd like to work together with him and Bell Advanced Communications to submit a combined proposal, like the one U of T prepared with IBM and INSINC in the early days. But this one would recognize that this was not an area that U of T would stay involved in for the long term. I think Rayan was hopeful UUNET would pick up some of the business itself. Bell was speaking with a forked tongue here and said, "Sure, sounds good," even though they had every intention of booting UUNET out as soon as they could.

UUNET was adding its experience as an Internet service provider to Bell's bid. And, of course, Zachariassen's experience as a senior systems architect on the original CA*net network was duly noted in the response.

The BAC consortium submitted a detailed proposal in response to the RFI that included an ATM-based backbone across Canada. The CA*net board accepted the proposal in March 1995.

A month later, BAC expanded the proposal to include two virtual networks on the same backbone. One network would be CA*net, with an immediate upgrade to 10 Mbps and then to at least 45 Mbps by 1998. The second network would be CANARIE's National Test Network (NTN), which had been created in 1995 as another ATM backbone linking provincial networks that supported the development and testing of advanced broadband applications. BAC promised to have the ATM spine and the initial upgrades to CA*net in place by September.

The job BAC had taken on was formidable. Not only was BAC converting CA*net's links from INSINC's lines to Stentor-owned lines, but they also had to switch the NOC gradually from UTCS to BAC. Siciunas recalls that the process got off to a bumpy start with his own staff at U of T:

> When our staff got wind of all this they said, "Oh shit," and just took jobs elsewhere... This left us with a network and no staff to run it. Bell tried to throw in whatever expertise they had, which wasn't all that great, and UUNET actually jumped in and provided Guy Middleton, a very knowledgeable guy they had on staff. In panic mode, UUNET was helping put the fires out and BAC was cranked up, getting ATM deployed across the country.

Many observers wondered if Bell could do it, including Siciunas:

They did pull off quite a trick, and they did get an ATM network coast to coast, and I think we were done by October, which amazed me more than most. I don't know how many people realized just how monumental a task that was. They proved they did have the resources to throw into it to pull it off after all. It was amazing... But I still think we ran it better than they do.

TOUGH CHOICES

Several important developments in the mid-1990s forced the CA*net board to take stock and finally decide on the future of the network. BAC was now operating the CA*net backbone, but CA*net Networking Inc. still retained ownership and control – and the question "why" began to haunt the group. Ultimately, that question came down to how best to ensure dependable, affordable connectivity for CA*net's members.

The Internet at the end of 1995 was unrecognizable compared to early 1993 when CA*net first received funding through CANARIE from Industry Canada. Private sector and home use were now the fastest growing sector of the Internet in both Canada and the US. CA*net itself had taken small steps to expand its client base beyond the public sector and was providing services to a few telcos and related companies.

South of the border, the NSF, CA*net's long-running networking partner, had left the production network business altogether in April 1995. The NSF sold off its production network services (Sprint and MCI took over most of them) and focused instead on high-speed test networks. Rather than one backbone network serving the whole nation, network interconnections were then handled through Network Access Points (NAPs) and a commercial Internet exchange (CIX) – a critical link to the smaller, private networks that previously had not had access to the NSFnet.

Canada's own commercial ISPs continued to grow as a strong community of their own. Throughout 1995, Fockler found himself in the midst of yet another debate with UUNET Canada about the peering issue, but this time in person with Zachariassen and Glinos. In a moment of shared clarity, the two sides began to understand each other and work towards a mutually acceptable solution. What they planned in partnership with iSTAR and others came to be known as the Canadian Internet Exchange (CDN-IX), and would benefit the ISP industry as a whole.

All these developments came to a head at a CA*net board meeting in

December 1995. The group reduced their options to four:

1. Preserve the status quo. CA*net continues to recruit and serve customers from the R&D, educational and quasi-private sector communities as an independent entity, possibly in cooperation with CANARIE and Bell through a joint marketing agreement.

2. Form a virtual company. Like the status quo except that BAC would take on the financial management of CA*net. The backbone service would run as a BAC operation, but with a CA*net "front window."

3. Become a for-profit company. Reincorporate and go full blast into competitive marketing.

4. Fade to black. BAC would provide transit service to CA*net customers until CA*net Networking Inc. wound itself down. Existing CA*net customers would then have the option of becoming BAC customers, or not, as they chose.

Options two and three (especially three) were daring, putting CA*net firmly into commercial territory. CA*net might well have a role to play in the commercial market as an AUP-free backbone that would serve as the point of interconnection and become the main thoroughfare for internal Canadian traffic. However, CA*net was not created to be a business, and serious doubts existed as to whether it was possible for the organization to pursue this path.

For CA*net to pay for the technology required to remain a competitive valuable resource, it would have to open its network up to a much broader user base. But pursuing such a move would have changed CA*net into something it wasn't intended to be. Complicated plans to have two tiers of membership and rates, one for commercial and one for non-profit organizations, were also considered. The overriding sentiment in the end was that the universities had not intended to get into the carrier business, and the potential benefits from transforming CA*net into a commercial operation were not worth the effort, risk and potential disadvantages.

Option one, trying to carry on, had its own advantages and dangers. The most obvious downside was the strong possibility that the CA*net group would lose ownership and control over the network itself if revenues did not match costs, which was highly likely if CANARIE and the government refused further subsidy to what would then have been a commercial venture.

This option was the one closest to their hearts and the toughest to let go of. In the best of all possible worlds, they would have continued to operate their network themselves, as they saw fit. But they also felt they couldn't ignore the pressing realities of the world they were operating in. CANARIE and the government were not going to continue to fund network upgrades if CA*net chose to operate as a cooperative, non-profit entity providing Internet services to Canada's R&D. CA*net, ultimately, would not be seen as providing its constituents with any benefits that could not be equally well delivered by the private sector.

By a painful process of elimination, that left option four as the only real possibility. Nevertheless, the final decision to wind down the corporation was a source of controversy that had dragged on for months. In one of the earlier meetings on the subject, Jack Leigh, who was as passionately against the decision then as he is now, recalls thinking he had actually turned the tide:

> I didn't think it was right that we put the network into a commercial sphere. If Bell had wanted to build a transit service they had all the means in the world to do so. No reason why the universities couldn't continue to operate for their own best interests and for the future development of the Internet. And most of the people around the table actually spoke in favour of my point of view.

Those who finally voted to wind down the network were sympathetic to Leigh's arguments; but reason won out over passion. A motion to accept option four was passed in February 1996, and a tumultuous era began drawing to a close.

"CA*net Networking Inc. ceased to have any responsibility for the operation of the national backbone on April 1, 1997. It was entirely up to each of the original regional networks to decide whether or not to use the new transit service offered by BAC. As it happens, they all did."

THE PASSING OF THE MANTLE

The CA*net crew couldn't just turn out the lights and leave, however. Many loose ends had to be taken care of before the organization could formally dissolve, a step that was planned for the board meeting and member's meeting on June 23rd, 1997 in Halifax.
Among these there was one that had a significant silver lining.
Gerry Miller explains:

> The question became, How do we get something out of all this? A lot of people had put a lot of work into CA*net. It's non-profit, so we couldn't sell it for money. We had a board meeting and said to BAC, you can't give us money but you can create a foundation to support the original aim of CA*net: to help build a Canadian Internet for the social good and the economic good. And they agreed. We used the money to create a thing called the CA*net Institute.

Since it was created in 1997, the CA*net Institute has funded twelve deserving projects and is still operating. In 1999 it became a stand-alone program operated by CANARIE, with Gerry Miller continuing as the program chair. This history is one of its projects.

CA*net also had a hand, again in partnership with Canada's larger commercial ISPs, in the founding of the Canadian Association of Internet Providers (CAIP) in March 1996. CA*net helped the association get off the ground by contributing both money and the time of its own president, Ken Fockler, to serve as the association's first chair and president. CAIP continues to serve as a group promoting the interests of its members on issues such as censorship and fair competition.

In many ways, however, the original mandate of CA*net lives on most directly through CANARIE. Over the years since 1990, the organization that was once viewed as a competitor to CA*net has been transformed. The same week as the board of CA*net Networking Inc. formally voted to dissolve, CANARIE inaugurated CA*net II, an ATM-based "intranet" for the research community in Canada, combining features of both the National Test Network and the CA*net backbone. This was a world leader. That network has since been replaced again with CA*net 3, the first all-optical, research network in the world. Many of the people mentioned earlier in this story worked on these efforts, as implementers, architects, policy makers and leaders. Gerry Miller is currently treasurer of the CANARIE Board and chair of CANARIE's policy committee for the network. Andy Bjerring continues as President and CEO.

CANARIE has also provided financial support to the provincial research networks, including several that were part of the original CA*net community: ONet, BCnet, the NB/PEI ECN and RISQ. Others are newly created versions of the original CA*net networks, now focusing not on providing connectivity between institutions and the commercial Internet, but on providing very high bandwidth network services in support of research and educational collaboration and innovation.

The history of CANARIE is still being written. Moreover, some of the scars earned in the battles between OSI and IP, and between the vision of a network owned and operated by and for the academic community and one developed by the private sector, are still fresh. Nonetheless, CANARIE's most important contribution might yet come to be seen as its ability, when the time was right, to build on the foundation provided by the CA*net community. Academic networking in Canada has always worked that way, from NetNorth, to CA*net, to CANARIE. When the time is right, the mantle is passed.

But, for all that, CA*net was special, and none of its founders and builders can help but feel that something special ceased to exist on June 23rd, 1997.

EPILOGUE

For the founders and builders of Canada's first computer networks, the winding down of CA*net certainly marked the end of a very special era. The end of CA*net didn't mark the end of the story, however.

Canadians in the research community are perhaps those receiving the most direct legacy from CA*net, with CANARIE's advanced technology successors to the original network, currently CA*net 3 together with the provincial advanced networks that provide the actual institutional connections, continuing to provide sophisticated network services for this sector. Moreover, CA*net 3, which went into coast-to-coast operation in 1999, was the world's first nation-wide optical R&D network, and since its inception has been emulated around the world. The planned CA*net 4 is expected to extend Canada's leadership into the next generation of technology and architecture. Quite a legacy for a network that in 1990 was five years behind the Americans!

Of course, the staggering increases in the performance of microchips lies at the heart of these technological possibilities. According to the "law" published 30 years ago by Gordon Moore, technological developments will continue to drive a doubling of the number of transistors that can fit on a silicon microchip roughly every 18 months, which in turn will drive improvements in virtually every aspect of performance (e.g. clock speed, storage capacity, etc.). In fact, with the migration from electronic to photonic technology the capacity of networks is growing at an even faster rate than Moore's Law would predict. The doubling time for network capacity is down to every few months, due to advances in optical networking technology as well as increases in microchip performance. If this trend were to continue, and it appears it will, the cost of provisioning a network will soon be largely independent of bandwidth.

In a wider context, CA*net can be seen as Canada's introduction to the Internet, and one sign of the contemporary influence of the Internet lies in the extent to which many Canadians now take being connected for granted. At the time this is written, more than half of all Canadian households have Internet access, and Canadians are among the most avid users of the Internet anywhere on the globe. The speed at which online technologies have reached 50% penetration of Canadian homes is truly remarkable: some five or six years compared to decades for the telephone, radio, TV and automobile! Of course, many of these online Canadians would be surprised to learn that the their pokey 56 kilobit-per-second desktop modem offers as much bandwidth as Canada's main Internet backbone did for the entire research community a mere 10 years ago!

With all this technical progress and its undoubted benefits, however, has come a more challenging legacy for Canadian citizens, scientists, business leaders, policy-makers and others. This legacy includes questions that continue to be debated in Canada and elsewhere – questions touching on how the Internet should be used, how it should be governed and how conflicting interests should be accommodated.

One dimension of this more challenging legacy has been the many efforts over the last half dozen years on the part of governments and public agencies like the CRTC to deal with the social and economic ramifications of the Internet's growth and development. These concerns have led to a range of public consultations as government has tried to come to grips with how the Internet is affecting its role. At the national level, the Information Highway Advisory Council (IHAC) started the process in 1994, to be followed by IHAC II in 1996, the Advisory Council on Health Infostructure in 1998, the Advisory Committee on On-Line Learning in 2000, and, most recently, the National Broadband Task Force, which started its deliberations in January, 2001.

Provincial governments have also been active, with numerous consultations and initiatives of their own. Most notable among the latter have been the federal and provincial/territorial collaborations in the areas of SchoolNet and the Community Access Program, both of which stand out on the world stage for their accomplishments.

Members of the original CA*net community have been involved in most of these deliberations and initiatives.

At root, however, the most challenging legacy of the builders and founders of the Internet is the question of its future evolution. The original Internauts were from the academic and research communities and they simply wanted to develop tools for those communities. Theirs was a "public interest" orientation based in large measure on the traditional "public good" nature of the research and education being produced by their home institutions. Along the way, and partly because of the robustness of the "rough consensus and running code" approach to protocol development that led to the eventual triumph of the TCP/IP suite over the ISO standard, they just happened to create the infrastructure that is now expected to underlie the future economy of the world, the so-called "networked economy". But this evolution is still in its early days, and the disruptive power of the Internet, and especially its "distributed" model of control, is still being digested by the more traditional business, commercial and political interests of the world. For their part, governments in Canada and the US seem intent on distancing themselves from any ongoing role in terms of control or regulation of the Internet.

One illustration of this continued tension over the Internet's future can be seen in the difficult birth of, and frequent skirmishes regarding, the Internet Corporation for Assigned Names and Numbers (ICANN), the body that the White House and US Dept of Commerce helped to create with a view to delegating governance over domain names and Internet numbers to a non-governmental body. Fair representation of non-commercial interests on the ICANN board and committees, as well as the implementation of dispute resolution mechanisms, especially over the use of domain names, are just two of the fronts on which the new body has had to do battle with public-interest advocacy groups of all stripes, and to try to balance commercial and non-commercial perspectives of the Internet's future.

Canada seems to have escaped the worst of these battles, with the protracted debate over how best to administer the ".ca" space having at last been settled. The Canadian Internet Registration Authority (CIRA) is now up and running and about to hold its first public elections. Again, members of the original CA*net community remain active in the operation and governance of CIRA.

The underlying concerns at issue in these debates, however, defy easy resolution. They include the right to individual privacy; the protection of digitized intellectual property; free speech online; fair and open access to networks by both service providers and consumers; and what has come to be known as the digital divide between those with online access and skills, and those without. While these issues have roots that go far beyond the Internet, most are also given fresh meaning in the tug-of-war over the Internet's future between the public interest on the one hand and commercial interests on the other. The founders and builders of CA*net could not possibly have foreseen the extent to which a part of their legacy would involve debates over privacy, free speech and open access that are likely to embroil governments, the courts, huge corporations, Internet governance bodies and other interested parties for years to come.

This glance at the controversial side of CA*net's legacy cannot diminish the very real achievements of those who conquered geography, limited financial resources, commercial ambivalence and widespread skepticism to create an organization that, through its offspring, is still transforming Canadian society. CANARIE and its provincial counterparts, CIRA, CAIP, the National Broadband Task Force, and the CA*net Institute are all, in their way, inheritors of CA*net's mantle.

For its part, the CA*net Institute, which now functions under the auspices of CANARIE, provides funding to projects related to the Internet and the original vision of CA*net – including all the work involved in researching, writing and publishing the particular piece of Canadian history you are holding in your hands.