ESnet Connectivity To Germany
G. Chartrand
3/3/94

ESnet Connectivity To Germany
G. Chartrand
3/3/94

INTRODUCTION

Several years ago, ESnet installed a leased line into Germany to address requirements identified by several programs in ER. The two most prominent programs, HEP and OFE requirements were driven by the activities at DESY for HEP, and ITER for OFE. ESnet explored its options and decided that working with DFN in Germany was the most rational and cost effective solution at the time. In fact, DFN was the only entity that was interested in sharing the costs of a Germany-to-USA circuit. ESnet explored other solutions, but because of DFN's dominance, bypassing them was not practical. Bypassing was desirable however because DFN uses X.25 as its primary network fabric, and it was obvious to us based upon our own experiences that linking sites together with X.25 was not very effective.

The line was installed and shared with DFN. It was agreed to allow DFN's national infrastructure traffic to run on the line and be directed to NSFnet once the traffic reached the USA. The link performed reasonably well for quite a while, but performance issues dictated upgrades to the transatlantic line several times up to its current rate of 512 kbs. In spite of the upgrades, performance issues have persisted for a long time, particularly to DESY. HEP has been complaining about performance to the ESSC for about two years.

CURRENT STATUS

The line to DFN has been saturated for some time. In fact, the last upgrade in speed of the line only provided a performance improvement for a very short period of time. The line saturated after a number of days, and remains so today. The performance degradation is now best measured in percentage of packet loss. The range of loss appears to run between 10% and 40% depending on the time of day *in Germany* [1]. Packet loss is a critical measure because it spawns re-transmissions which in turn generates more network traffic in the process, thus creating a self destructive situation.

This problem is reported to be so bad that the measured throughput drops to about 300 bps and response times exceed several seconds for most of the of the day. *Given this reported level of performance, one can do far better with telephone modem connections!*

PROBLEM ANALYSIS

The problems associated with the connection to Germany is multifaceted. In the very first case, this line is shared with DFN to provide those on the German national and educational network (WIN) access to the USA Internet. Although difficult to measure, the WIN portion of the shared traffic has grown in percentage of the total load on the transatlantic link from about parity (50%) when initially installed, to something over 90% today. Because the performance is so poor, one can assume that many of the DOE users probably have given up on using the link, and thus the percentage of DOE usage during peak hours is probably less than 5%.

It appears that the problem is not necessarily directly related to the bandwidth of the transatlantic circuit between the USA and Germany. Figure 1.0 shows the "ping" response time for the various network routes between a node in the USA and a ITER node in Garshing Germany. The graphs provided *represent* the actual situation which exists. The "ping" and "showpath" utilities are crude tools to use for analysis of complex problems like this. However,

[1] Limited measurements made during the week of 2/20/94

many hours were spent collecting data and averaging results to assure the data represented the actual situation being observed.

As seen in Figure 1.0, there is an obvious difference in the latency distributions between day and night *in Germany*. If only the night time distribution is considered, one would be led to believe that the majority of the problem is associated with the transatlantic link. The day time distribution is entirely different. During the day, the latency associated with the WIN network between the first German router and the node at ITER in Garshing *exceeds* that of the transatlantic link. This obviously is due to the day time load of the intervening WIN network. The day to night differences to DESY appears not to be as drastic (Figure 2.0), however, the increased latency during the day is clearly obvious.

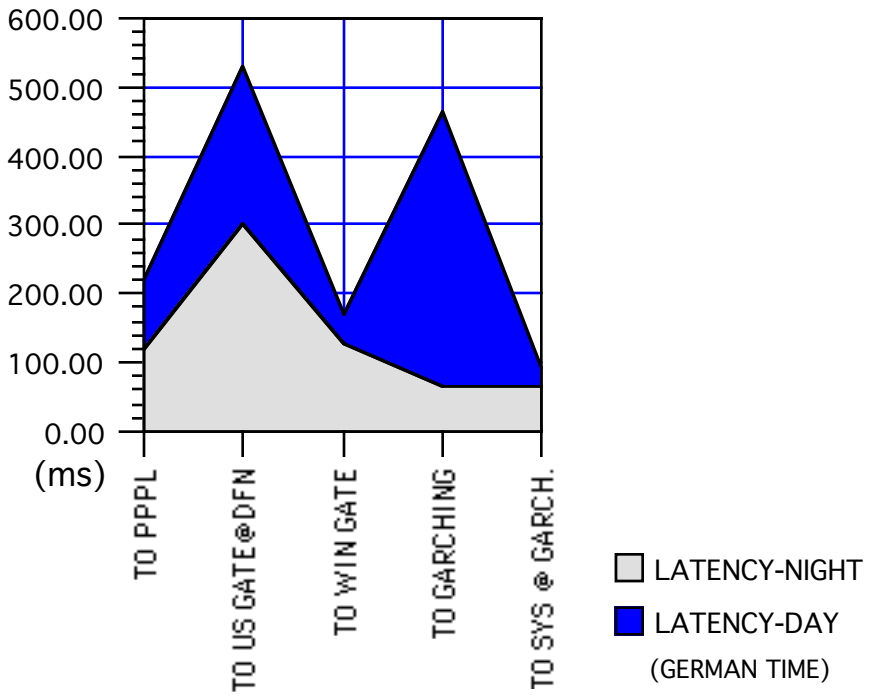

Figure 1.0 Analysis of Garshing Latency

Due to the pricing monopoly of the German PTT, the half circuit sharing "agreement" with DFN forces them to pay about 70 to 80% of the total line costs. Given that disparity in costs, ESnet should be receiving 20 to 30% of the circuit throughput. However, current analysis indicates that our share is more like 1% to 10% depending on the time of day in Germany which related to the domestic traffic load on the WIN network. Additionally, it is safe to assume that our desired use of the network connectivity to Germany is when it is most busy, i.e. during normal business hours at both ends of the circuit

The conclusions one can draw from this analysis is that DOE is no longer receiving its fair quantum from this sharing arrangement with DFN. Furthermore, increasing the bandwidth of the transatlantic circuit will not result in a corresponding increase performance to the locations of interest to DOE. In fact, *it is highly probable that the actual percentage of DOE utilization will decrease by increasing the circuit bandwidth.*

RECOMMENDATIONS

Under the circumstances, ESnet should not increase the bandwidth of the German circuit with the current situation. Doing so would mostly provide additional throughput for WIN traffic to the USA, and DOE would not receive proportional benefit from the upgrade.

The analysis and reported performance also suggests that both the DOE DESY and ITER activities could be better served by having dedicated 64 kbs circuits going to them thus bypassing the WIN network. The process of procuring dedicated circuits should be initiated as soon as possible. If other solutions prevail, orders can always be canceled or put on hold, but they can never be made retroactive. This alternative may not result in increased costs to ESnet compared to the current and proposed upgrade obligations.

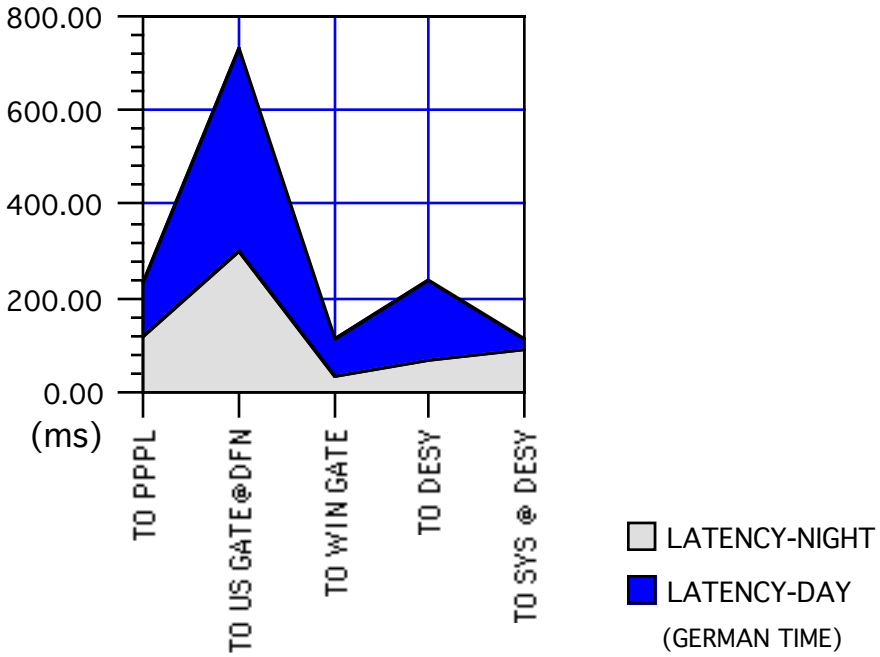

Figure 2.0 Analysis of DESY Latency

Discussions with DFN should be initiated immediately to address the latency across WIN to the locations of interest to DOE. DFN should be given a set time period to fix the problems (two months for example) and if they are unable to reduce the latency to a tollerable level on a 24 hour basis to our sites of interest, we should proceed with implementing dedicated circuits to our two sites and terminate the shared circuit with DFN. It is understood that DFN plans to upgrade the WIN network fabric from X.25 to ATM. At that time, WIN's network infastructure should no longer be an issue. However, our current requirements force us to seek effective alternatives for our needs. *Solutions other than multiple leased lines might be acceptable, but only if they result in adequate network throughput to the sites of interest related to our DOE program requirements.*

Further analysis is needed to determine more details of the problem. For example, it would be interesting to discover which sites within Germany receive better connectivity than those of interest to us. However, any additional analysis should not result in delaying actions to resolve the problem and the ordering of dedicated circuits. It is understood that having dedicated circuits to these two locations is contrary to the traditional ESnet principals of line sharing for

multi-program use, and provision by individual programs for specific requirements. However, under the current circumstances, consideration should be given for non-traditional solutions to address the immediate situation.

CONCLUSIONS

The ITER site in Germany is very important to the OFE program. Much of the activity associated with the program requires login's to systems at both ends of the link. This is the most network demanding application, and the one that suffers most from network "overload".

DESY is obviously an important site for the HEP community, but is also key in providing some of ESnet's connectivity to Russia. DESY supports a 256 kbs link to Moscow State University (MSU) which now has a functional 2 Mbs link to ITEP. Both MSU and ITEP will soon be on the Moscow-wide Fiber infrastructure which will begin operation this summer. In past discussions, DESY has agreed to allow us to use their Russian connectivity. This clearly is a key factor in addressing the multi-program Russian requirements of our DOE research programs.

The current ESnet connectivity to DOE sites of interest in Germany is unacceptable. Definitive actions are necessary to fix this situation as soon as possible.