DEC 21 '92 08:47AM ICC P.1

MFS-ICC
8100 BOONE BOULEVARD, SUITE 500
VIENNA, VIRGINIA 22182
TEL (703) 506-2000
FAX (703) 847-9426

TELECOPY COVER SHEET

DATE: 12/21/92

TIME: 8:15AM

TO: Scott Yenser

FROM: Kathleen Davis

RE: WUNet MSA, ANS Perf. Specs (DS-3), FYI - MAP WUNet intends to use in their slide show (I will have Doug make it up).

PHONE NUMBER OF RECEIVING PARTY: 713/236-9637

TELECOPY NUMBER: 713/225-3616

NUMBER OF PAGES INCLUDING COVER SHEET: 16

INSTRUCTIONS: Stan

Read WUNet ANS-spec's & Kathleen's comments please

MFS-ICC FAX #: 703-847-9426 OR 703-356-5066

CONTACT SENDER IMMEDIATELY IF ANY PAGES ARE MISSING

703-209-8000

5100 MS/km
-mili -1,000 $\cdot 10^{-3}$
micro 1,000,000 10^{-6}
NS. 10^{-9}

Date: December 20, 1992
To: Scott Yeager
From: Kathleen Davis
Subj: UUNET M5A ISSUES

#1 10 Mbs aggregate between 9 cities? Aggregate per connection or aggregate per cloud between 9 cities. Take out aggregate or specify.

#2 Requesting end to end connections or not beyond 50-100' cabling for UUNET responsibility otherwise too ambiguous regarding cost for interface e.g. beyond 100', another floor, a remote MFS hub 5 miles away and now we are back into local loop charges. "Costs should be factored in at the beginning not after the fact" and "AUI interface attached to the UUNET router." The MFS Datanet demarcation point will be the female Ethernet AUI interface to a UUNET which will be the demarcation point.

AL & Bill

#3 The MFS Datanet backbone will be provisioned to provide a diverse "full mesh backbone" path to reroute packets...

#4 To be implemented on a 6-8 week internal basis per city once MFS opens up a city for 10 Mbs connectivity. UUNET has the option of opening up city with 1.5Mbs bandwidth at fractional rates. The language should reflect a reasonably definitive starting date to start the aggregate $63k price. Liability on this discounted pricing is an issue here.

P9

Implementation Schedule

Atlanta	6wks > Hous
Boston	last city 6wks > Atl
Chicago*	fract'l 1/4/93 pending Datanet upgrade
Dallas	Collocation in/ pending Datanet upgrade
Houston	6wks > Santa Clara
Los Angeles*	1/22/93 fract'l upgrade pending Datanet
New York*	fract'l in pending Datanet upgrade ASAP
Santa Clara	6wks > LA 3/93
San Francisco**	possible for coll only by July '93
Wash, DC*	fract'l in/ pending Datanet upgrade ASAP

#5 Usage sensitive component doesn't make sense.... why can't we count packets?, this could end up costing him big $$$$'s

#6 "The charges" for the usage sensitive component ....

#9 A) In the event that the "Virtual" National 10 Mbs Ethernet...
*** B) $7k per city is an issue in terms of when he will turn up cities. Rick does not want to be penalized for turning up a city

MFS-ICC

Post-It™ brand fax transmittal memo 7671 | # of pages ▸ 10

To Bill Emola	From Scott
Co.	Co.
Dept.	Phone #
Fax #	Fax #

prior to the aggregate billing date/liability on $63k by paying $7k per city before he gets full benefit of the nine city full mesh diverse backbone network.

Insert #10 Performance Measurements - Error Rates/Packet Loss/Latency (Addresses routing on the network) Delay factor? Cross country T-1 averages 15-20 ms and the MFS link from DC to NY is taking 40 ms - is this intentional broadcast for testing or is it a routing issue? Worse performance from DC to NY than DC to BOS. (FYI, The circuit to NY is less than acceptable and Rick is not going to pay for it until performance comes up to at least T-1 standards for latency. Rick has asked Sprint for the same kind of performance rating and that a circuit is operational at this level for at least a 24hr period.) UUNET is performing some evaluations (200 ping sampling) to determine latency averages. They will try to come up with T-1 latency figures but expect latency on 10 Mbs to be better and will require some language in the MSA to accommodate these standards and credits for poor performance as well as easement on the termination liabilities.

I also have a copy of some ANS specs on latency that were given to Rick that I will fax along to you. Potentially take ANS @ DS-3 figures and add 30% for 10 Mbs performance.

"Establish performance requirements so that UUNET doesn't "pay" for debugging."

FORMULA - (Expectation of better than T-1 Latency) Round Trip time for a 32 Byte packet shall be no longer than 3ms* + .025 x the circuit mileage e.g. DC to Santa Clara shall be no more than 60 ms. When the numbers get small there is a loss of accuracy, A+BxZ doesn't work when trying to fit the curve @ zero.

* 2ms locally (in one room) 1ms into the equipment, 1ms out of the equipment + mileage e. g. Local DC Ethernet to Ethernet/Sprint to PSI is 5 ms.

MASTER SERVICE AGREEMENT

#1 MFS Datanet, Inc. will provide a Virtual Ethernet backbone to UUNET Technologies, Inc. (UUNET) that ~~can burst to 10 mbs aggregate between 9 MFS on net cities~~ at the charges indicated in this agreement. This agreement represents special terms and conditions that address the unique requirements of UUNET.

#2 The MFS Datanet demarcation point will be the female Ethernet AUI interface on the MFS Datanet equipment. It will be UUNET's responsibility to provide the male Ethernet AUI interface into MFS Datanet's equipment, and any associated costs. In each city, UUNET will co-locate with MFS at either the MFS node or an MFS POP, to be determined as soon as possible. MFS Datanet will be responsible for providing 7 x 24 hour maintenance and network monitoring of the service as part of the monthly recurring charges.

#3 The MFS Datanet backbone will be provisioned to provide a diverse path to reroute packets in the event of a circuit interrupt occurring between MFS cities. In the event UUNET would lose connectivity in any city on the backbone, UUNET would have connectivity between all other cities on the Virtual Ethernet backbone.

#4 The Ethernet backbone will interconnect Chicago, Boston, New York, Atlanta, Washington, D.C., Houston, Dallas, Los Angeles and San Francisco at a burst data rate up to 10 mbs for the aggregate cost of $63,000/mo. plus usage sensitive component as defined below to be implemented by July 1, 1993. Non-Recurring Charge is $1,000/one time per city at the time of installation for each city.

#5 The usage sensitive component will be calculated in the following manner. UUNET will be able to burst up to 10 mbs at any time during the month but will pay a premium above the monthly recurring $63,000 rate under the following conditions. UUNET will pay $4.80 per 128 kilo bits of bandwidth above 2 Mbs as measured on an average of any 20 minute time interval in the month. Only 400 intervals of 20 minutes that burst above 2 Mbs data rate will be counted in any month. The 400 intervals will be those 20 minute intervals with the largest average usage above 2 mbs in that month.

#6 The usage sensitive component of this service offering will not begin until July, 1993.

The prices offered in this agreement are confidential and proprietary. UUNET agrees to keep these confidential per the attached confidentiality agreement.

This agreement will supersede the prior agreement for service between N.Y., Chicago and Washington D.C. as of March 1993 when the National 10 Mbs Ethernet Backbone becomes available.

#9 In the event that the national Mbps Ethernet backbone becomes available prior to the March time frame, and UUNET chooses to upgrade the three city service (NY-CHI-DC) to 10 Mbps, the rate for this service will be $14,000/month until the other cities are added. If UUNET chooses to roll-out the cities over several months, the rate will increase by $14,000 for the first west coast city, and $7,000 for the rest of the cities until all cities are added. UUNET shall assign an implementation schedule for each city as part of this agreement as identified in Attachment A. UUNET has until July 1, 1993 to add all nine cities to the Ethernet backbone to guarantee the aggregate cost of $63,000/month

#10 The term for this agreement is eighteen months.

★ Termination Liability for services installed between the date of this agreement and June 30, 1993 shall be 100% of the monthly recurring charge for the remainder of the contract term.

MASTER SERVICE AGREEMENT (continued)
Page Two

- If UUNET does not reach the (9) city commitment for the $63,000 monthly aggregate price, MFS reserves the right to retroactively bill back UUNET for the difference between the special price offered and the standard national network price

- Termination Liability for services installed as of July 1, 1993 will supersede prior liabilities and will be 100% of the aggregate total of $63,000 per month for the remainder of the contract term or until July 1, 1994.

This agreement is entered into between MFS Datanet, Inc. and UUNET on this 18th day of December, 1992.

UUNET TECHNOLOGIES, INC.

BY: _____
President

MFS-ICC, INC.

BY: _____
Vice President

MFS DATANET, INC.

BY: _____
Vice President of Sales and Distribution

ATTACHMENT A

Virtual 10 Mbs Ethernet Backbone

NATIONAL NETWORK IMPLEMENTATION SCHEDULUE

CITY	IN SERVICE DATE
ATLANTA	
BOSTON	
CHICAGO	
DALLAS	
HOUSTON	
LOS ANGELES	
NEW YORK	
SAN FRANCISCO	
WASHINGTON, D.C.	

NON-DISCLOSURE AGREEMENT

This Agreement, made and entered into this 18th day of December, 1992 (The "Effective Date") is by and between MFS Datanet, Inc. (hereinafter "MFSDN") and UUNET Technologies, Inc. (hereinafter "UUNET").

WHEREAS, UUNET and MFSDN have entered into discussions and negotiations during which the parties may disclose certain information related to their operations and business; and

WHEREAS, UUNET may request MFSDN and MFSDN may request UUNET to disclose certain information relating to its operations and business; and

WHEREAS, the parties acknowledge that their discussions and exchanges of information could lead to some form of business relationship in the future; and

WHEREAS, MFSDN and UUNET desire to protect the confidentiality of certain information that it may disclose to the other party;

NOW, THEREFORE, for an in consideration of the mutual promises contained herein, and of other good and valuable consideration, the receipt and adequacy of which is hereby acknowledged, the parties hereto, intending to be legally bound, do agree as follows:

1. Definitions. For purposes of this agreement, the following definitions shall apply:

   (a) "MFSDN" shall include MFS Datanet, Inc., its parent company, subsidiaries, and affiliate, if any.

   (b) "UUNET" shall include UUNET Technologies Inc., its parent company, subsidiaries, and affiliates, if any.

   (c) "Confidential Information" shall mean any information, including but not limited to trade secrets, know-how, formulas, processes, data, network configuration and rights-of-way, drawings, proprietary information, customer lists, prices, any non-public information which concerns the business and operations of a party to this agreement, except for that information specified in Section 6 of this agreement. Confidential Information, when disclosed in written, machine-readable, or other tangible form by one party to the other party, shall be clearly marked as "Confidential". Information which is disclosed orally and is of an apparent confidential nature shall be treated as Confidential Information by either party and used only according to the terms of this Agreement. Each disclosure of Confidential Information shall be documented by the disclosing party, and in the case of oral disclosures such documentation shall be provided to the opposite party within ten (10) working days in the form of a written notice, summarizing those matters designated as Confidential Information. This written notice will be the best, but not the only , evidence of such disclosures, and other oral disclosure may later be proven by clear and credible proof and evidence in the event that either party shall consider that there is in fact a need for such proof.

2. Use of Confidential Information. During the course of the discussions between the parties, MFSDN and UUNET may disclose certain Confidential Information, either by verbal or written communications. These disclosures will be made upon the basis of the confidential relationship between the parties and upon their agreement that, unless specifically authorized in writing by the other, they will:

(a) use such Confidential Information solely for the purpose of evaluating the proposed business relationship between MFSDN and UUNET;

(b) promptly return to each other, upon request, any and all tangible material concerning such Confidential Information, including all copies and notes, whether such material was made, compiled by, or furnished by UUNET or MFSDN.

3. Nondisclosure. Each party agrees to receive the Confidential Information in confidence. Each party agrees that it will treat such Confidential Information in the same manner as it treats like information of its own that it does not wish to disclose to the public, but in all events it shall use at least a reasonable degree of care. To that end, neither party will make any copies of Confidential Information of the other party will that is in documented form except for use by employees with a need to know.

Each party further agrees not to distribute, disclose or disseminate in any way to anyone, except its employees or subcontracts who have such need to know, except as authorized in writing by the other party, any of the Confidential Information of the other party in any form whatever. Each party agrees that is disclosure of Confidential Information to its employees who have such a need to know shall be limited to only so much of such Confidential Information as is necessary for that employee to perform his/her function.

4. No License. Nothing contained in this Agreement shall be construed as granting or conferring any rights by license or otherwise in any Confidential Information disclosed to the receiving party.

5. No Obligations. The furnishing of Confidential Information hereunder shall not obligate either party to enter into any further agreement or negotiation with the other or to refrain from entering into an agreement or negotiation with any other party.

6. Exceptions. The obligations imposed herein shall not apply to Confidential Information:

(a) Which becomes available to the public through no wrongful act of the receiving part; or

(b) Which may be published prior to the date hereof; or

(c) Which is received from a third party without restriction and without breach of this Agreement; or

(d) Which is independently developed by the receiving party; or

(e) Which is disclosed pursuant to a requirement or request of a government agency.

Nothwithstanding the above, all materials, including without limitation, dcouments, writingt, designs, drawings and specifications furnished and that are designated as Confidential Information shall remain the sole property of the disclosing party and shall be returned promptly to the disclosing party at its request with all copies made thereof.

7. Termination of Obligations. The obligation to protect the Confidential Information received hereunder shall continue for five years following the return of the Confidential Information pursuant to subparagraph 2(b) or paragraph 6, or two years following provision of the Confidential Information, unless a specific request is mutually agreed upon to maintain the information on a confidential basis for a longer period of time.

8. Entire Agreement. This Agreement and any Consulting Agreement constitutes the entire agreement between the parties and supersedes any prior or contemporaneous oral or written representations with regard to the subject matter thereof. This Agreement may not be modified except by a writing signed by both parties.

9. Enforcement. All provisions hereof for the protection of each party are intended to be for each party's benefit and enforceable directly by each party. Each party agrees that any remedy at law for any actual or threatened breach of this Agreement by the other w ould be inadequate, and that each party shall be entitled to specific performance hereof or injunctive relief or both, by temporary remedy, writ or orders as may be entered into by a court of competent jurisdiction in addition to any damages that the harmed party may be legally entitled to recover, together with reasonable expenses of litigation, including attorney's fees incurred in connection therewith as may be approved by such court, and each party further agrees to waive any requirement for the securing or posting of any bond in connection with obtaining any such injunctive or equitable relief.

10. Notification. Any notices, requests, consents and other communication under this Agreement shall be in writing and shall be deemed to have been delivered (a) on the date personally delivered, (b) on the date mailed, postage prepaid, by certified mail with return receipt requested, or (c) telegraphed and confirmed. All communications under this section shall be addressed to the respective parties hereto as follows:

If to UUNET TECHNOLOGIES, INC.

Richard L. Adams, Jr
UUNET Technologies, Inc.
3110 Fairview Park Drive, Ste. 570
Falls Church, VA 22042

If to MFS Datanet, Inc.

F. Scott Yeager
MFS Datanet, Inc.
55 S. Market
San Jose, CA 95113

11. This Agreement may be executed in counterparts, each of which shall be deemed an original, but all of which shall constitute the same instrument.

12. This Agreement shall be construed and enforced under and in accordance with and governed by the laws of the State of California.

MFS Datanet, Inc.

By: _____

Name: F. Scott Yeager

Title: V.P. Sales and Distribution

UUNET TECHNOLOGIES, INC.

By: _____

Name: Richard L. Adams, Jr.

Title: President

closed local ENSS locations. The tail circuits of the existing network attachments to this POP will be reduced to local access circuits only.

5) The installation of a new CNSS in Atlanta is scheduled for 9/26 to reduce the GA Tech T3 tail to local access only, and provide expansion capability in the southeast.

T3 Network Performance Enhancements
==========

The general approach to engineering the T3 network has been to prioritze enhancements that improve stability rather than performance. Since the T3 network RS960 upgrade in May '92, the stability of the network has become very good, and we have been able to spend more resources focusing on the performance of the network, which has also improved significantly. With the upcoming deployment of the new RS960 FDDI adapter, we expect to observe higher peak bandwidth utilization across the T3 network, and higher aggregate packet traffic rates. In anticipation of this, we have conducted some baseline performance measurements on the T3 network that serve as a basis for continued tuning and improvement over time.

T3 Network Delay

In order to analyze the delay across the T3 ANSNET, we start by measuring the delay incurred by each T3 router hop, and then measure the circuit propagation delay across all backbone circuits. We have MCI T3 circuit route mileage figures which can be calibrated with PING measurements to determine how much each hop through a T3 router adds to the round trip time.

A set of round trip delay measurements was made using a special version of PING that records timestamps using the AIX system clock with microsecond precision. The technical details of the measurements may be described in a future report on the subject. The end result is that the round trip transit delay across a T3 router was measured to be about 0.33 ms (0.165ms one way delay), with a maximum variance between all samples on the same router of 0.03 ms. The T3 routers currently experience very little variance in delay at the current load on the T3 network. The T3 router transit hop delay is therefore negligible compared to the T3 circuit mileage propagation delay.

It turns out that the round trip delay between the Washington POP and the San Francisco POP can be 77ms for packets traversing the southern route (Washington->Greensboro->Houston->Los Angeles->San Francisco) or 67ms for packets traversing the northern route (Washington- >New York->Cleveland->Chicago->San Francisco).

During the timeframe of "Hawthorne" technology routers, it was appropriate to choose internal routing metrics that balanced load across redundant T3 paths, and minimized transit traffic on the routers. However now with RS960 technology, the requirement for load-balancing, minimizing transit traffic and hop count, and maintaining equal cost paths is no longer justified. With the introduction of the new Atlanta CNSS, we will explore adjustment of the internal T3 link metrics to minmize round-trip latency ENSS<->ENSS. This will improve overall network performance as perceived by end users. The summary on T3 network latency is:

(1) Delays due to multiple hops in the T3 network are measurable, but not large enough to matter a whole lot. The observed T3 ANSNET one way delay associated with a single T3 router hop is 0.165mS per router (1.35mS cross country one way delay due to 8 router hops). This is neglible compared with the cross-country propogation delays (e.g. 35ms one way). It would require the addition of 30 T3 routers to a path to add 10 ms to the unloaded round trip time, given constant circuit mileage. Delays introduced by extra router hops are negligible compared to circuit mileage delays.

(2) For small packets, like the default for ping and traceroute, the round trip delay is mostly dependent on circuit mileage, and is relatively independent of bandwidth (for T1 and beyond, at least).

(3) All [illegible] network are maintained at equal cost ..... metrics regardless of physical mileage. This was designed during the timeframe when RS/6000 routers were switching packets through the system processor, and hop count, and transit traffic through the router were important quantities to minimize. With the introduction of pure adapter level switching (e.g. no RS/6000 system processor involved in switching user datagrams), minimizing hop count and router transit traffic become less important. Minimizing overall ENSS<->ENSS delay becomes more important.

(4) The T3 ANSNET maintains two different physical circuit routes between Washington D.C. and Palo Alto. Each of these routes represent equal cost paths, and therefore will split the traffic load between them. However one of these physical routes is about 600 miles longer than the other. This can introduce problems involving asymmetric routes internal to the T3 network, and sub-optimal latency. The T3 ANSNET circuits are physically diverse to avoid large scale network failures in the event of a fiber cut. Compromising physical route diversity is not planned. However some reduction of real T3 circuit mileage (and therefore about 5mS of delay) might be possible on the ANSNET with the installation of the Atlanta POP CNSS in September. ANS is conducting a review with MCI to determine whether the Washington->Greensboro->Houston->Los Angeles->Hayward physical route can be reduced in total circuit miles without compromising route network diversity. This might be possible as part of the plan to co-locate equipment within Atlanta.

T3 Network Throughput

The RS960 adapter technology will support up to five T3 interfaces per router, with an individual T3 interface operating at switching rates in excess of 10K packets per second in each direction. The unit and system tests performed prior to the April '92 network upgrade required the CNSS routers to operate at 50KPPS+ aggregate switching rates, and 22Mbps+ in each direction with an average packet size of 270 bytes on a particular RS960 interface. The router has also been configured and tested in the lab to saturate a full 45Mbps T3 link.

The performance that is currently observed by individual end users on the T3 network is largely determined by their access to the network. Access may be via an ethernet or an FDDI local area network. Many users have reported peak throughput observations up to 10Mbps across the T3 network using ethernet access. Several of the T3 network attachments support an FDDI local area network interface which, unfortunately, does not result in more than 14Mbps peak throughput across the T3 backbone right now. With the new RS960 FDDI adapter to be introduced in September, end-to-end network throughput may exceed 22Mbps in each direction (limited by the T3 adapter). The initial RS960 FDDI card software will support a 4000 byte MTU that will be increased later on with subsequent performance tuning. Further performance enhancements will be administered to the T3 backbone in the fall and winter to further approach peak 45Mbps switching rates for end-user applications.

From: mak DEC 21 '92 08:54AM ICC
Subject: August Backbone Engineering Report
Date: 11 Sep 92 03:07:16 GMT

ANSNET/NSFNET Backbone Engineering Report
August 1992

Jordan Becker, ANS becker@ans.net
Mark Knopper, Merit mak@merit.edu

T3 Backbone Status

The system software and routing software for the T3 routers has stabilized. The new RS/960 FDDI card has completed testing and deployment schedules are in progress. A new system software build with support for 10,000 routes maintained locally on the smart-card interfaces is being tested on the T3 Research Network.

Planning is now underway for dismantling of the T1 backbone which is targeted for November. Several steps to be completed prior to dismantling the T1 backbone include support for OSI CLNP transport over the T3 backbone, and the deployment of the redundant backup circuit plan for the T3 ENSS gateways at each regional network.

Further activities in support of the Phase IV upgrade to the T3 backbone are in progress.

Backbone Traffic and Routing Statistics

The total inbound packet count for the T1 network during August was 3,903,906,145, down 17.9% from July. 298,961,253 of these packets entered from the T3 network.

The total inbound packet count for the T3 network was 13,051,979,670, up 1.3% from July. 129,835,094 of these packets entered from the T1 network.

The combined total inbound packet count for the T1 and T3 networks (less cross network traffic) was 16,527,089,468 down 3.1% from July. Reports on T3 backbone byte counts for June, July and August were incorrect due to SNMP reporting problems. These will be corrected soon in reports available on the nis.nsf.net machine. The totals for June, July, and August are 2.279, 2.546, and 2.548 trillion, respectively.

As of August 31, the number of networks configured in the NSFNET Policy Routing Database was 6360 for the T1 backbone, and 5594 for the T3 backbone. Of these, 1233 networks were never announced to the T1 backbone and 1102 were never announced to the T3 backbone. For the T1, the maximum number of networks announced to the backbone during the month (from samples collected every 15 minutes) was 4866; on the T3 the maximum number of announced networks was 4206. Average announced networks on 8/31 were 4817 to T1. and 4161 to T3.

New FDDI Interface Adapter for ENSS Nodes

We have a new RS960 FDDI adapter for the RS/6000 router that provides much improved reliability and performance. It was our hope that the new RS960 FDDI interface adapter targeted to upgrade the older 'Hawthorne' technology FDDI adapters in the the T3 ENSS routers would be ready for deployment in early August. However several serious bugs were encountered during testing in late July, and the upgrade has been delayed by more than a month.

Fortunately we have corrected or worked around all of these known remaining bugs. We are re-running our full suite of regression tests, and a full set of stress tests on the T3 test network during the labor day weekend. Pending the successful completion of our tests, we expect that the first set of FDDI adapter upgrades on the production T3 ENSS nodes could begining during the week of 9/7. We would like to

begin placation of these new interface adapters at ENSS128 (Palo Alto), ENSS135 (San Diego), ENSS129 (Champaign), and ENSS132 (Pittsburgh). We will develop plans for any further FDDI deployments after these first 4 installations have been successfully completed.

Dismantling the T1 Backbone

The current target for dismantling the T1 backbone is November '92. This can be accomplished once the remaining networks using the T1 backbone have been cut over to the T3 backbone (these are: ESnet, EASInet, Mexican Networks at Boulder, and CA*net); an OSI CLNP transport capability over the T3 backbone is in place; the T3 ENSS nodes are backed up by additional T1 circuits terminating at alternate backbone POPs; and the network-to-network source/destination pair statistics matrix is available on the T3 backbone. These activities are described below. Since the RCP nodes on the T1 backbone are experiencing further congestion and performance problems due to the growth in networks, we are planning to reduce the number of networks announced to the T1 nodes by the T3 interconnect gateways. This will eliminate the use of the T3 to back up the T1 for those networks yet to cut over, in the event of a failure in the T1 network.

Remaining Network Cutovers

The ESnet cutover is waiting for a new version of software to be configured for the ESnet router peers at FIX-West and FIX-East. The Mexican autonomous system will be cut over soon, pending communication with the folks in Mexico. We are developing a plan that will allow EASInet to peer directly with the T3 network. The plan for CA*net is to remove the RT from the token ring on the NSS nodes at Seattle, Princeton and Ithaca, configure them to run the CA*net kernel and gated, and peer directly across the ethernet to the T3 ENSS at these sites.

OSI Support Plan

In order to dismantle the T1 backbone, we need to support the transport of OSI (CLNP) packets across the T3 network. Because we would like to target dismantling of the T1 backbone for sometime in late 1992 and the T3 backbone software for support of OSI is still in test, we would like to proceed with a phased (multi-step) migration for support of OSI switching over the T3 network in order to ensure network stability as we introduce OSI software support. The migration plan involves several steps:

1. Convert RT/PC EPSP routers that reside on the shared ENSS LAN into OSI packet encapsulators. This would be done at the 8 or so sites where there are regionals that currently support OSI switching services. OSI traffic is encapsulated in an IP packet on the RT router and forwarded as an IP packet across the T3 network to a destination RT de-encapsulator. This software already exists and can support the migration of OSI traffic off of the T1 backbone, with no software changes required to the T3 backbone. This software is entering test now and could be running in production by early October.

2. Introduce new RS/6000 OSI encapsulator systems that are the running AIX 3.2 operating system with native CLNP support. These machines will replace the RT OSI encapsulators on the shared ENSS LAN. As the CLNP software gets more stable, the RS/6000 system can begin to support non-encapsulated dynamic OSI routing. There are still no changes required to the production T3 network software in this step. This step could occur sometime in the mid-fall.

3. Deploy the AIX 3.2 operating system and native CLNP switching software on the T3 routers across the backbone. The experience gained in step#2 above will facilitate this migration. This step is expected sometime in January 1993.

T1 ENSS Backup Circuits

The T1 backbone is currently providing backup connectivity ir

the event of a problem with the T3 backbone. Since the T3 ENSS nodes are currently singly-connected to a CNSS at an MCI POP, the single T3 circuit and CNSS node represent a single point of failure. As a backup plan, each T3 ENSS will be connected to a new T1 circuit which terminates at a different backbone POP CNSS. This will allow bypass recovery in the event of circuit or CNSS failure. We are executing a test plan on the test network to measure internal routing convergence times and end-user observations during a backup transition. These circuits are being ordered now and are expected to be in place by late October.

Network Source/Destination Statistics

During the migration to the smart card forwarding technology (RS960/T960) we temporarily lost the ability to collect network source/destination pair traffic statistics. This is because packets were no longer passing through the RS/6000 system processor where the statistics collection application software ran. We are now testing new software for near-term deployment that will allow us to continue to collect statistics for each network source/destination pair. These statistics include packets_in, packets_out, bytes_in, and bytes_out. The statistics will be cached on the RS960 and T960 interfaces and uploaded to the RS/6000 system for processing and transmission to a central collection machine.

Increase Routing Table Sizes on T3 Network

We continue to experience an increase in ANSNET/NSFNET advertised networks, (see Backbone Traffic and Routing Statistics, above) The current on-card routing table size on the T3 router RS960 card (T3/FDDI) and T960 card (T1/ethernet) supports 6,000 destination networks with up to 4 alternate routes per destination. The current on-card routing tables are managing on the order of 12K routes (including alternate routes to the same destination).

We are now testing new software for the RS960 and T960 interfaces that will be deployed shortly that supports up to 10,000 destination networks with up to 4 alternate routes per destination. This software will be deployed on the T3 network in the near future.

We also continue to work on support for on-card route caching which will significantly increase the upper limit on the number of routes to unique destination networks that may be supported. This software will be available with the AIX 3.2 operating system release of the router software in early 1Q93.

Phase-IV T3 Network Upgrade Status

The scheduled upgrades to the T3 backbone discussed in the July report are continuing on schedule and will allow the dismantling of the T1 backbone. The major features of this plan include:

1) T3 ENSS FDDI interface upgrades to new RS/960 card. This is currently being scheduled at 4 regional sites.

2) T3 ENSS backup connections are being installed. A T1 circuit will be installed at each T3 ENSS to allow a backup connection to a different CNSS. This will provide some redundancy in the case of T3 circuit or primary CNSS failure. These circuits are scheduled for cutin in October.

3) T3 DSU PROM upgrades. A problem was uncovered in testing the new DSU firmware. The new firmware supports additional SNMP function and fixes a few non-critical bugs. Since this problem was uncovered, a fix has been provided. However the testnet has been occupied with FDDI and other system testing since then. Therefore the upgrades to the DSUs that were scheduled to begin on 9/14 will be postponed until early October.

4) The existing set of CNSS routers in the Washington D.C. area will be moved to an MCI POP in downtown Washington D.C. on 9/12 for

UUNET ~~MFS~~ DATANET

~~UUNET National Network Pricing - 1.544 Mbps~~

UUNET 10 mbps Backbone
Fully Operational Summer, 1993

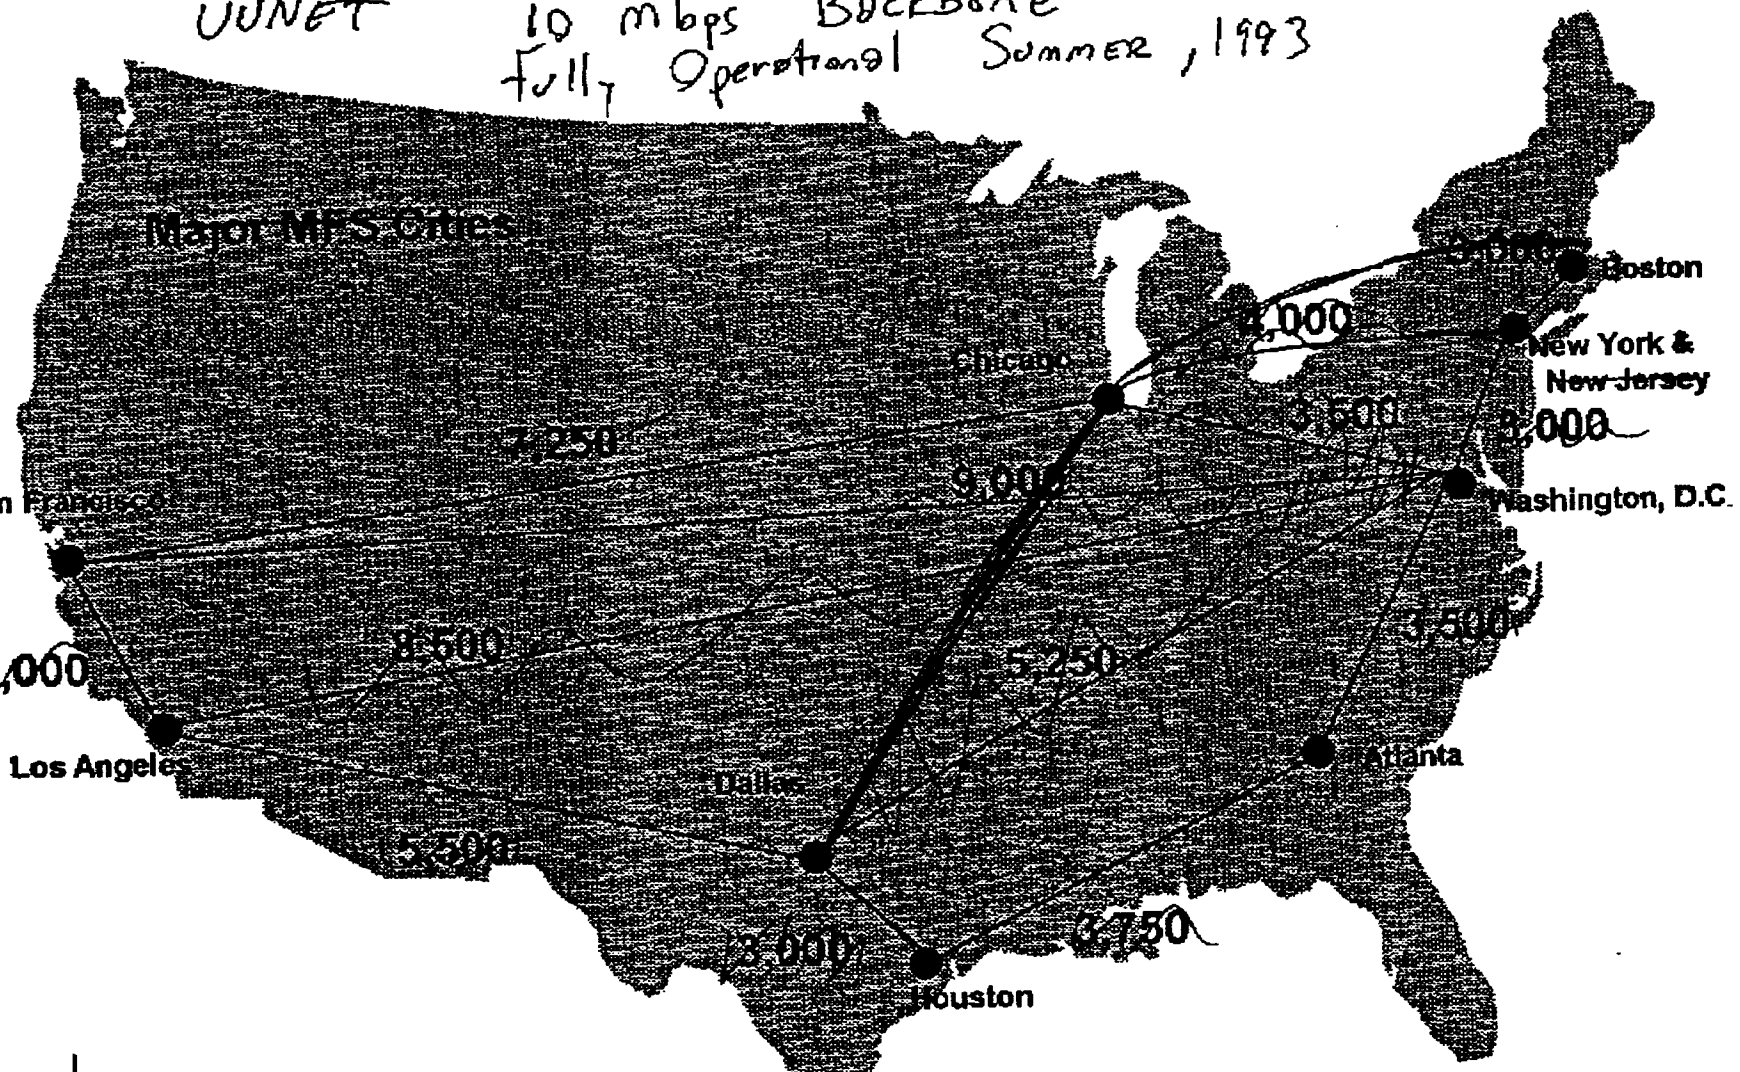

* All paths shown are Point-to-Point ATM Virtual Circuits, not physical routes.

* Add $900.00 Per Month for each Local Loop drop point in a city.

CARRIER: MFS DATANET