Barker

MEMORANDUM

To: IMP Guys

From: McQ

Subject: The Recent Network Crisis IG # 7

Date: December 22, 1972

The last few days have seen a network crisis of large proportions. The problems started on Monday Dec. 18 when we had several difficulties which we labelled as store-and-forward lockups. They continued into Tuesday when we had great difficulty getting out the release of IMP 2725 and experienced new lockup conditions. Finally, we had more trouble on Wednesday with several different sites, which turned out to be unrelated hardware problems.

We have known for a long time that the IMP program has a design weakness: retransmissions are not rerouted. That is, when we choose an output line for a packet and it fails to get through the first time, the IMP program continues to retransmit the packet on the same line. This strategy is best for most cases (phone line errors, momentary storage congestion, etc.), but it is not a good strategy when the receiving IMP, for one reason or another, is never going to accept the packet. This retransmitted packet should be rerouted when it appears to be necessary. We knew that this deficiency could, under unusual circumstances, lead to a global lockup because all IMPs would get filled with packets going into congestion and never route them away from the congestion.

When we had problems on Monday, it looked like a store-and-forward lockup. We succeeded in releasing IMP 2725 on Tuesday, but the problem was still there. Bernie suggested that there might be a "bad IMP" and it began to look like GWC was poisoning the network. It was not until Wednesday that Joel and Steve finally isolated the problem. GWC was getting modem inputs on both its lines with spurious DMC pointers. That is, the

IMP program read the memory channel pointers to determine the length of each new input packet and sometimes received erroneous data. Then it attempted to forward these packets to its neighbors. Sometimes the packets were too long, longer than the maximum length the other IMP would accept, and so they were retransmitted forever. Slowly, GWC filled up with "bad packets", all the while continuing to exchange routing information with its neighbors. Eventually, GWC was full (all 8 output slots on both lines contained bad packets) and the network congested quickly and irreversibly soon thereafter.

To the best of my knowledge no version of the IMP program, past, present, or contemplated, has ever been protected from this kind of error. Further, it is a very difficult problem to protect a network from malfunctioning hardware or software at a particular site or sites.

The other problems that persisted from Tuesday to Wednesday and caused the network to be partitioned were a collection of hardware problems. UCLA had power troubles, blew a fuse and burned out some other gear. The AMES IMP had a flakey memory which was fixed by adjusting the 24 volt power supply. The FNWC machine never showed a sign of life. These and other fires made it seem like the GWC problem was more pervasive or damaging than it actually was.

Finally, it is worth describing our proposed solutions to the general problem pointed up by the GWC experience. If a packet is retransmitted more than 100 times, it will be discarded, on the assumption of a hardware or software failure at one IMP or the other. The line will be held dead for long enough to allow both IMPs to resynchronize things and reroute any other packets on the line. A special trouble report will be sent to the NCC. Notice that we will discard the packet that is not accepted rather than rerouting it. This action is necessary for the GWC problem because that packet is not acceptable by any IMP over any line. This is the first instance of an IMP knowingly deciding to discard a packet which did not have a hardware error indication. It seems such precautions are

necessary to preserve the integrity of the network in the face of unusual malfunctions.

McQ/ph