

The History of the Internet in South Africa

How it began

Author: Mike Lawrie

1997

Acknowledgement

The author would like to thank a number of people who have kindly provided some information about the early days of networking in South Africa. These include Dave Wilson, Francois Jacot Guillarmod, Justin Jonas, Pat Terry, Victor Wilson and Sean Haffey.

Overview

This document sets out to describe how the Internet was started in South Africa. It is a ramble around some of the background events that caused this to happen. There might be little apparent logic, nor any good sequence, in this document - perhaps time will permit some tidying up to be done. So be it, but everything described in this text had an influence on the startup of the Internet in South Africa.

In brief, a group of three persons, viz Francois Jacot Guillarmod, Dave Wilson and the author, found a way to establish a sustainable email link to the Internet in 1988. A number of others helped with this exercise, particularly Randy Bush and Pat Terry. The link was established between Rhodes University, in a small city called Grahamstown, and the home of Randy Bush in Portland, Oregon. The email link went into production for general campus at Rhodes in February 1989, although email was flowing across at least two Fidonet links in 1988. The author was the Director, Computing Services at Rhodes University, and Jacot and Dave were Systems Programmers in the Computing Centre. Pat Terry was the head of the Computer Science department, and the author is not too sure of a correct description for Randy Bush, but we all knew that, while he had not the slightest obligation to help any of us, he was a person who gave willingly of his time and shared his experiences. At the time that he started to help us, he was a compiler writer for Oregon Software.

The email link used the Fidonet mailing system as a transport mechanism to exchange email between the Control Data Cyber computer at Rhodes University and a Fidonet gateway run by Randy Bush of Portland, Oregon, in the USA. The Fidonet system in the

USA had a gateway into the Internet, and to many people's amazement, including the group at Rhodes, the system worked and stayed working.

One thing led to another, and about a year later there was a uucp gateway in parallel to the Fidonet link. This was also a dialup system. The uucp gateway provided vastly superior facilities to Fidonet, because the standards used for email on the uucp transport protocol were in line with the Internet standards as described in RFC 822. At that time (circa 1989/1990), the TCP/IP protocols were in use internally at a number of South African universities, in particular at Rhodes University, the University of Cape Town (UCT) and the University of Natal, Durban (UND). In that era, use of a 9600 bps channel was obtained via the fledgling Uninet network to connect Rhodes and UCT, and an attempt was made to connect up a small internet between the LANs at these two institutions - this was in (?March 1990). Routers were to all intents and purposes unobtainable, so we tried to use Win/Route - a commercial product sold by Wollongong that was supposed to make a PC act like a router. This simply did not work. Having wasted good money on that product, and about to give up, we then tried the public domain PCRoute, and then traffic started to flow. Work on getting email to flow across this link took place in May 1990. The technical work at the UCT end was done by Chris Pinkham, support at a policy and strategic level at UCT was from Fred Goldstein.

With the success of that internet link between Rhodes and UCT, a second internet link was established to UND. The technical work at that end was done by Alan Barrett. Note the lower-case internet, this was not part of the world-wide Internet yet.

In due course, the uucp dialup link was replaced with a full Internet connection that operated across a leased line (to Randy Bush) at 9,600 bps. This was towards the end of 1991. For whatever reason, we used the KA9Q package to handle the routing at each end of the circuit - quite why we did not use PCRoute consistently, or KA9Q consistently, has not yet been explained to the author by Jacot or Dave or Randy. It matters not, the packets flowed.

When sanctions against South Africa were lifted by the USA in 1992, a pair of Cisco routers was used on the link to the USA. In due course, Ciscos replaced the PCRouters which at that stage were connecting about 15-odd Uninet sites together.

Now this is all pretty simple to tell in a summary such as this. But that does not tell the entire story. There were many things that had to happen in order to get things to work, and to keep them working. This is the story about those things. Very few of them were planned or even thought about, things had happened that contributed to the success, and the co-incidences are quite surprising.

The Design

It is just as well that there was no formal design process, and that the design of the system was not spelt out, because everyone would have laughed. This is what it looked like:-

- Start with a Cyber 60-bit word computer that has a native character set in upper-case, and where lower-case ascii characters have to be kludged.
- Write an internal mailing system in Cobol for the Cyber, catering for upper- and lower-case characters.
- Use the Cyber NJEF package to link this mailing system to host computers that can handle IBM's RSCS protocol.
- Use kermit to move mail spool files to hosts that cannot handle the RSCS protocols.
- Use the Cyber as a gateway between the RSCS-hosts and the non-RSCS-hosts.
- Use copper twisted pair cabling to link terminals to the this Cyber for distances of up to 2 Km in order to allow Rhodes staff to access the email.
- On the Cyber, extract mail destined for the USA and further afield, and use a kermit-like transfer to move this mail into a Fidonet system running on a scrounged 286 PC with a 20 MB disk. Filter off mail for South African hosts that, for political reasons, were not allowed to use the international link.
- Make the mail on the PC look as if someone had typed it directly into the Fidonet system. Merge this mail with any genuine Fidonet mail in preparation for transfer to the USA.
- Dial to just about the furthest point in the USA, to the home of someone who does this for fun and has no obligation to assist in any way whatsoever, and transfer the mail into his PC, and receive incoming mail from his PC.
- In the USA, distribute the mail into a UUCP dialup mailing network, and then into the Internet.
- Incoming mail from the Internet will follow the reverse path until it reaches the 286 PC at Rhodes University, when it will be intercepted, the addresses examined, and most of that mail will move via a kermit-like transfer into the Cyber.
- Use the Cyber as a mail switch to deliver the email via a variety of protocols (RSCS, kermit) to hosts at other universities in the country.

That, in broad outline, was what was implemented in the first gateway, and what ran for about 12 months or so as the only email link to the USA for the researchers and academics in South Africa. It went into full-scale production use at Rhodes in February 1989. It was a relatively simple technical step to change this to become a UUCP gateway, and then to become the Internet gateway, but it was a bit tricky to convince other universities and research computer centres in the country to agree that using the Internet protocols was a good idea.

Theorem: There is a lot that they don't teach in Computer Science.

Early Computing at Rhodes

Rhodes University was one of the first universities in South Africa to install a computer. This was an ICT 1301 computer, installed in November 1965. The author was the customer engineer for this computer, which in itself could be a longish story. One of the most important programs that this computer ever ran was one to calculate "Ray Tracing

for Initial Values as Follows", which was the work for Pat Terry's M.Sc degree in Physics. This program used to run overnight, and on more than one occasion Pat was found in a sleeping bag on the floor of the computer room - but that's another story. One of the "fun" things run on the computer was a computer dating program for the Arts and Science Ball in September 1967, where the idea was that the computer would give you a shortlist of possible partners to invite to the Ball. There were a number of other fun-but-useful things done, like a number of very useful extensions to the MAC compiler, in particular adding a TEA statement so that the running program could be dumped to the drum store while everyone went off to have tea. This was of particular use in that Pat's long-running program could be interrupted while short jobs were run, and then re-started again. The gist of this is that the author and Pat Terry built up a very strong relationship over this particular computer, and had this not happened the Fidonet link might not have survived the "C:\JUNK" incident of November 1989.

Towards the end of 1969, the 1301 computer was replaced with an ICL 1901A. In July 1971, the author was appointed to the post of Computer Manager at Rhodes, and a few years later an ICL 1902T was installed to replace the 1901A. The 1902T ran the Maximop operating system that had been developed at Queen Mary College, London. The staff at Rhodes, in particular the systems programmer Martin Urry, rapidly became the South African experts on Maximop. The first terminals that were used were teletypes running at 110 bps, and they were very popular in spite of the slow speed and the noise, because they allowed interactive computing. Indeed, they were so popular that in 1975 the card reader and card punches were decommissioned, and all computing was done via the network.

Moo

The author is convinced that the efforts that were put into implementing the MOO game helped considerably to further cement relationships that took strain during the development of the Fidonet gateway, so the story is given below.

During the days of the teletypes, Pat Terry had returned to Rhodes with a Ph.D from Cambridge ("the" Cambridge in the UK), and had encountered a program called MOO. This was a player-against-the-computer game, but there was a wider issue in that each player's score was on a public league. It gave great status to be at the top of the MOO league. An article in Software Practice and Experience (SPE 1(2), 201-204, 1971, author's name was given as "Aleph0") showed by means of Information Theory that if the player could use all of the game response information perfectly, the long-run average score would be about 4.3 (ie 4.3 turns on average to reach the end of a game). An alternative way to get to the top of the league was to crack the storage location of the league table and put your name there. More exciting was to put a name like "gremlin" at the top of the league so that the maintainers of the program were left guessing as to who gremlin was as well as how this had been done. So, the challenge was thrown open by Pat Terry (main programmer of MOO) and the author (main protector of the computer files) to all and sundry to crack MOO by any means whatsoever, but heaven help whoever was caught trying to crack any other part of the computer system.

Well, one student by the name of Sean Haffey went through a phase of being at the top of the MOO league for far too long with far too low a score. Pat Terry and the author spent considerable time together to find out how this had been achieved, because neither Sean Haffey nor anyone else could play so well so consistently. This cemented relationships between Pat Terry and the author, particularly when dealing with technical computer issues.

For the record, and perhaps Sean Haffey does not know this, in desperation we spent quite a bit of time scratching through the waste-paper bins of the terminal room looking for the abandoned teletype printouts that were from Sean Haffey's interactive sessions, and eventually found enough to give us a good lead as to how he was doing what he did, and thus we patched the security gap.

The game itself required the player to deduce a 4-digit number that the computer had generated at random for that particular game. The number could not start with a zero and had no repeated digits. The player typed in a 4-digit number (of whatever choice), and the computer responded with a count of the number of bulls (ie correct digits in the correct position) and the number of cows (correct digits but in the wrong position). This is a bit like the commercial game called Mastermind. One significant difference was that it had league table that was publicly readable, so anyone could see how well anyone else was doing.

The First of the Video Terminals

Teletypes gave way in due course to video terminals. ICL's video terminals were wickedly expensive, used synchronous communications, and were not suited to the academic environment at Rhodes. (maybe this is unkind to say so, but there really was not a lot of money to go around). So, we went for cheaper async terminals. Some low-cost async video terminals were purchased from Newbury Labs in the UK, and speeds moved up to 300 bps and to 1200 bps. There were also three Tektronix T4010 graphic terminals.

There was a bit of a problem in that neither the staff at Rhodes nor ICL knew how to connect these video terminals to the ICL computer. Eventually, we at Rhodes worked out that this could be done with two wires, viz the ones that carried the transmit and receive signals, provided that we could get a good local earth at each end of the link, and provided that certain control signals were strapped. For the technically minded, on the RS-232 25-pin connectors we would strap pin 4 to 5 and strap together pins 6, 8 and 20. Pin 7 went to the local earth at each end of the connection, and pins 2 and 3 were connected to pins 3 and 2 resp of the remote connector. By whatever weird accident, we found that this would work error-free at 1200 bps over some 2 Km of Telkom's copper local-leads, and the author (and in due course many others) used this to get interactive access from home. We put lightning suppressors between pins 2 and 7 and between 3 and 7, and put a 500 mA fuse in the circuit to protect the Telkom cable from blowing up during thunderstorms. Some years later, we used a 4-wire circuit driven by RS-422 line-

drivers that we developed at Rhodes, and this would work at speeds of up to 38.4 Kbps, on the shorter circuits, without any errors. These were still in production in 1997.

In due course (??year?? 1978??) the 1902T was replaced with an ICL 1904A. By then the network had expanded, and the faster computer was needed. Maximop was still the operating system in use, and we made use of ICL's George 2 and Queen Mary's College's Maxibatch.

The point about all of this is that in due course the technical computer staff at Rhodes were automatically provided with a connection to their homes, as was the author. Mind you, it was very difficult to get members of the Computer Steering Committee to agree that PCs should be provided as well, but in due course this happened. So, when the Cyber/Fido gateway was being developed and was being expanded, the computer staff who were doing the development could access these facilities from home, in the evenings. Given that time zones are such that the South African evenings are the normal working hours of most of the USA, we could communicate with our contacts in the USA at times that suited all of us (no job-type disruptions for us, and normal working hours for our friends in the USA). It was relatively easy to have several exchanges of email in an evening, and this boosted the productivity of the efforts that we put into developing this system.

There was a teletype circuit to the ICL computer at the University of Fort Hare, where Maximop was also in use, but this was used only to connect a teletype at one end to the ICL host at the other - there was no attempt to link the two computers directly.

The Cyber Computer

Rhodes University installed a Control Data Cyber model 825 computer in November 1981. This computer was an excellent number-cruncher, and it had a very flexible Network Processing Unit that handled the data communications. The operating system on the Cyber was called NOS, which stood for Network Operating System. The simplest way to put jobs into the computer and to control them was via a VT100 or similar terminal.

The Cyber's storage unit was a word of 60 bits. This was ideally suited for scientific applications because floating point operations were very fast indeed. Compile times for Fortran programs were incredibly fast as well, a moderately-sized program would compile in a matter of seconds. The Cyber's 60-bit word was treated as 10 6-bit characters whenever it was necessary to process character strings. Because the Cyber was designed for scientific work, the programming language was Fortran, which at that time required only upper-case characters, so a 6-bit character set was very efficient. However, the full ascii 7-bit character set could be handled by the NPU, which was hard-coded to treat "ordinary" 6-bit characters as their 7-bit ascii upper-case equivalent, but if the character was a circumflex (ie "^"), then the following character was converted to a lower-case ascii 7-bit character. That was for characters leaving the Cyber via the NPU, and there was an equivalent conversion when ascii characters were received via the NPU

and stored within the Cyber. A parity bit could be handled as well, and the full 8-bit character set could operate through the NPU in ascii or ebcdic coding. In all, the NPU was an extremely flexible and fast device, and could handle just about any communications protocol that was in use by any other mainframe manufacturer. It is significant that, unlike IBM with their SNA network and DEC with their DECNET network and other manufacturers with their own proprietary protocols etc, Control Data did not invent any proprietary communication protocols, but had packages for the Cyber's NPU that allowed it to handle such protocols.

So what, you well might ask. Well, for sure it made programming a mailing system on the Cyber a bit of a problem if you wanted to handle the full ascii character set. And that had to be done if one wanted to communicate with the rest of the world, because that was the lingua franca. We had quite a bit of work on our hands as a result, but all of this was completed long before we had even heard of Fidonet.

So much for the character set and protocols, there were other matters that had to be addressed. One of them was that the Cyber did not have a mailing system, so we had to write one if we wanted one. Much of the programming of the Rhodes mailing system on the Cyber was done in Cobol. Yes, that scientific oriented Cyber had an excellent Cobol compiler. Cobol was chosen as the programming language for the Rhodes mailing system because it was so easy to make use of the disk routines, particular the Direct Access routines. Handling ascii characters, though, was something else, and that gave us a headache. The main programming work on that mailing system was done by Hugh Murrell, in the early 1980's. Ian Dore and Eugene de Jager were involved in the changes that were necessary later on in order to interface with the Fidonet gateway.

The stand-alone mailing system had moderate use at Rhodes. Each registered user had a mailbox, some users made use of them sometimes. There was considerable email exchanged between the three of us in the Computing Centre, so if nothing else we could see the benefits of expanding this facility on to cover a much wider geographic region.

Why is any of this significant? Well, the full source code of the Cyber mailing system was at hand, as well as the total design, because it was developed as this author's design, and thus when it came to writing gateways to move mail in and out of the Cyber, it could be done relatively easily.

Attempt to Use Modems

There was a serious attempt to "legalise" the situation regarding the use of Telkom's copper cabling that ran around Grahamstown (1989?). It somehow seemed to be wrong that we were not using modems, so we made a very serious attempt to rectify the situation. An order was placed with Dimension Data for their LDM2000 limited distance modems that had been on demo at a Telkom conference in Cape Town. The order fortunately spelt out part of the specification that was on the glossy blurb, something along the lines of "These units shall operate at 9,600 bps or faster over a distance of 5 Km". Given that the Rhodes' homebrew devices ran with 100.00000% reliability at 19.2

Kbps from the Cyber to the author's house (about 2 Km of wire), this circuit was chosen as the test-bed for these modems. In short, nothing worked on these LDM2000 modems. Telkom technicians were called in to check the line (there was nothing wrong with it), the LDM2000 design engineer flew in from Johannesburg to attempt to fix the modems, all to no avail. Well, not quite, a Telkom technician did a "hot change" of a modem board in the modem rack (this is what the specs said could be done) and blew a 5 cm dirty black hole in one of the boards in the Cyber. The modems were returned to Dimension Data in July 1989. Well, a year or so later we eventually got our money back from the vendor, but we'd learnt a lesson. Don't use modems unless you have to, rather drive the wires with your own devices because they work.

Theorem: Vendors of products and users of products are quite distinct, and the one does not understand the needs of the other.

Corollary: If you buy a product for a particular purpose, it will not work to its specification.

Plato

Rhodes University ran the Plato computer based education system on the Cyber. Rather, on one of the Cybers, the one called RUPLA. Three of these computers stood in the computer room, one for "normal" computing of research and administration (RURES), one for Plato (RUPLA), and one that was acquired from JCI as a spare that could be cannibalised in case the sanctions against South Africa got worse before they got better.

A number of institutions connected to this Plato system, as far afield as Cape Town and Johannesburg. This was quite tricky, and had to be done by means of communication across lengthy (600 mile) analog circuits that were not very reliable. The lessons gained in keeping these communication links working were extremely valuable when it came to working across an analog circuit to the USA. Some of these links used a private (ie non-Telkom) X.25 PAD at each end in order to multiplex a number of Plato terminals on one circuit - again, valuable X.25 experience was gained.

But there was another benefit that arose from Plato. Professor Fred Hoffstetter from the University of Delaware visited Rhodes University (year??--mid '80s??). One thing led to another, and the author was allowed a signon on one of the Udel computers (??1988??). Access to the Udel computer was via the exorbitantly expensive and inefficient Telkom international X.25 service. The big plus was that Udel was on the Internet, and the author could thus send and receive email directly into the Internet. The value of this cannot be overstressed - when the Cyber/Fidonet gateway was being developed in 1988, it was relatively simple to check that email did or did not arrive correctly at an Internet site, and that it could be sent from an Internet site into the Cyber. It was also possible to download RFCs from the Internet via this link, albeit an expensive thing to do, so we could rapidly obtain the specifications that were needed. Getting these specifications electronically was very helpful indeed. The procedure to get them in paper form was horrendous. Firstly, non-existent funds had to be found in an impossibly-tight budget, then ordering from the

USA was a long-winded process requiring clearance in advance for a foreign exchange payment, which in itself required the vendor to provide a pro-forma invoice in advance, and then when/if the order was ever fulfilled by the USA vendor (remember, sanctions were in place, some vendors refused to supply goods) it would arrive in South Africa and incur a hefty import clearance charge and import duty, never mind the lengthy delays of this last process.

For the record, the author's signon at Udel was lawrie@vax.oit.udel.edu.

JNET

Well before the Foundation for Research Development (FRD) started the Uninet project in 1987, it was apparent to the group at Rhodes that it was essential to find a way to link the Rhodes Cyber to the network of IBM computers that linked the Universities of Potchefstroom, Witwatersrand and Pretoria together, and that also linked into the CSIR. That small network was connected at 9,600 bps on leased analog lines, and by all accounts there was no great use made of it. However, the network existed and was significant in that it connected different and autonomous institutions together. The technical work of maintaining this network was done by the networkers at Potch, so Potch was the obvious connection point for linking to the Rhodes computers. The cost of a leased line between Rhodes and Potch was simply unaffordable, so other means had to be found. In practical terms, the only option was to use a dialup connection, the only modems available worked at 2,400 bps. The communication protocol on this network was the very protocol that was used by Bitnet, viz IBM's RSCS, which used a synchronous communication protocol.

It was quite out of the question to attempt such a dialup link without first gaining some experience with the RSCS protocol. This link would have linked computers of different vendors and under separate autonomous control. The potential for chaos was simply far too high.

So, we started work at Rhodes to link a local computer to one of the Cybers by means of the RSCS protocol. This local computer was on our doorstep in the Physics department at Rhodes, in the form of a small VAX run by Justin Jonas.

Gateway Cyber/VAX

In the Physics department at Rhodes, there was a small Vax 11/730 computer. Some time in 1987 the author came across a package called JNET, which enabled the VAX to run the RSCS protocol. It interfaced well into the VAX operating system, and in particular into the VAX's email system. Thanks to the vision of the Control Data SA account manager for Rhodes, viz George Campbell, Rhodes had obtained the Network Job Entry Facility (NJE) for the Cyber, which enabled the Cyber to run the RSCS protocols. Funds (about USD900) were found by the Physics department to purchase the JNET package from Joiner Associates in the USA. The political credentials of Rhodes were provided in order to deal with the sanctions issues which were then rife in the USA. The datestamp on the

JNET tape, which was still at Rhodes in 1997, was 11 October 1987. By the 18th January 1988, good progress had been made with the connection between the Cyber and the VAX and we could make use of all of the facilities that it offered.

The linkup between the Cyber known as RURES and the VAX known as RUPHYS was established, and it worked very well. Modifications had to be made to the Rhodes mailing system that ran on the Cyber in order for mail to use the RSCS protocols as a transport mechanism. This was not too difficult for the Computing Centre staff to do. The exercise taught us a great deal about how to set about building a network of computers that ran different operating systems, and were controlled by different departments. The link operated at 9,600 bps on Microcom modems that had been obtained cheaply from a commercial company (name is forgotten) that had recently changed from analog to digital circuits and thus no longer needed analog modems.

In due course, when the Cyber/Fidonet gateway was working, it was a simple matter to include the RUPHYS host in the tables that routed email to the correct hosts. By an extension, carried out by Justin Jonas of the Rhodes Physics department, a microvax, called RUCHEM, in the Chemistry department could also be included in this network. A means was found to get email to a host, called RUBIS, in the Business Information Systems department - this being motivated by Malcolm Sainsbury. Simple? Well, not quite at the time, but simple in retrospect. One has to bear in mind that this was the only "Bitnet" style network that ever existed in South Africa, and we had to find our own way ourselves in order to make it work.

South Africa's only BITNET-Style Network

In due course, a number of computers in South Africa were connected either directly or indirectly via the Cyber computer at Rhodes. Without setting out to do it, the group at Rhodes had developed a network that operated in a manner virtually identically to the Bitnet network. The network used IBM's RSCS protocol to exchange email, it was possible to do remote job entry (RJE) from one host to any other host on the network, printing could be shared, files could be transferred, etc.

The earliest archived message that the author can find that flowed across this network is a message dated 8 March 1988 from Vic Shaw logged into the IBM computer at Potch (VSIRSHAW@PUKVM1). However, it is clear from a message dated 12 February 1988 to Vic Shaw from the author, that at that stage email could be exchanged between a computer at the CSIR (CSIRVM) and any of three computers at Rhodes (RURES, RUPLA, RUPHYS).

This section gives some information about that network, which proved to be a stepping stone on the way to the full Internet connection

Link to IBM at Potch

Having cracked enough of the local networking problems at Rhodes so that email could be exchanged between different hosts that were controlled by different authorities, and having got to grips with RSCS protocol that was used by the Bitnet network, it was now time to attempt the next big step. This was to link the Rhodes Cyber to the IBM computer at Potchefstroom University, and thence to Wits, Pretoria and the CSIR in such a way that all of the Rhodes hosts could use this link. The only option open was via dialup working, using synchronous modems, because the cost of a leased line was prohibitive. The IT Director at Potch, viz Philip Welman, was extremely enthusiastic about such a linkup, and we had his full support and the keen interest of his networking staff (Jan van Rooy, Henry van der Walt and Sakkie Larney).

For the uninitiated, Potch and Rhodes were at the opposite ends of the political spectrum in the country. Rhodes had a very liberal left-wing anti-government reputation, and Potch (to some people, anyway) had a reputation of being somewhat verkramp (loosely, that means "less liberal than others"). There was also a language barrier, which might seem strange to persons not brought up in South Africa, in that Rhodes was an English-language university, and Potch was an Afrikaans-language university. There was also the not-so-minor matter that the name Rhodes, which was taken from Cecil John Rhodes, a British Colonialist of not such good repute, had an unfortunate association with the Jameson raid of the 1890's, and this raid was an attempt to instigate an uprising in the old Boer republic so that Britain could seize the then new and lucrative gold mines of the Witwatersrand. Potchefstroom's forefathers were very much part of that republic. Well, the bottom line is that none of these differences seemed to matter to anyone who was trying to get the network to operate. There is a lesson there somewhere.

A really complicating factor was that Potch and Rhodes were separated by about 1,000 Km (600 miles). The link did not work at first (or second or third or Nth) attempt, and it took several visits by staff from both sites to make some kind of progress.

In the end, the dialup worked. This was at only 2,400 bps, but that did not matter too much. We could exchange email. So what? Well, some of our friends also said "so what, if we need to communicate with someone at Potch we can phone or fax them". There was no quick reply to this, but when these sceptics returned from visits to the USA, they changed their criticism to support.

For the record, the name of the IBM host at Potch was PUKVM1.

That Cable <blush>

In the process of trying to get the dialup link working to Potch, we learnt a very valuable lesson.

Some of us had been into the Potch computer room, and we saw little else but "true blue" IBM equipment. Even the modems and even the cables to the modems were made by IBM. Coming from the thin edge of the net at Rhodes, we inevitably made our own

cables, and often tweaked them to minimize the number of strands in the cable, thereby reducing the cost. So, we were somewhat bemused by what Potch was doing.

The cable that we used at Rhodes to connect the modem to the Network Processing Unit (NPU) on the Cyber was a "proper" cable made by some cable-maker, and was not the usual home-made cable that we used at Rhodes. After an extended number of unsuccessful attempts to get the modems to lock onto each other and produce a carrier detect signal, and having eliminated all possibilities, including that "proper" cable, of the problem being at the Rhodes end of the link, we were totally convinced that the problem lay at the Potch end. To cut what would be a long story short, we eventually found out that the cable at our end was quite unsuitable for a synchronous modem, and when we put a proper "proper" cable then everything worked. Whether or not the folk at Potch ever discovered our mistake we know not, because they were also making changes at their end of the link.

This was a very valuable lesson to us, viz no matter how certain you are that the problem with a communication link is not at your end, for sure it always is. Beyond question, this lesson helped no end when we had to deal with a network reaching into the USA.

Theorem: The problem is always at your end of the connection.

Corollary: You will always believe with absolute certainty that the problem is at the other end of the connection.

Access to Wits/Pretoria/CSIR

One spin-off of the dialup link to Potch was that there was now communication between staff at Rhodes and the network in the northern region of the country. The author made use of this extensively in the work that ensued as part of the CUP networking committee, and at the same time he exchanged an enormous amount of email with Vic Shaw, then in charge of the CSIR's computer division (and later manager of Uninet). This led to Rhodes acting as a node on the CSIR's national network, thereby gaining more networking experience. More importantly, it built up an ethos of email communication between Vic Shaw and the author, which continued long after Uninet became a formal entity in 1987.

Jessie Gwynne

The person at Joiner Associates who dealt with the sale of JNET to Rhodes was a charming lady called Jessie Gwynne. The author exchanged a large amount of semi-private email with her, and in due course met her on 15th October 1990 at the Educom conference in Atlanta, and visited her and her husband (Bob) at their home in Madison, Wisconsin. This was well before the TCP/IP packets were flowing into South Africa. It brought home to the author, who at that stage did not really need any more convincing, of the power of email communication and the friendships that it generated at the time.

Fidonet Issues

The Fidonet system played a major role in the first email gateway at Rhodes, so it is worth spending a bit of time discussing issues surrounding it.

As far as the author knows, Fidonet was developed by Tom Jennings out of a PC-based bulletin board system (BBS) system. This was in the 1980 era. By 1985 there were some 400 nodes in operation around the world, but mostly in the USA. BBSes were popular with computer hobbyists. In essence, a registered user could load software to and from such a system via a dial-up modem, and could leave messages for the system operator (sysop) or for other users of the same BBS system. Tom Jennings extended this concept on Fidonet, such that messages could be left for users of other BBSes, and thus a system of interconnections between BBS systems took place. Given the cost structure of phone calls in the USA, where calls within most cities were free, Fidonet was built around an hierarchical system where PCs in one city would communicate directly with a hub PC via a local (free) phone call, and the hub PCs would communicate with other hub PCs in other cities using efficient aggregation and compression techniques. By carefully planning which hubs contacted which other hubs, the cost of the long-distance calls could be kept as low as convenient, and messages flowed via a store-and-forward mechanism.

A very important aspect of Fidonet is that it allowed some bright and enthusiastic persons to get involved in networking. The efforts that these persons put into this exercise were enormous, and it was done with full self-motivation. The costs were met out of their own pockets, so there was every reason to get the last ounce out of the computers and out of the network. This led to implementing all kinds of cost-saving schemes. Eg, all files of any significant size that were exchanged across the communication links were compressed by the most efficient compression algorithms available, and there were restart points during transmissions that catered for spurious disconnections. A scheduling system was in place that allowed for cheaper time-of-day telephone rates to be used. And there was more and more of this type of thing. Not to be forgotten is the amount of help that was readily forthcoming from Fidonetters world-wide. It is very difficult to overstate the amount of help that we received in this regard from Randy Bush, who put in an incredible amount of effort to help us.

Fidonet does not scale very well. Every Fido node needed to be registered in a central file, and this central file had to be installed on every Fidonet system in the world in order to ensure reachability. Maintaining this file consumed significant effort, and consumed precious bandwidth. But beggars could not be choosers, we were very grateful for what Fidonet offered us.

Fidonet was developed at a time when the Internet was to all intents and purposes limited to the research and academic community in the USA. Before the end of the 1980's, it was apparent to all that it was highly desirable that the various email systems of the world should interconnect, by means of email gateways. Such an interconnection would, for example, allow users of the UUCP and Bitnet networks to exchange email with one another and with users of the Internet. By some means or other, and no later than by 1988, the Fidonet system could exchange email with these systems as well. Someone had developed an "add-on" package called UFGATE, which ran on the Fidonet PC and acted

as a gateway to/from a host running the UUCP communication protocols. There was a large UUCP community in the USA in that era, and the UUCP network had several gateways into the Internet itself. For some reason or other, UFGATE did not appear to catch on in South Africa, and the group at Rhodes were unable to get it to work successfully.

The design of the Fidonet email addressing system, and the method of holding addressing information, was not directly compatible with the standards for email exchange between two computers as spelt out in the Internet standard known as RFC 822. Just for one thing, the Fidonet system did not support the concept of cc:ing a message to others, and this was added as a kludge (as were other features, like addressing email to someone on the Internet).

In itself, the vagaries of Fidonet were not a major problem for the team at Rhodes. For one thing, the Rhodes mailing system on the Cyber had a number of shortcomings as well. However, the standards of the Rhodes mailing system could readily be changed, those of Fidonet and RFC 822 had to be adhered to. The author drew up a specification for moving the email between the Cyber and the Fidonet system at Rhodes. This Fidonet system was called "Settler City", after a nickname for Grahamstown. (A large number of British settlers had moved to the area in 1820). The author's specification was in a document named MAIL000. Jacot and Dave did the implementation, with Jacot handling what had to be done to the Cyber in order to move email into and out of the mailer there, and Dave doing likewise on the Fidonet system. Fortunately, the experience that we all had in making the Cyber exchange email with the Bitnet-style JNET package that we purchased for the VAX in the Physics department, and in communicating with the IBM host at Potchefstroom University, stood us in very good stead.

Fidonet as a Carrier

It must be understood that we used Fidonet merely as a carrier for our email. Very few users at Rhodes used Fidonet directly. Typically, email was created by users of the Cyber (or the VAX in the Physics department), and read by users on those systems. The Cyber recognized that the recipient address in the email was not for local delivery, and so it put it in a spool file that was subsequently transferred to the Settler City Fidonet system by means of a file transfer. Similarly, email arriving into Settler City would be examined to extract email that was destined to a Rhodes computer (and eventually to any of the computers at Uninet sites), and it would be re-formatted and stored in a spool file for transfer to the Cyber.

Apart from email, there were some usenet newsgroups that were somehow brought in, as well as some Fidonet "echo" groups. There was no transfer of newsgroups between the Cyber and Fidonet, any distribution that took place had to be done to other Fidonet systems directly. This in itself was not too limiting, email was a very strong means of communication, and there were many things that could be done from a "mail-only" system. Some listserv hosts could distribute newsgroups, and there were ftpmail servers

as well. We never really understood why and how the sites that ran these things did so for free, but we were extremely grateful that they did so.

Addresses within the Fidonet system, when seen from the Internet, consisted of a user-id built up of `firstname.surname@{fidonet-system}`. On the Cyber, the user-id's that we used at Rhodes consisted of four characters, being a hangover from the Maximop era. Eg Jacot's user-id was `ccfj`. The Cyber that was used by researchers at Rhodes was named RURES. The kludge used within the Fidonet system to give an email address to Jacot was then `ccfj.rures@{Settler-City}`. The `{Settler-City}` part was quite convoluted but met the standards of Fidonet.

For the record, Settler City was initially assigned the Internet-style address of `f19.n490.z2.fidonet.org`. Soon after June 1989, at of the whim of some Fidonetter, the address changed to `f4.n494.z5.fidonet.org`. The disruptive factor of this change was enormous. This taught us a good lesson, viz never rename a host to have a different Internet address unless you really have to do so, and even then don't do it.

For the record, within the Fidonet system the addresses were of the form "Firstname Surname of <zone>:<net>/<node>.<point>", for example "Victor Wilson of 5:7106/22.3". Not all addresses made use of the ".<point>" option.

The attempts at Rhodes to link to the Internet using Fidonet started in about July/August 1988. The contact with Randy Bush was from September of that year. Mail was flowing to and from Randy Bush by about 10 September 1988, but the process to set up a more formal system started with a letter to Randy Bush dated 14 September 1988. Email could be exchanged between Rhodes and the Internet by 16th December 1988. The earliest email that the author has in his archives that exchanged between RURES and the Internet via the Fidonet gateway is dated 15 February 1989.

It was a significant day in our lives when the Fidonet gateway was still operating in production three weeks after it was opened for general campus use at Rhodes, in February 1989. This was because we had heard that quite a few months earlier, UCT had operated a UUCP connection to UUNET in the USA, and that this link had been closed down for political reasons after three weeks. We really could not believe that we had done better than this - we had broken through the sanctions barrier on a zero budget operation without having to resort to any cloak and dagger exercises. We had also beaten the South African government of the day, which was desperately trying to control every last form of communication in and out of the country.

Linking to Botswana

By some strange quirk of fate, somewhere in Portland, Oregon was the son of a Fidonetter called John Case of Botswana. Prior to Rhodes linking Fidonet directly to Randy's system, the Fidonet routing from the USA took place via a very long-winded path that included someone in Amsterdam. It made sense, once Rhodes had a direct link with Randy, to act as the forwarding system between John Case and his son. Email

records show that steps to do this were initiated at the end of January 1989. So, from the time that there was general international email on the Rhodes campus, Rhodes was acting as a gateway for email to another African country. This trend was to continue into the future, and Rhodes built up a considerable reputation in this regard.

Pat Terry and Modula-2

A major boost was given to the Fidonet linkup by Pat Terry. Pat, a professor in the Computer Science department at Rhodes, was heavily involved in drawing up the ISO standards for the programming language Modula-2. He attended one of the ISO committee meetings in August 1988, and realised how out of touch he was with the rest of the committee, who had been doing a great deal of communicating with each other by email. Randy Bush was a member of the same committee, and encouraged Pat to set up a Fidonet system himself so that he could participate more actively in the work of the committee. Pat returned from that August 1988 trip with a great deal of enthusiasm now that he could see a real use for email. He brought with him the Fidonet version known as BinkleyTerm, and a promise of a gateway link via Randy. The team in the Computer Centre then switched from the flavour of Fidonet that was currently under test, viz Fido 11, and whatever convoluted routing (via Amsterdam) that was being tested, and started working with Randy Bush.

One cannot overstate the help that we got from Randy, it was quite incredible. This was not only when the Fidonet link was being developed, it continued for very many years afterwards, through the stages of the UUCP link and the first TCP/IP link. Pat Terry's comment puts it very precisely, viz "It was an unbelievable coincidence really that of all the possible Fidonetters any of us might have met, we just happened to get a personal contact with Randy who acted as zonegate from the US to several other zones."

Zonegate Issues

It became apparent by March 1989 that the Cyber/Fidonet gateway worked in production, could handle a reasonable load of international traffic, and could be extended to handle the email of the various Uninet sites subject only to political constraints. Under such conditions, the Uninet project of the Foundation for Research Development (FRD) in South Africa was willing to fund the costs, in particular the cost of the phone calls. It was quite clear that the volume of Fidonet traffic that would flow in and out of Rhodes/Portland connection would exceed the previous Fidonet/South African traffic by many orders of magnitude. So, an offer was made to the Fidonet community for the gateway at Rhodes to become the main gateway for all Fidonet traffic in and out of Southern Africa. Thus Settler City was offered to act as the zonegate for the region, and the costs of the international traffic would not be passed on to South African Fidonetters.

For some strange reason, the Fidonet community in South Africa viewed this with considerable alarm. Clearly, something had been misunderstood, because tests had shown that the route via Amsterdam was very unreliable, apart from having to be paid by

members of the Fidonet community. After a great deal of emotional email discussion, things were sorted out.

In due course, the distribution of Fidonet into Southern Africa took place via Settler City. One of the longest lasting and significant distributions was to the University of Zambia in Lusaka. This started up when the politics of South Africa started to change for the better and sanctions were crumbling - a meeting between Mark Bennett, who was an expatriate from the UK doing sterling networking at the University in Zambia, and the author, took place in Harare in November 1991. That meeting led to the establishing of this link, which ran for at least three years.

The growth of email volumes was very difficult to cope with. The Uninet project of the FRD paid for the cost of the phone calls, and provided two 386 PCs (then considered to be cutting edge), one for Fidonet and the other for development work using SCO Xenix. In the last few months of the life of Fidonet and UUCP as the main carriers of email (ie shortly before there was a leased line running TCP/IP protocols to the Internet in November 1991), the cost of the telephone calls were at least double the cost of a leased line. Telkom initially refused to install this leased line on the grounds that there would be "third party" traffic flowing across it, and some bureaucrat inside Telkom insisted that the "common interest group" recognition did not apply to the international circuit. That the research and academic community had achieved a major breakthrough to the benefit of the country meant absolutely nothing to Telkom. Somehow or other Vic Shaw sorted that one out.

Once the leased line was in place, and had connected Uninet to the Internet, there were the diehards that would not give up Fidonet. There was also an active Fidonet community in the country, and Uninet was not going to forget the favours that it had received. Somehow between Dave Wilson and Randy Bush (others??), a means was found to reverse the original transport technique and to make the Fidonet traffic use the Internet as a transport mechanism. In essence, Fidonet traffic at the one end of the Internet circuit was gathered into a file with suitable headers and moved across the Atlantic into the Fidonet PC at the other end, and then was unpacked and proceeded within Africa as Fidonet traffic. As far as the author is aware, this is still in operation (in 1997). It was this type of cooperation and willingness to help each other that made things work, and we all felt that it would be a very sad day were this spirit to be killed off for some reason.

Fidocon '90

Given the importance of Fidonet to Rhodes University, and the interest in Fidonet among its adherents in the country, we organised a Fidonet conference at Rhodes on 8th September 1990. Guest star was Randy Bush, whose trip was sponsored by the FRD. A number of persons who were to play a role in the Internet in South Africa attended this conference, eg Victor Wilson of Telkom. The author has not been able to track down the list of delegates to this conference.

The C:\JUNK Directory

By late 1989, the growth in traffic finally caught up with us, and the flow of traffic through the Fidonet system underwent a reorganisation. For whatever reason, the redesign was such that the DOS .BAT files did very little checking before taking action. In particular, there was a directory named C:\JUNK. Hidden in the depths of the gateway process, there was a .BAT file that had the command CD C:\JUNK followed by a DEL *.*.

On the evening of 7th November 1989, the author was trying to sort out some major problem with the gateway, which was jammed. In spite of the recent reorganisation, it was running out of disk space. It is not clear why Dave Wilson was not involved in this exercise, but it matters not, the author removed the C:\JUNK directory which was, by assumption, full of junk. In due course, the CD C:\JUNK command did not change directory, and so the entire contents of another directory was wiped out. This was a key directory, and so the system was unusable.

Pat Terry somehow got very involved in recovering the system. Fortunately, the disk had recently been defragmented, and most of the files could be recovered by deft use of an UNERASE command. But it was tough going, at 3am, with Pat about to leave town for a few days within the next few hours.

The lesson was clear. More people had to understand what was happening on the gateway, and it had to be ruggedised. The author prepared a document MAIL017 that tried to do this very thing. This was on 21 November 1989. This helped, but at no stage could the Fidonet gateway be considered to be totally reliable.

KA9Q on the USA Link

The first leased line to the USA, between the computer room at Rhodes University and the home of Randy Bush in Portland, Oregon, was driven by Penril modems at 14.4 Kbps and had a 386 PC at each end. The software that performed the routing functions was Phil Karn's KA9Q program. This was in November 1991 (?? to be checked??).

Long response times

One of the "features" of the KA9Q routing software was that it appeared to be reluctant to drop packages when there was congestion on the circuit. Whether or not this was desirable is not discussed here, but for sure when there was congestion, the ping times between the two routers would rise to anything up to 300 seconds. This was way beyond the timeout limits of the TCP/IP protocols, and so the hosts that were connected across this link would themselves cause a timeout and a re-transmission. Given that the routers were still doing their best to transmit the original packets, the congestion simply worsened. The effect was that the routers at each end would hang, and the only way out was to phone to the USA and ask for a reset. One can understand that people like Joan Bush (Randy's charming wife) did not like to be woken up at, say, 3 am to be asked to reset the router. Besides, Randy had a number of such routers, and it was all too easy to reset the wrong one. Dave Wilson designed, and Brian Kemp (of Rhodes's Electronic

Services divisions) built, a circuit that monitored the RS-232 interface between the router and the modem, and if it went idle for too long, it would do an automatic reset on the 386 chip, and the router would then restart itself. This was installed at each end of the link, and helped a great deal.

The long response times caused a large number of problems that would not occur on a network that had a fast response time. Considerable time was spent trying to isolate these problems. One such problem that was eventually isolated, but could not be fixed permanently, was that domain storms would readily occur. One technical example of this was the case of the University of the Orange Free State, domain uovs.ac.za. This had its primary nameserver within uovs.ac.za, and secondaries at hippo.ru.ac.za, daisy.ee.und.ac.za, bertha.ee.und.ac.za and rain.psg.com. One weekend, the uovs.ac.za computer was unreachable for some reason or other. Using FTP Software's LanWatch program to monitor the ethernet traffic, we could see a DNS query arrive from the USA that had a request for resolving a uovs.ac.za hostname. The congestion on the USA circuit was such that even had the uovs.ac.za nameserver host been able to answer the query, the host that originated the request would already have sent out queries to daisy, bertha and hippo. Neither of these three hosts had information on the uovs.ac.za domain (for some reason or other at that particular time, the primary nameserver was unreachable and the DNS caches had timed out), so all three of these secondary nameserver hosts would operate in recursive mode and try to get the request resolved, which could only be done by querying each other and by querying rain. Now rain.psg.com was in the USA (in Randy's house), and rain itself could not do the resolving either so it would fire off recursive queries to the uovs.ac.za host (which was unreachable), and to daisy, bertha and hippo, which in turn would.... and so on and so on. Under those conditions the link rapidly degenerated into a state of carrying no productive traffic whatsoever. Once this situation had been identified, Jacot found a way to kludge a resolving of the uovs.ac.za queries on hippo, and things got back to normal. However, it took the entire weekend to track this particular problem down, but we could not find a way to prevent this from occurring again at some point in the future.

University Networking

This section endeavours to give some background information about the developments of university networking in the 1980's. This is necessarily incomplete, nothing was centralised, there was no forum for exchanging information and ideas, so no one had the overall picture. That in itself describes pretty accurately the mess that things were in.

This was still the era of the mainframe in South Africa. IBM PCs were around, Apple II computers were around, these were to be found mainly in student labs. Any networking that was to be done would have to be done via the mainframes.

There was a variety of computer manufacturers that sold computers to the universities in the 1980s. Names like IBM, ICL, Burroughs, Sperry/Univac, DEC/VAX, Hewlett Packard, Control Data, Data General, Prime and others all had some of the business. There were geographic regions of influence. In the Transvaal, IBM seemed to hold sway

(hence the SNA network linking Potch/Wits/Pretoria/CSIR. In the Western Cape, there were VAXES linked between the Universities of Cape Town (UCT) and Stellenbosch (SUN). Hewlett Packard seemed to have the business in the Natal area, although there were VAXES as well. In the Eastern Cape, Rhodes had a Control Data Cyber, the University of Port Elizabeth had a Burroughs.

Even after the Uninet project had been running for some years, this difference in the mainframes (which implied differences in networking protocols) still dominated. Eg, at a Uninet conference in October 1988, the author presented the only paper on the TCP/IP protocols and the benefits thereof. A much later Uninet Board meeting took the author's advice that "Uninet should get to the standards of the Internet and follow those standards as they evolve". This acceptance put to an end the protocol wars that were smouldering. That differences of opinion should exist was hardly surprising. In the early 1990's there were protocol wars in the USA about the benefits of TCP/IP vs OSI.

To the best of the author's knowledge, the wide area university and research networks that existed in South Africa in the 1980s were these. Corrections and extensions to this list will be very welcome.

- The CSIR had a network that linked IBM and Cyber computers in Pretoria to research units across the country. This used Case multi-protocol multiplexors. It did not link the IBM to the Cyber, though.
- Potch/Wits/Pretoria/CSIR, and IBM SNA network using 9,600 bps leased lines.
- UCT/SUN, a Decnet network linking VAXES
- University of Natal, linking HP computers at the Durban and Pietermaritzburg campuses.
- ...other.....

CUP Computing Subcommittee

The Committee of University Principals (CUP) was a body that was formed (by statute) to oversee the interests of universities in South Africa. In the 1970's, growing out of efforts by the university IT staff for a forum to discuss common issues, the CUP formally established a computer subcommittee. This subcommittee, at a meeting at Rhodes University on 6th June 1986, formed a networking subcommittee. This was chaired by Philip Welman of Potchefstroom University, the other two members being Willie Fryer of the University of Pretoria and the author. This committee met on several occasions, but with the best will in the world could not produce ideas for solving the "protocol war" problems, nor produce acceptable proposals for installing and running a network because of the high costs of Telkom circuits. With the best will in the world, the suggestions were to start with the RSCS protocol in a BITNET-style interconnection, and then migrate to the OSI standards. Not a word of TCP/IP.

Uninet Starts Up

In 1987, the President of the Foundation for Research Development, Dr Rein Arndt, took an active interest in the formation of a computer network. This was done in consultation with interested parties, including the CUP's computer committee (at a meeting at RAU, (when?? 1988??)), and with the networking subcommittee. In due course, seed money was made available by the FRD to purchase networking equipment and to hire circuits. One of the motivations for a network was that South African scientists had a need for access to a supercomputer, and without a network there was no way to make such a computer accessible to the scientists. As an aside, it is interesting to note that some ten years later there is still no supercomputer for these scientists.

Vic Shaw formerly of the CSIR, was appointed to the FRD as manager of Uninet in October 1987 (?) to co-ordinate the efforts of the computer staff at those universities that took an interest in establishing a national research network. A study into networking needs was conducted, and surprisingly while it mentioned the TCP/IP protocol, that protocol was not identified as even being the way forward. This may well have been a byproduct of the sanctions that had isolated South Africa from the mainstream of scientific developments elsewhere in the world.

One of the problems facing Uninet was that of Telkom regulations. The "third party traffic" rule was notorious, and had been misused time and again to close down such things as PC based BBSes. The rule applied only to leased lines, not dialup connections, and in a nutshell prohibited traffic from organisation A from being transferred across a leased circuit that linked organisations B and C. Clearly, Uninet needed to do this very thing when it linked up the various research and academic institutions of the country. Vic Shaw applied to Telkom for a concession under an escape clause about "common interest group". This application was made on 7th January 1989, and permission was granted in due course.

The networking equipment that was purchased for Uninet was Case multi-protocol multiplexors, which had been used very successfully on the CSIR's network. The circuits that linked key hub sites of Uninet operated at 64 Kbps on Telkom's Diginet network. Uninet was granted special permission by Telkom to operate this network under the "common interest group" escape clause of Telkom's "no third party traffic on leased lines" regulations.

All of this took time to get into place. (When were these circuits operating? They were in place by Feb 1989 when the Rhodes Fidonet gateway started working).

The muxes of Uninet provided virtual circuits on top of the physical ones. Any port on one mux could be allocated synchronous or asynchronous bandwidth to any port on any of the other muxes. Muxes were located at all of the sites that wished to participate in the Uninet project. These sites included UCT, Rhodes, Stellenbosch, Potch, Computer Centre for Water Research, University of Natal Durban, Wits and a few others (which??).

It was now possible for the VAXES at UCT and Stellenbosch to communicate with, say, the VAX in the Physics department at the University of Natal Durban by means of a

Decnet channel, and the IBM at the CCWR to join the network of IBMs at Potch, Wits, Pretoria and the CSIR. No one had a ready solution for the problem of getting the VAXES to communicate directly with the IBMs, for example, quite why the JNET package was not run on the VAXES defies any understanding on the part of the author.

The only commonality within the early days of Uninet was that the Cyber at Rhodes had a mechanism to exchange email with just about any of the computer brands that existed on Uninet, and it soon became a key element in what was to follow.

Six times through Wits, or bust

There was a big difference between physical and virtual routes of the dataflows across the early Uninet network.

By 1989, Rhodes was acting as a protocol converter gateway for email of Uninet. The Cyber at Rhodes could handle the RSCS protocol, and thus could exchange email with the IBM hosts of Uninet, and at the same time it spooled email into a file to be collected by a program like kermit for hosts like the Hewlett Packards and VAXES in Natal and the Western Cape. All of these connections, including the kermit ones, ran across virtual circuits of Uninet.

It so happened that the route taken by email between the HP on one floor of the computer building at the University of Natal Pietermaritzburg and the IBM mainframe in the basement of that building was very convoluted. Looking at the virtual circuits, the HP communicated "directly" with the Cyber using kermit protocols, the Cyber accepted this email and converted any that was addressed to the IBM host into the RSCS protocols, and sent it "directly" to the IBM at Pietermaritzburg. Straightforward, until one looks at the physical circuits. From Pietermaritzburg, the 64 Kbps diginet circuit went to Wits. Potch was on a stub link to Wits. Rhodes was on a stub link to Wits. So, the kermit virtual circuit between the HP at Pietermaritzburg went physically via Wits. The Cyber then converted it to RSCS protocols and sent it to Potch (via the physical circuit to Wits). Potch sent it on its physical circuit back to Wits by means of the SNA protocol, and Wits sent it to Pietermaritzburg. Well, it was something like that, but a bit more complicated because among other things there was more than one mux at Wits.. All told, we counted six transits of the email through Wits in order for email to flow from one floor to another in the same building at Pietermaritzburg.

Search for an Internet Gateway

In spite of convoluted physical routes, Uninet worked well. Email exchange between hosts run by separate autonomous institutions was taking place. This was indeed a very big step, but the bigger need was for email to flow between Uninet and the Internet. The efforts to get an international link for Uninet started in earnest in 1989. The problem was political far more than technical, but that is not to say that the technical problems were trivial.

Some time before 1989, UCT had established a UUCP connection to UUNET in the USA. This ran across the international X.25 service of Telkom. Unfortunately, after about three weeks the folk at UUNET pulled the plug on the grounds of South Africa's apartheid policy. They felt that they could not be seen to be co-operating in any way with a South African organization, no matter how liberal that organization might be. A number of individuals, including the author, had private and personal signons on hosts in the USA and Europe, granted out of a spirit of generosity. A researcher at the University of Natal Durban, Dave Levy, had connected the VAX there to the NASA SPAN network on this basis, and it was operating in 1989.

Letter from Bitnet

At about the time when JNET was being brought up to speed to link the Cybers in the Computing Centre to the VAX in Physics, the author wrote to the Bitnet organisation and requested a connection. A simple request, but one has to bear in mind that sanctions against the country were in full force. Even though it was not against the USA's Comprehensive Apartheid Act to deal with South African academic institutions, it would have to take a very brave institution in the USA to deal with any South African university. Organisations like SPSS Inc simply "pulled the plug" and refused to renew the license for a liberal anti-apartheid university like Rhodes to run their software.

That said, an Elizabeth Kilkoynne of Bitnet wrote back saying that we could connect to Bitnet. The author could not believe his luck, but could not afford the required 9,600 bps leased line to the USA. The matter was taken to the networking group (?? Cannot recall whether this was the then Uninet Board or the CUP networking subcommittee, but Vic Shaw was present), who in effect laughed this out of court on the grounds that this was permission granted at a "technical level" and still needed the approval of the Bitnet governing body. There were too many people in that committee who believed that the governing body would reject the matter. The author, much to his disgust, went with the flow, and vowed never again to listen to the advice of a committee of "experts".

Theorem: Any expert can prove that something cannot be done.

Corollary: If you wait for experts to agree on something, you will wait a long time.

It is of much relevance that the Fidonet design and initiative was never put to any committee. The three of us in the Rhodes Computing Centre "went it alone", without budget, with scrounged hardware, but with a firm belief that we could do it. The same applied to the links into Africa.

Letter from Sol Lederman of SRI

Having had the Bitnet acceptance rejected by local "experts", the author had nothing to lose by writing directly to the organisation that at the time seemed to be the body that had any kind of control over the Internet, viz SRI International in Menlo Park, California, and asked various questions about connecting to the Internet. This was on 6 May 1988.

A reply was received, dated 28th July 1988, from a Sol Lederman of SRI. In essence, a connection to the Internet was unlikely to be approved, but he suggested that we investigate four other networks. These were

- TWICS (Two Way Information Communication System) serving mainly Japan but with a link to South Africa;
- FidoNet, with a reference given to Dave Pedler based in Johannesburg;
- Telenet, a commercial Public Data Network linked into South Africa's SAPONET;
- Tymnet, a commercial Public Data Network.

Of all the things that gave inspiration to the author, it was this letter. Someone out there had clearly taken considerable trouble to search out this information for us, and to give some kind of guidance. It came at a time of fairly deep depression, when it seemed that there was simply no hope of getting international email for many years, and that there were too many people in this country who simply accepted that nothing could be done about this situation.

TCP/IP at Rhodes University

The Internet uses the TCP/IP protocols, so it is relevant to attempt to shed some light on how these protocols were introduced at Rhodes. There was no strong motivation, because the network that evolved worked well enough over twisted pairs of copper wires. All that was needed to connect a PC to the network was that the simple and low-cost cabling should extend to the location of that PC, and that the PC should have a serial card. Initially, this was fairly simple, but the policy of allowing (indeed, forcing) each department to find the funds from wherever for purchasing the PCs that they needed (as opposed to having a limited central pool of such funds) meant that expansion was rapid, and wiring up departments even with simple wiring became a major exercise.

Also, after a disastrous attempt to use Microsoft's LAN manager in the PC labs, we switched to Novell and then the installation of a campus-wide LAN became essential. Such a LAN was an ideal test-bed for experimenting with TCP/IP.

PCIP

The first version of TCP/IP to be used was that of the PCIP package that was developed at MIT and put on public release. In this particular case, a complete set of disks was purchased from Austin Code Works (??date??). This package supported only a few types of LAN cards, and so one of the specific cards was purchased at a wicked price. An IP number was obtained (192.42.99.0) and an attempt was made to try to get something to run.

FTP Software

The natural evolution from PCIP was to move (??when??) to the successor commercial product, viz PC/TCP, that was developed by FTP Software in Wakefield MA in the USA. Rhodes University's political credentials were important in securing this purchase because all too many US-based firms would have nothing to do with any institution based in South Africa. This package provided plenty of flexibility, and we updated things when the packet driver technology emerged. We also purchased the LanWatch package so that we could monitor ethernet traffic - this package alone played an important role in that it allowed us to get to grips with TCP/IP as low as at the bit-level when necessary.

SCO Xenix

Rhodes had acquired a 386 computer with the SCO Xenix operating system some time in 1989. This was a donation arranged by Dave Barnard of Central Data Systems. This acted as a vehicle for learning about the Unix-style of operating systems. This computer was initially given the name RUCCX1, and subsequently when we decided to use an "African animal" naming theme, was renamed QUAGGA. This Xenix host had been augmented by a second one by October 1989.

CDCNet

The NPU of the Cyber had been augmented with a CDCNet box. This had some simple TCP/IP facilities, and allowed telnet and FTP. It needed a hostname table file to drive it. By the 22nd May 1989 there were some 5 entries in the Rhodes' hostname table, and this had grown to 13 entries by the 3rd of October 1989.

IP Number Assigned

Rhodes received its first IP network number from the Internic on the 16th September 1988. This was the number 192.42.99.0, and was assigned by Sue Romano of SRI, (whom the author was to meet in October 1990 in order to discuss .ZA registration issues.)

NCSA and CUTCP

Helpful and supportive as FTP Software was, and as good as their software was, it was simply going to be beyond Rhodes's budget to install this package on each and every PC on campus. In due course, we discovered the NCSA and CUTCP public domain packages. These were fairly good, but were buggy. Some work was done at the University of Port Elizabeth by Darryl Anderson and/or Peter Lock that removed a timeout bug in the underlying software, and this helped a great deal to improve the performance. This fix was to the CUTCP package, but apparently the original authors took little interest in this fix - maybe it simply was too late in the day, and Windows™ packages were becoming the rage at the time.

TCP/IP Workshops

A number of Internet workshops were run at Rhodes University. These were specifically to deal with issues of internetworking, and it was required that the attendees should have some knowledge of how to install and run a TCP/IP package on a LAN.

The first workshop was held on 6th to 8th August 1990. This was specifically to plan a TCP/IP link between Rhodes and Potch, but the concept expanded to become a general sharing and pooling of TCP/IP knowledge and experience.

Attendees at this workshop were Alan Barrett (UND), Angus Hay (Wits), Jacot Guillarmod (RU), Gavin Melville (UPE), Chris Pinkham (UCT), Jorge Rabacal (CACDS/CSIR), Henry van der Walt (Potch), Jan van Rooy (Potch), Hennie Smit (Medunsa), Mark Stobbs (UFH) and Dave Wilson (RU). Of course, the author attended as well. The idea behind the workshop was literally to bootstrap the efforts to get TCP/IP connections running on some of the Uninet circuits. By that stage, the TCP/IP links to UCT and UND had been operating for two months, using the PCRoute package, and there was a growing use of this small internet.

More Workshops

There was little doubt that a core of Internet knowledge had built up at Rhodes, and there was also little doubt that there was considerable interest in this topic from various persons in the country. So, we (Jacot Guillarmod, Dave Wilson and the author) decided that it would be a good idea to share this knowledge, and to get some return on our efforts. Workshop #2 was hardly advertised, word was put out via the network as it then was towards the end of 1990. In no time, the workshop was oversubscribed. The names in the author's diary for that workshop are, in no particular order, Paul Nash, Paul Ducklin, Mark Elkins, Hans van Staden, Dennis Cronje, Sean Naude, Johan Strydom, Tony Putman, Heidi de Wet, Kuyper Hoffman, David Lurie, Aris Stathakis, Vic Shaw and Hennie Rautenbach. It is not known to the author at this point in time whether these all attended, or whether others attended in their place. This course ran from 7th to 9th January 1991.

Because this second workshop was oversubscribed, we organised a third one. This ran from 29th to 31st January 1991. Regrettably, the author does not have the names of the attendees (??help will be appreciated??).

Value of a boozy lunch

The last day of these workshops included a lunch at a local restaurant, which was but a short walk from the Rhodes campus. Most South African universities seem to have a reputation for alcohol consumption, and Rhodes is no exception. The "last day lunch" of these workshops rapidly degenerated, and no formal coursework was done on the last afternoon. Yet, somehow, it seemed to us (ok, we were looking through the haze) that somehow more was learnt about networking as a result of that lunch. Perhaps it was because the folk on these courses were indeed the real networkers of the country, who

were going to make things happen, and they got to know each other so much better. Who knows.

Theorem: There is a lot about networking that is never taught in any Computer Science course.

Still to be written/refined

TCP/IP on Uninet

There can be little doubt that the TCP/IP protocols were being used internally at a number of Uninet sites in the late 1980's. Rhodes was running SCO Xenix before 3rd October 1989. Running the TCP/IP protocols at a Uninet site was a pre-requisite for any attempt to link Uninet sites into a small internet.

This section describes some of the experiences in getting TCP/IP working between Uninet sites. What is surprising is that Uninet was cleared to connect .ZA to the Internet by 25th November 1990, yet it took a further 12 months before a circuit was installed and the connection established. The roll-out of TCP/IP at the South African universities was not done with any alacrity at all, and even today (1997) there is not a good depth of knowledge of TCP/IP at all of the Uninet sites.

Rhodes-UCT Linkup

The first link that took place was between Rhodes and UCT. Vic Shaw arranged that the Uninet multiplexors were reconfigured in a way that provided a 9,600 bps async link between the two sites. Routers were to all intents and purposes unobtainable, but we managed to purchase a package developed by Wollongong called Win/Route. This simply did not work, and the setback that this created nearly killed off further attempts. This must have been in about May 1990.

PCRoute

Given the failure of Win/Route, the attempt to establish a linkup between Rhodes and UCT almost died. Someone (?Chris Pinkham? Of UCT) came up with the idea of trying to use a public domain product called PCRoute. (The earliest reference to PCRoute that the author has is in a directory listing that he obtained from a USA site on 3 Dec 1989).

It made sense to get the networkers who were closely involved in this linkup to speak to each other face to face. Jacot and John Stevens (of Rhodes) attended a networking meeting at the University of Stellenbosch 4th and 5th June 1990, and spent some time at UCT. This had an immediate benefit, and work was in progress to use PCRoute on the

Rhodes/UCT link by the 7th of June 1990. It took only a few more days to get TCP/IP traffic to flow between Rhodes and UCT.

Later that month (June 1990) a similar link had been established between Rhodes and the University of Natal Durban, with the technical work there being done by Alan Barrett.

Expansion of this type of networking was surprisingly slow. The next link in the country was from the Rhodes gateway to Potch. This was some time after the TCP/IP workshop that was held at Rhodes on 6th to 8th August 1990. The oldest SMTP mail from Potch that the author has in his archives is from Jan van Rooy and is dated 20th August 1990, but there is a message implying quite clearly that the dial-up modem was no longer in use on the 27th June 1989 - possibly because the email was flowing across one of the Uninet virtual circuits as RSCS traffic, it was most certainly not as SMTP traffic.

From then on, the roll-out to other sites speeded up, as the method of getting IBM hosts onto the TCP/IP network had been established.

As an aside, Russell Vincent of the University of Natal, Durban, came up with an excellent scheme of allowing PCRoute to operate at 19.2 Kbps in asynchronous mode across Telkom's 64 Kbps synchronous Diginet service. This merely involved making up a suitable cable to connect the PC's serial port to the X.21 interface of the Telkom Network Terminating Unit (NTU). This was in operation in July 1992, it may have been used a year earlier than this date.

[The wiring of the cable is described here.](#)

Those !@\$(Default Routes

PCRoute used the RIP protocol for keeping track of routes to the various networks of Uninet, and in due course for the link to the USA as well. In essence, there had to be but one default route (ie route to network 0.0.0.0) announced on but one router within Uninet, and this had to be done by the router at Rhodes, because that was where the router to the USA was located.

Again and again what would happen, particularly when an internal network was added at a Uninet site, was that a default route would get set into the configuration file of the PCRoute that provided the connection. This would cause tremendous problems, as it would cause a failure at other points on Uninet when that incorrect default route propagated across Uninet. This would happen at sites that did not have a great deal of TCP/IP experience, and at sites that mistakenly thought that they had experts and this whole routing issue was a pushover. One highly undesirable side-effect of this is that failures of international routing were inevitably and unfairly blamed on the networking staff at Rhodes.

Interop '90

8th-12th October 1990, San Jose.

Visit to SRI

(18th or 19th) October 1990 - Sue Kirkpatrick (nee Romano) and Doug MacGowan. Led directly to the registration of the .ZA domain a month later.

Other Networking in South Africa

Lest anyone get the wrong impression, there were other networking activities in South Africa at the time. The ones that are of interest and relevance are those that interconnected different independent organisations, because this required a completely different approach to networking compared to the situation where one person or group controlled all of the computers that interconnected (eg in the case of a network of a bank, or of the government).

The Fidonet network has already been mentioned. Regrettably, the author does not know of the date when this network started, nor its history or extent.

One such network, which used the UUCP, protocol was called SANET, pronounced "essay-net". There was a bit of dispute over the ownership of that name, somehow this sorted itself out. This network also connected, via dialup, to the home of Randy Bush in Portland, Oregon. The author is uncertain of the timescales, but believes that this was after the UUCP to Rhodes had been operating for some months (ie in or after 1990). Further information on this network would be welcomed.

In September 1988, the FRD set about establishing a facility whereby the administration of grant funding could be done online. The concept was to have a microVax hold a database of the grant information for FRD-funded researchers, and that the various universities would be linked into this system via network technologies of X.25 or the CSIR's asynchronous network. It was also suggested that there could be links with NetLAN boxes and/or via IBM 3270 terminals connected by ways still to be defined. The author was involved with the service at the Rhodes University end, which was always problematic (for one thing, the system did not cater at all well for terminals that were not made by DEC). What is remarkable about this approach is that it did not dovetail into the Uninet design.

Hopefully, this history of the Internet will inspire others to publish similar stories about these networks.

Sanctions, Politics

Background

Politics of Rhodes University

Difficulties in obtaining Hardware and Software

Use of email communication

Un-closeable Gateway

The Fidonet gateway at Rhodes was established at a most difficult time in the history of South Africa. The apartheid government was under enormous threat from the majority of the populace, and was doing its utmost to control the flow of information out of the country. Never mind the government of the USA putting sanctions in the way of contact with South Africa, the South African government was imposing its will regarding communications, and was doing so with an iron fist. The government also viewed Rhodes University as a hotbed of liberal opposition, and had its spies infiltrate many aspects of the University's activities.

The author thus had solid grounds for concern that the gateway might be summarily closed down by the South African government. Fortunately, the design was such that the gateway could have been operated from anywhere near a telephone that could dial through to the USA. It would have been difficult to do this, but it would have been possible should the need have arisen. Fortunately, the need did not arise, but there were times when the author thought that it might. There was a goodly amount of anti-government information being sent out of that gateway.

Another coincidence still puzzles the author. It was almost a year to the day after the Fidonet gateway went into general operation at Rhodes (Feb 1989) that President F W de Klerk stood up in parliament and announced that Nelson Mandela would be freed, the ANC unbanned, and a host of other reforms (2 Feb 1990).

APPENDICES

Here are some of the documents that were used during the development, implementation and extension of the mailing systems at Rhodes University. These are the working documents as used by the group that did the development. Unless stated otherwise, specification documents were all prepared by the author of this article.

- [MAIL000: Specifications for Action to be taken by the Mailing System.](#) This was the specification of 1 March 1989, that showed how mail was to flow between the computers at Rhodes, onto Uninet and to the USA.
- [MAIL001: Some test addresses that were used.](#) The mailing systems had to produce the correct results in the appropriate files.
- [MAIL002: Comments made by Jacot on one of the specifications.](#) One that was recorded, out of many that were made.

- [MAIL003: Address formats for mailing](#). A description of how the user would type in addresses. The document is dated 7 February 1989, which is very close to the date that the Fidonet gateway was released for general use at Rhodes.
- [MAIL004: Sample mail between RURES and RUPHYS](#). This mail flowed on 10 January 1989, and showed the headers. The text of the message was used to confirm that the message lines were stored in the correct place on the Cyber.
- [MAIL005: Sample of kludged mail from RURES to RUPHYS](#). Even in January 1989 there was concern about spamming.
- [MAIL006: Another sample of kludged mail from RURES to RUPHYS](#). This example is here because it was found in the author's archives. It adds little to what is in MAIL005.
- [MAIL007: Changes required to the Cyber mail system](#). By 21st March 1989, it was clear that the home-brewed Cyber mailing system was creaking, and that changes were necessary.
- [MAIL008: Testing for weird patterns in the message](#). This is the text that was used to check that the text inside a message could look like the contents of a message header.
- [MAIL009: What to do with message headers](#). This document spelt out how we were to handle the headers of the messages. At the time (12 January 1989) some confusion had arisen and thus clarity was needed.
- [MAIL010: Test messages from RUPHYS to RURES](#). This tested that RURES would bounce badly addressed email back to the sender.
- [MAIL011: Some message headers from the early days](#). These message headers were used to show us what we had to cater for when developing the Fidonet gateway. It also gives some dates of the early dial-up Bitnet-style messages that flowed within South Africa.
- [MAIL012: Specification for logging Fidonet messages](#). This was quite likely not implemented. The early statistics showed that the cost of a message was only a few cents each, and it was not worth the effort to build an accounting system at that stage.
- [MAIL013: Provisional document to link a Unix system](#). This was an attempt in March 1989 to get a uucp gateway operating via a donated NCR Tower computer in the Rhodes Computer Science department. This could not be made to work, the project fell through.
- [MAIL014: Proposal for a major mail hub at Rhodes](#). It was clear by June 1989 that the Fidonet/Cyber gateway could and would be expanded to become quite big. The matter had now to be taken more seriously.
- [MAIL015: The specification to connect a computer to RURES](#). This document, dated 27th June 1989, set the standards to be used in order for other computers around Uninet to connect to the RURES mailer. From this mailer, it was possible to exchange email with any other of the Uninet hosts of whatever manufacturer, as well as with hosts on the Internet.
- [MAIL016: Mail flows at Rhodes University](#). A document by Jacot, dated 16 June 1989, giving an overview of mail flows at Rhodes.

- [MAIL017: Changes to the Fidonet Gateway](#) By the 21st of November 1989, the Fidonet gateway was creaking quite badly. This document spelt out what changes were needed.
- [Some info about the first TCP/IP workshop at Rhodes.](#) This has the preliminary announcement, as well as the actual timetable.
- [Common Interest Group Application.](#) This was made to SAPO (predecessor to Telkom) early in 1988.
- [South Africa cleared to connect to the Internet.](#) Here is some correspondence on that topic. It was to be a full year before the Internet connection was established.

Some Traffic Figures for Fidonet

Here are some traffic figures for Fidonet, extracted from a note by the author to Pat Terry on 16th July 1990:-

| Month | Kbytes Compr. Mail | Telephone Units (1.75 secs) |
|----------|--------------------------|-----------------------------------|
| October | 9,688 | 23,000 +- |
| November | 10,612 | 20,000 +- |
| December | 19,573 | 15,273 |
| 1990 | | |
| January | 14,630 | 25,842 |
| February | 17,979 | 14,872 |
| March | 26,127 | 23,944 |
| April | 28,958 | 31,553 |
| May | 39,745 | 28,446 |
| June | 25,906 | 41,028 |

Uninet Sites as at 13 September 1989

(Source: Original email from Alan Barrett)

Electronic Mail
13 Sep 1989

Wed,

=====

This document is an attempt to provide sufficient information to permit users of the UND EE VAX and HP9000 to make use of the mail facilities that exist between these computers and various other computers and networks in the RSA and overseas.

Alan Barrett, Dept of Electronic Eng, University of Natal, Durban
UniNet (RSA): Barrett@UNDEE
Phone: (031)-816-2758

... [cut] ...

| Node Name | postmaster user name | Node Description |
|------------------------------|---------------------------------|------------------|
| ===== | ===== | |
| CSIRVM machine | uninet@csirvm | CSIR - VM |
| JES2 machine | (not active) | CSIR - TSO |
| PMB1UN Pietermaritzburg | kure@pmb1un | Water Research, |
| PUKVM1 University | rkdjlvvr@pukvml | Potch |
| RGNVM1 Research Council | (not active) | Human Sciences |
| RUBIS department | jhobson@rubis (ADDS/Mentor) | Rhodes Univ BIS |
| RUCS01 Comp Sci Dept | postmaster@rures (NCR/Unix) | Rhodes Univ |
| RUPHYS Physics Department | postmaster@ruphys (VAX/VMS) | Rhodes Univ |
| RUPLA Comp Cent - Admin | postmaster@rures Plato (NOS) | Rhodes Univ |

| | | |
|--------------------------------|-------------------|-----------------|
| RURES | postmaster@rures | Rhodes Univ |
| Comp Cent - Research (NOS) | | |
| UCTCS | postmaster@ucthpx | Univ of Cape |
| Town, Comp Sci Dept. | | |
| UCTGEO | postmaster@ucthpx | Univ of Cape |
| Town, Geochem Dept. | | |
| UCTHPX | postmaster@ucthpx | Univ of Cape |
| Town, Info Tech Services | | |
| UCTVAX | postmaster@ucthpx | Univ of Cape |
| Town, Info Tech Services | | |
| UNDEE | postmaster@undee | Univ Natal, |
| Durban, Elec Eng (generic) | | |
| UNDEEB | postmaster@undeeb | Univ Natal, |
| Durban, Elec Eng (VAX/VMS) | | |
| UNDEED | postmaster@undeed | Univ Natal, |
| Durban, Elec Eng (HP9000/Unix) | | |
| UNDHPD | postmaster@unphpd | Univ Natal, |
| Durban, HP Desk. | | |
| UNIS | (not active) | University of |
| South Africa - TSO machine | | |
| UNPHPD | postmaster@unphpd | Univ Natal, |
| Pietermaritzburg, HP Desk. | | |
| UPVM2 | bdr0015@upvm2 | University of |
| Pretoria - research | | |
| WITSVMA | 074spock@witsvma | Wits University |
| - research | | |

FOR CHECKING

At Rhodes:

- installation date of 9,600 bps circuit to Randy Bush

[NOTE of Nov 2001., based on old files at the NRF]

There is a letter dated 24 June 1991 from Vic Shaw to the Assistant General Manager, International Services (3C31C3), SAPT that applies for a 9,600 bps

dataline to the USA. This letter was also sent to the Postmaster General, Commercial Section (3K92), SAPT on the same date.

The request was acknowledged by the SAPT on 22 July 1991 - ref 3C31C31/OS PW 49, D van der Walt -, for the circuit between Grahamstown (ie Rhodes University) and Portland Oregon (ie Randy's home). The serial number of the application serial number was quoted as 26997.

The author well recalls the delays that SAPT/Telkom introduced in installing this circuit. The cost of the dialup calls exceeded the cost of such a leased line by a factor of about 3 for a number of months, and this hurt all of the Uninet sites. There is a letter from Telkom (formerly SAPT) dated 18 December 1991, ref 3C31C3/91, signed by or on behalf of the General Manager International Services, that denies being aware of the key issues relating to the application, and stating "we arranged in good faith for the provision of the leased circuit to be delayed pending a clearup of uncertainties that surrounded the matter.". Regrettably, Telkom did not see fit to be pro-active in involving Vic Shaw in clearing up these "uncertainties", as is obvious from the later paragraphs of this letter. This letter states that the leased line was operational on 15 November 1991.

[End of NOTE of Nov 2001]

- date of letter from Elizabeth Kilkoyne of Bitnet
- date of purchase of PCIP
- date of purchase of FTP Software package
- date of purchase of WinRoute package (1988/9??)
- date of donation of 386 + Xenix by Central Data Systems (Dave Barnard) (1989??)
- date of installation of TCP/IP on VAX

In my mail archives

- Mark M Dickey (Interop, 1991) re prices of Ciscos
- date of first email to/from Potch