

MEMORANDUM

To: F. Heart, D. Walden
From: A. McKenzie
Subject: Report of Major Network Failure on March 22
Date: 23 March 1973

On March 22 at about 7 p.m. eastern time, the ARPA Network began to experience serious difficulties. From the point of view of the NCC operators, these difficulties were being indicated by many traps throughout the network. Machines began going down and there were also some user complaints. The night operator called several members of the software staff so that by 11 p.m. the following crew was on hand working on the difficulty:

Walden
McKenzie
Levin
Crowther

Also, John McQuillan, Bernie Cosell and Steve Butterfield had been contacted and it was arranged that they would come in early the next morning. At the time of the failures, we had advance warning that the Cambridge Electric Company was going to turn off the power to BB&N from about 2:30 a.m. until about 6:30 a.m. on Friday, March 23, so that repairs to the network would have to be made before 2:30. By about 2:30 most of the network was up and running, however, Lincoln was down with hardware problems which had been reported earlier on Thursday and the following sites were isolated from the remainder of the network:

UCSD
Rand
Both Ames Machines
Hawaii
Stanford

Memorandum
Page 2
23 March 1973

It should be noted that prior to 2:30 we had attempted to connect from SRI to the Ames machine and had immediately lost SRI. We therefore suspected that the difficulty with the network was somewhere in the two Ames machines or at Stanford. It therefore seemed prudent to leave these two sites isolated from the network. The UCSD and Rand link was also isolated, but this isolation was due to the lack of time to work on it before the scheduled Cambridge power shutdown. The power was restored shortly after 5 a.m. and by 6 a.m. Steve Butterfield was in the Network Control Center assisting the operator in attempting to diagnose the faults remaining and bring up the remaining sites. This was not accomplished until about noon on Friday the 23rd, but at that time all sites were operating normally.

The cause of the difficulty appears to be a single bit which was lost in either the Ames IMP or the Stanford IMP. The bit that was dropped was dropped from the OP code-field of one of the instructions which generates IMP-to-IMP routing messages. As a consequence, incorrect routing messages were being generated and propagated through the network (with correct checksum). This had the effect of an IMP which was sending bad routing information poisoning its neighbors and it was for this reason that it took a long time to bring the network back to an operational state. Once again, this is a case of a random hardware failure having a serious deleterious effect on the entire network. During the night and morning the software staff has had a chance to consider possible remedies to this type of situation only briefly. There appear to be two remedies, one is to checksum the code which generates the routing messages and the other is to incorporate new software which performs consistency checks on the routing tables either as or after they are generated. This should be able to catch not only bits dropped in the memory, but also bits dropped in the processor, the index register, etc., and is the

Memorandum
Page 3
23 March 1973

solution currently favored by the software staff although more investigation is obviously required.

I telephoned ARPA at approximately 1:30 p.m. on Friday the 23rd and finding both Roberts and Crocker out of the office, reported the above facts to Bruce Dolan.

AAM/jm