

1.4.2. Meeting Notes - 1 February 1978

Agenda

- 8:30- 8:40 Introduction and Objectives - Cerf
8:40- 8:45 Arrangements - Cohen
8:45-10:30 Status Reports
1. PR to Xerox PARC delivery schedule -
 2. LCSNET status -
 3. UCL X.25 implementation status -
 4. MIT gateway implementation -
- 10:30-11:30 Gateway Software Plans - Strazisar
Alternate routing proposal
Flow control proposal
Software implementation schedule
- 11:30-12:00 Gateway Monitoring Formats - Cole
12:00- 1:00 Lunch
1:00- 2:00 Discussion of Gateway Plans
2:00- 3:00 Naming, Addressing and Routing - Cohen
3:00- 4:00 Checks and Balances/Access Control - Ray McFarland
4:00- 4:30 Agenda for next Meeting - Cerf

Arrangements - Cohen

Danny gave a quick overview of the plan for the day and a brief description of the physical facilities.

Introduction and Objectives - Cerf

Vint primarily presented a summary of the results of the TCP meeting which took place the preceding two days.

The main result is a decision not to extend the TCP to include alternate modes of service, and instead to provide for a set of parallel protocols that all utilize a common "datagram" type internet protocol. The functions of the internet layer would be addressing and fragmentation. It was decided that every network must be able to handle an internet packet of 1008 bits without fragmenting it. The fragmentation will be done in terms of units of 64 octets. The fragment offset field was set at 8 bits, leading to a maximum internet segment length of 2^{14} octets. There is also provision for a bit indicating that fragmentation is not allowed. There are implications about routing and error reporting mechanisms needed in the gateways that should be explored.

Other TCP meeting results were a decision to revise the TCP3 specification to correct minor points, and sometime in the future to

undertake a major revision of the documentation. Another activity is the study of the inter-relationship of various mechanisms and services of TCP to understand the dependencies between them.

Status Reports

1. MIT gateway implementation - Dave Reed

A local net interface is being tested as a ring interface, three more are to be shipped to MIT within a week or two. The "Etherization" of the ring is in design. There will be a gateway between the LCS ring and the ARPANET based on a PDP-11 Unix (MIT version).

2. PR to Xerox PARC delivery schedule - John Shoch, Vint Cerf, and Ron Kunzelman

A Packet Radio unit is now installed at PARC and acting as a repeater. The first experiment is to use the PRNET as a transit path between two Ethernets. A second experiment may be to have a mobile Alto access an Ethernet via the PRNET. The first interface between an Alto and the PR unit is scheduled for June.

3. UCL X.25 implementation status - Peter Kirstein

UCL is working on X.25 interfaces for terminals based on both a LSI-11 and a 8080. Also working on X.25 interfaces between networks, both via gateways and directly. UCL is also and quite independently working on some interfaces to EURONET. The LSI-11 X.25 terminal will utilize the EURONET "Virtual Packet Terminal Protocol".

Peter fills us in on EURONET -- To understand EURONET one must understand TRANSPAC since the hardware & software are being prepared by the same supplier. The PTTs view EURONET as the international net that interconnects the individual national nets, i.e. EURONET has as users only the PTTs. EEOC which sponsors EURONET does not share that view. EURONET to national net connections will probably of the X.7.X form. There is some inconsistency in the EURONET specifications of terminal access and host access, as well as between the specifications and the proposed applications, yet to be resolved.

There is to be a major demonstration of EPSS in the spring, and UCLs efforts will be concentrated on that until it is completed.

Gateway Software Plans - Strazisar

Alternate Routing and Congestion Control

Ginny has two papers to be distributed: "Gateway Routing" (available at the meeting) and "Gateway Distributed Routing" (being mailed).

BBN is working on a MiniGateway, the current system is a PDP-11/40 ELF, the next step is a PDP-11/40 MOS, and finally a LSI-11 MOS version.

Routing Design Goals

- Low Delay
- High Thruput
- Rapid Response to failure and recovery
- Stable and Low Overhead in the Steady State

Constraints

- Distributed
- Allow for reconfiguration
- Simple

The current view is that each gateway knows about all other gateways in the same net, each gateway is in at least two nets.

Routing

What information should be exchanged (examples)

- Connectivity
- Hop Count (each net = 1 hop)
- Delay
- Marginal Delay

How information should be exchanged (examples)

- Reliability -- Events
- Unreliability -- Status

When information should be exchanged (examples)

- Event Driven
- Periodic
- Both

How should information be used (examples)

- What is the algorithm ?

Current Tentative Plan

Use Hop Count, can later be refined to any associative measure.

This implies that each gateway must have a table of all nets in the internet system.

Congestion Control

The gateways measure the performance across the networks they are connected to, for example, percentage of packets lost, or queue lengths. Performance information is reported to neighbors and propagated. The information is used to split the traffic between equal length paths, and if poor performance information is propagated all the way to the source, the source is quenched.

Gateway Monitoring Formats - Cole

The gateway monitoring message format is as follows:

Local Net Header

Internet Header

Format = 4

Default Destination = Net 10 (8 bits), Host 5 (24 bits)

Code (8 bits)

Bit 8 (high order set) = byte reversal (for pdp-11's)

Bits 7-0 Code indicating Thruput type

Number of Nets (8 bits)

Sequence Number (16 bits)

For Each Net:

Thruput and status information, trap information.

The collection facility running on a host (e.g. a Tenex) receives these messages and puts them in a data base. Programs can be written to abstract messages and summaries of messages. The ability to selectively disable active probes of gateways is requested.

The collection facility may on occasion need to send a message to the gateways. It is suggested that a "GATESTAT" program be made available that would portray some simple status summary of currently active gateways.

It is pointed out that the current SATNET monitoring via the SIMPs has this organizational model.

Checks and Balances/Access Control - Ray McFarland

In the military the managers of networks are very concerned about the accountability, self protection, and efficiency of gateways. Accountability measures may include flow analysis (number in compared to number out). There is some confusion about current practice and its merits, Ray will report in more detail on this at a future meeting. The following items might be useful to collect for accountability: time sent, time received, number discarded, number through, gateway source-destination pair, number of packets returned. In the area of self protection, possible strategies are: Restricted Routing, Entry Protection, and Access Control. One efficiency

concern is whether retransmission is between hosts or between gateways.

Review of the Internet Header - Ray Tomlinson

Ray reviewed the state of the internet header before the TCP meeting and the new fields to be added. The new fields are a 16 bit ID field, a 8 bit fragment offset, a not-end-of-segment bit, and a don't-fragment bit. Ray reviewed the fragmentation procedure and pointed out that uniqueness of a segment is determined by the concatenation of the four fields: ID, Source, Destination, and Format. Ray also noted that (in principle) the Version field is not part of the internet header, but part of what came before the internet header, and tells what kind of an internet header follows. Jon Postel is to prepare a revised internet header format and circulate it to this group.

Irwin Jacobs suggests that the TOS field should be at least 8 bits, to allow for the SATNET types of service.

Naming, Addressing and Routing - Cohen

John Shoch made a very useful presentation on the difference between Names, Addresses, and Routes, and made available a memo on the subject (IEN 19).

A Name is what we seek, an address is where it is, and a route is how to get there.

Names are independent of location, "unique" in some name space, used to identify resources, and may be generic. There is a mapping from names to addresses, a name may map to more than one address, and a resource may have many names.

An address defines an addressible unit. The address space may be flat or hierarchial. There is a mapping of addresses into routes. A hierachical space allows multiple address assigning authorities.

Routes may be of one or many hops, the path may be assigned at the source, determined on a hop-by-hop basis, or by a hybrid scheme. The control of routing decisions may be isolated, centralized, or distributed. The routing decisions may be made with fixed or dynamic information. The binding of name to address to route is to be defered as long as possible to maintain the maximum flexibility (this is a lesson from operating systems and high level language development). When confronted by an error, back up, that is choose a different binding of address to route, or, if that also fails, of name to address, etc.

Hierarchical addresses aid in issuing new addresses, aid routing, aid moving resources, but there are interactions between topology and routing and addresses that must be carefully investigated. If strict use of fields is enforced information may be lost. To use shortcut routes it is necessary to use something like the phone systems "6 digit translation". The possibilities for high use shortcuts, reaching into partitioned nets, and off net routing should be preserved.

Hybrid routes contain a few intermediate points but leave open the route between the named points. There are two ways: to list all the named points in the header, and to list the first in the header and the rest in the data. In this second case, each time a message reaches a named point a forwarder process there replaces the address in the header with a new address obtained from the data field. An interesting, but perplexing, case is a negative route -- don't go through any of these named points.

Much discussion followed with examples from the telephone system and the ARPANET. One question brought out was how to determine the the meaning of the first field, is it a network, a host, or a process address?

Danny is to prepare a memo on the addressing issue.

Schedule of Events - Cerf

- 1-Feb ARPANET/EPSS paper available - Kirstein
- 8-Feb Internet Header Specification - Postel
- 15-Feb Preliminary Spec of Gateway Monitoring Functions - Cole
- 1-Mar Tops-20 traffic generator for SATNET - Kirstein
- 1-Mar ARPANET - X.25 interface for Terminal - Kirstein
- 15-Mar Revised Specification of Gateway Monitoring Functions - Cole
- 1-Apr "Gatestat" program available - Cole
- 1-Apr Informal Specification of Gateway - Strazisar
- 1-Apr Schedule for Gateway Implementation - Strazisar
- 1-Apr Memo on Accountability in Networks - McFarland
- 1-Apr EPSS - X.25 interface for Terminal - Kirstein
- 1-Apr ARPANET - X.25 interface for FTP - Kirstein
- 8-Apr Memo on Addressing Issues - Cohen
- 1&2-May Internet Meeting - UCL
- 1-Apr Demonstration of ARPANET-EPSS at meeting - Kirstein
- 1&2-Jun TCP Meeting - MIT-LCS
- 1-Jun PARC PRNET-Alto Interface - Shoch
- 1-Jun Minimal Gateway at MIT - Reed

Agenda for next Meeting - Cerf

Status Reports

- MIT gateway - Reed
- PARC gateway - Shoch
- BBN gateways SATNET/PRNET/ARPANET - Strazisar
- UCL gateways - Kirstein
- Routing/Addressing/Naming - Cohen
- Gateway Congestion Control & Routing - Strazisar
- Gateway Monitoring - Cole
- Accountability - McFarland
- Gateway Performance
 - UCL - Kirstein
 - BBN - Cole/Strazisar

Memos Distributed

- 1) A packet of updates for the internet notebook - Postel
- 2) IEN 19 on Naming, Addressing, and Routing - Shoch
- 3) IEN 20 on Fragmentation - Shoch
- 4) On Names, Addresses, and Routings - Cohen
- 5) Gateway Routing - Perlman
- 6) Gateway Monitoring Messages - Cole

Attendees

Nick Abel	LINKABIT	ABEL@ISIE	(714)453-7007
Richard Binder	BBN	BINDER@BBNE	(617)491-1850
David Boggs	XEROX-PARC	BOGGS@PARC	...
Ed Cain	DCA	ECAIN@BBNE	(703)437-2247
Vint Cerf	ARPA	CERF@ISIA	(202)694-3049
Wesley Chu	UCLA	CHU@ISIE	(213)825-2047
Danny Cohen	ISI	COHEN@ISIB	(213)822-1511
Jon Cole	BBN	COLE@BBNE	(617)491-1850
Sunil Das	UCL	KIRSTEIN@ISIA	UK: (01)387-7050x817
Gunnar Fredrikson	NTA	AAGE@SRI-KA	Norway: (02)739211
Jim Forgie	LINCOLN LABS	FORGIE@BBN	(617)826-5500x7191
Estil Hoversten	COMSAT	HOVERSTEN@ISI	(202)554-6092
Irvin Jacobs	LINKABIT	JACOBS@ISIE	(714)453-7007
Peter Kirstein	UCL	KIRSTEIN@ISIA	UK: (01)387-7050
Ron Kunzelman	SRI	KUNZELMAN@ISIE	(415)326-6200x4118
Yngvar Lundh	NDRE	YNGVAR@SRI-KA	...
Jim Mathis	SRI	MATHIS@SRI-KL	(415)326-6200x5150
Ray McFarland	DOD	DEEDWARDS@ISIA	(301)796-6290
Bill Plummer	BBN	PLUMMER@BBN	(617)491-1850x234
Jon Postel	ISI	POSTEL@ISIB	(213)822-1511
David Reed	MIT-LCS	DPR@MIT-ML	(617)253-6004
Michael Shen	UCLA	...	(213)825-2520
John Shoch	XEROX-PARC	SHOCH@PARC	...
Ginny Strazisar	BBN	STRAZISAR@BBN	(617)491-1850x632
Steve Treadwell	UCL	KIRSTEIN@ISIA	UK: (01)387-7050
Ray Tomlinson	BBN	TOMLINSON@BBN	(617)491-1850x363