INTERNET-DRAFT
Louis Breit
Expires: July 22,1997
January 22, 1997

Network Tuning for Efficiency and Throughput
<draft-breit-ntwrk-tuning-00.txt>

Status of this Memo

Abstract

Network tuning is usually performed after the network design is completed, but should also be done at intervals during the life cycle of the network. There are four basic areas that directly affect the efficiency of the network and its associated throughput:

Frame and packet size optimization
Segmentation avoidance or limitation
Minimization of device delay
Window sizing to avoid degradation

This draft has been written to document for the internet community a basic overview of network tuning and its benefits.

Acknowledgments

Much of this memo is taken from my personal experience in designing and analyzing local and wide area networks.

I would also like to extend my appreciation to my colleagues at work for their time and assistance; without their daily contributions, none of the knowledge we have gathered together could be applied as effectively as it has been.

Basic Network Tuning

1. Frame and Packet Size

In all packet switched networks, packet or frame size has a direct effect on overhead and throughput. Smaller packets (each containing a small amount of data) have more overhead than larger ones, since the ratio of overhead (header and control information, etc. - which tend to be fixed in size) is proportionally higher to the amount of data passed. As a result, the data throughput of the line decreases (throughput being defined as the quantity of data a user can pass across a given circuit or device). However, small packets have a reduced chance of retransmission, better response time, and are less likely to contain errors.

Larger packet sizes have a better ratio of overhead to data, which increases throughput. However, buffer and transmission delays and the resulting retransmissions can act to degrade throughput.

This leads us to the "U" curve, where throughput increases as packet size increases but then begins to degrade. The ideal packet size is at the crest of the inverted "U". For example, an IP network suffers from low throughput with very small packet sizes, but will also degrade when packet sizes are too large (due to segmentation). The goal is to increase packet size to just below the point where segmentation is needed, and increase buffers accordingly to compensate for transmission delays.

This will vary by packet service; in networks where each node must read an entire frame before transmission (frame relay, etc.), very large packets should be avoided. As always, the best approach is to "try and tune"; my suggestion is to begin with a packet size of 1500 bytes and assess performance as it is increased and decreased.

2. Segmentation

Segmentation should be avoided at all points in the network to increase throughput. Keep in mind that when IP addressing is used on higher level application protocols, file sizes can exceed the maximum data unit size permitted by IP. Fragmentation or the cutting of single files into multiple packets results in additional overhead. Ethernet, with an MTU of approx. 1500 bytes is similar to IP. However, applications which routinely generate large packets (NFS, etc.) can create delays. If possible, set higher level protocols to limit their packet size to what is supported at the network (IP) layer.

3. Delay

When vendors rate the PPS (or packets per second) supported by a specific device, they often do so using the minimum packet size to increase throughput. Whenever data is passed at a lower PPS than what is supported on the device, throughput is reduced by the difference. For example, if data is passed at 10,000 PPS through a device which offers a PPS of 20,000 PPS, the amount of overhead is 50%. Larger packet sizes will tend to account for some of the difference, but an attempt should be made to match the rated PPS as closely as possible for the appropriate packet size. If a packet size of 1500 bytes is ideal for the network, then the PPS rating of the device using 1500 byte packets should be the goal.

4. Windowing

Window sizes can be tuned within each node of a packet switched network, or within TCP (at the user access point). The window size can be increased on reliable networks and decreased on error prone lines in a packet switched node (X.25, etc.). TCP will dynamically adjust its window size based on throughput. By examining the trending in window size, network performance problems can be isolated and additional capacity added where needed.

It is important to note that window size should never be manually increased as a means to boost performance until a thorough study is done of hardware buffers, memory etc. Larger windows consume more resources.

Author's Address:

Louis Breit
34 Shadow Lane
Staten Island, NY 10306

Phone: 718-983-8472

EMail: louis_breit@jh.com

Expires: July 22, 1997