Appendix A Reading Performance Charts

Dialogues

This type of chart lists on the left those network stations acting as clients, and on the right those stations acting as servers. The list of servers is expanded to show their active services, described by their port numbers and application protocols.

A pair of thin or thick lines between each client and its server indicates by their thickness the volume of traffic from each station. The colour of a line indicates its error rate according to the legend shown in the top right-hand corner of the chart. A dialogue chart might highlight packet retransmission rate, the frequency of sequence-number gaps or the frequency of selective acknowledgements. Each rate characterises a different aspect of packet loss in the network or stress in a server. When packet-loss aspects are viewed across all the dialogues they help to identify where packets are lost.

If the chart is unable to fit all the relevant clients or services it gives preference to those dialogues with higher error rates or larger volumes. The chart's heading displays the number of relevant, and plotted, clients and services (e.g. 47 / 50 Clients and 84 / 142 Services). The heading also indicates the time range covered by the chart.

Bandwidth Use

The simpler form of traffic-volume bar-chart displays the total volume split by network protocol such as IP, IPX (Novell Netware), LLC (Logical Link Control) and ARP (Address Resolution). Traffic filtered out by NetData when processing capture files is shown as *excluded* (in black).

NetData recognises the loss of packets by the monitoring system as distinct from the loss of packets in the observed network. It labels data lost by the monitoring system as *MISSED* and displays an estimate of the missed volume in red. Missed data is often detected some seconds after it should have been captured. If much data has been missed all transaction volumes and transaction-throughput measurements will of course be understated. When data is missed all transactions in progress on the same connection are tagged as *questionable* and their measurements are excluded from the normal performance statistics. If the volume chart's traffic-type legends do not mention the MISSED category, then no TCP data packet was lost by the monitoring system.

Some analyses split all stations – clients and servers – into two groups determined by their network addresses. The first group defines a *region*, such as all the stations at the end of a link. The second group contains all the stations outside or remote from the defined region. In such cases the top half of the bandwidth chart displays the traffic volumes flowing out of the region (i.e. *outbound*), and the bottom half displays the traffic volumes flowing into the region (i.e. *inbound*). If a region comprises all the stations at the end of a link, then such a chart indicates the separate utilisations of the link's forward and reverse channels.

IP Bandwidth Use

This stacked-bar chart splits the total IP traffic volume in each averaging interval (usually 5 or 10 seconds) in two ways. Above the time axis the bars split the traffic according to its source (*outbound* from identified stations), and below the axis the split is according to destination (*inbound* to identified stations).

If the analysis defines a region comprising stations at the end of a link, this chart can take two different forms. In one variant the chart splits outbound and inbound parts into the stations at only one end of the link, to show, for example, how much traffic is received and transmitted by individual servers in a group.

The second variant is confined to either the outbound or the inbound channel of a link, to show, for example, how much traffic is transmitted by individual workstations and received by individual servers.

Transaction Performance

Response times of individual transactions (usually single network round-trips) are indicated by markers plotted at the time of day when the response started, and at a height proportional to the response time. The chart's title indicates whether the plotted response time is measured from the end of the request to the start of the response (*Start-Response*), to the end of the response (*End-Response*), or from the start of the request to the end of the response (*Overall*).

If a transaction has a horizontal bar it is plotted from the start of the request to the end of the response. If the request message has more than one packet a vertical tick on the bar indicates the time at the end of the request.

Response-Time Distribution

Yellow bars on this chart display the distribution of *server* times (from the end of a request message to the start of its response message). Brown bars plot the distribution of times measured to the end of a response message.

A blue graph plots the cumulative distribution.

A blue vertical dashed line marks the average response time. It may be enclosed by cream and pink rectangles that indicate the width of two confidence intervals. The heights of these rectangles indicate their confidence level (usually 80% and 95%). It can be said, for example, that the true average response time is within the span of the cream rectangle with a probability of 80%.

Transactions in Progress

The number of transactions in progress (TIP) through a server can be regarded as a measure of the length of the server's transaction queue, and any abnormal rise in this measure is usually a sign of stress or a performance problem. Graphs of TIP are sometimes overlaid on Transaction Performance charts, as are graphs of transaction rates (throughput) and traffic volumes.

A Transaction Performance chart cannot plot many more than 100,000 individual transactions, but NetData can search through millions of transactions for any abnormal rise in TIP. It displays on a single *activity-summary* chart the peak value of TIP for successive time intervals. In addition to plotting the peak value of TIP, bars on the chart indicate the average and lowest values in each time interval. Any abnormal rise in TIP, no matter how short-lived, is evident in such a chart.

An activity-summary chart may be overlaid with average response times and transaction rates for each server.

Transaction Mixes

One or more stacked bar charts present the contributions of various types of transactions to the overall transaction rate or machine holding time.

A stack labelled 'TIP' (Transactions in Progress) provides an estimate for the numbers of transactions of each type that might be found in progress at any instant. The heights of the bars in such a stack are therefore proportional to the amount of machine time held by the respective transaction types. Furthermore, provided the transactions don't spend most of their time waiting for locked resources such as database tables, machine holding time provides a useful estimate of resource consumption such as CPU utilisation. The transactions appearing at the top of a TIP stack are usually prime candidates for tuning.

A transaction-mix chart may present transaction contributions as either percentages or in absolute terms, of transactions/second or average numbers of transactions in progress.

Multiple stacks might compare transaction mixes in different servers or at different times.

Some charts split each stacked bar into a solid bar on the left and a pale bar on the right, separated by a black vertical line. The length of the solid bar is proportional to the transaction's average response time.

Packet Timing

The packets of a single TCP connection are plotted on two horizontal bands that represent streams of packets from the sockets at each end of the connection. Client packets appear on the lower band and server packets on the upper band. Each packet is indicated by a marker plotted at its time of day, and at a height within its band that is proportional to its length. The shape and colour of the markers distinguish different types of packets according to the legends on the chart. A grey line links the packet markers in chronological order. On some charts pink lines link Ack packets to the data packets they acknowledge, and curved lines link data packets that transmit the same data.

Yellow and blue bars on the socket bands indicate the server and message-transfer times of transactions that are conveyed by the plotted packets.

Orange and blue-green bars may indicate what portion of a transaction's overall time can be attributed to loop-delay, the time to propagate signals through equipment and transmission media, independent of message length and any congestion in the network. Loop-delay is estimated from the observed minimum time to set up new connections and for data to be acknowledged. Server loop-delay is the round-trip time for packets travelling from the recording point to the server and back. Client loop-delay is the round-trip time for packets travelling from the recording point to the client and back.

Web and other types of transactions often involve the concurrent use of several connections, and some packet-timing charts display the packet activity of many connections. Each connection is represented by its own pair of client- and server-socket bands, and there is a consistent spacing between the corresponding bands of each pair.

Data Sequence, Throughput, Window Size, Round-Trip Times

Advances in LLC2 frame-sequence numbers or TCP data-sequence numbers are indicated by a short vertical strip for each packet, plotted against the time of day at which the packet was transmitted. The height at which a strip is plotted indicates – against the y scale – either its LLC2 (Logical Link Control Type 2) frame sequence number or the span of its TCP data sequence numbers. A strip's colour indicates whether the packet left a sequence gap (due to packet loss or delivery out of order), filled a gap, or retransmitted data.

A blue or grey line under the strips represents the lower edge of the sliding window and indicates the time at which data was acknowledged by the receiving station. Blue and grey bands within the sliding window indicate ranges of sequence numbers specified in selective Acks. A blue band indicates acknowledged data.

The packet strips form a line whose slope indicates the data transfer rate. The transfer rate may also be indicated by a dashed line on the chart.

Short horizontal ticks at the top of each strip help to distinguish different packets. A longer horizontal tick indicates not only the end of a packet, but also the end of an application message block, such as an SSL record. If every burst of transmitted data packets completes a message block, flow is likely to be constrained not by any TCP window, but by a shortage of TCP send-buffer space.

The data- sequence chart may be overlaid with an area graph displaying the space occupied within the transmission window, increasing in steps as data packets are transmitted and decreasing as acknowledgements are received. When calculating window occupancy, and when plotting the lower edge of a sliding window, NetData delays the effect of an acknowledgement by the minimum round-trip time in order to present a picture of the window as seen by the transmitting station.

The data-sequence chart may also be overlaid with measured round-trip times for pairs of related packets. Trip times are indicated by the heights of packet markers, the same markers as those appearing on the packet-timing chart. For example, a black square indicates the round-trip time for a data packet and its acknowledgment. A blue diamond indicates the round-trip time for an Ack packet and the first data packet that fits through the widened window.

Green packet markers refer to packets that fill a sequence gap, and their heights indicate the time taken to fill the gap after the gap was first seen. Green markers sometimes fall into two bands: a band of very small times that indicate variations in path time when successive packets take different paths and arrive out of order. All times in the second band will be greater than the loop-delay, and indicate retransmissions prompted by Ack packets.

Time Summary and Transaction List

This chart typically characterises all the system activity involved in a single user transaction such as rendering a web page. The chart plots transaction bars without packet markers, on separate bands for each transaction, and the bars are overlaid with transaction descriptions to form a transaction list.

Columns on the left of the chart display the lengths of the request and response messages, excluding network headers. For HTTP transactions the length of a response message also excludes the length of the response header to show only the length of the payload, with the result that the length indicator is blank for 304-status and other responses that don’t have a message body.

Above the transaction list is a summary of the transaction’s response-time components, usually with three levels of detail. One or two horizontal bars above the transaction list show the totals for the different types of delays, against both a time and a percentage scale. The bottom bar breaks out the total loop-delay (propagation) to show how many loops were incurred by TCP waiting for acks, and the loops associated with individual requests for service. Both bars identify any delays associated with connection setups and SSL/TLS handshakes.

A table in the top-right corner of the chart summarises time in just four categories, of which one is loop-delay. If clients or servers are to be moved to a different region of the network, this should be the only component to change significantly, and the change will be in proportion to the loop-delays of the old and new network paths.

Web transactions in particular use several connections concurrently, and one connection might be transferring a message while another is waiting for a server’s response. To avoid counting a time period twice in the time summary, NetData chooses a single thread or “critical path” through the parallel activities, and counts time only from that path. The chosen path can be seen in the transaction list by tracing the transaction bars outlined in dark blue.

NetData similarly traces a single thread through the parallel activities during a message transfer to determine the number of times the loop-delay is incurred. The overall message transfer time is a function of the thread's sequence of loop-delays and the transmission times of the few packets between those delays, and excludes the transmission times of all the other packets.

Appendix B Terminology

Transactions

Server Transaction — A single round-trip from a user's workstation (the *client*) to a server, involving paired request and response messages.

User Transaction — All the activity between a user clicking on a button (or similar event) and the eventual rendering of a new display on the user's screen. It may involve a large number of server transactions, with a variety of servers, and those server transactions may run concurrently.

Latency

Transaction response time comprises server processing time, client (workstation) 'reaction' or processing time, and time spent in the network. Network latency has four major components:

Transmission Time — The time to transmit the bits of packets at the clock speed of each link in a network path. The time to transmit a long message is largely determined by the time to transmit the message over the slowest link in the path. Extra hops in a path increase the transmission time, particularly for the first packet of a message.

Network speed usually determines a link's bandwidth.

Network speeds in a modern network are usually so high that transmission times are rarely significant.

Congestion Delay — The packets from many clients and servers must queue for their turn to be transmitted over each network link, and congestion delay represents the accumulated waiting times for packets in a message.

Waiting time is largely a function of the average transmission time and the link utilisation. If links are not overloaded – i.e. utilisations are rarely above 50% – congestion delay should be less than the average transmission time and should therefore be negligible in a modern network.

Loop Delay — Electrical and optical signals take a finite time to propagate through a transmission medium. Propagation time is a fixed characteristic of a network path and the related speed is typically about 70% of the speed of light. Propagation time, equipment reaction times and equipment transit times (electronic and processing delays) are combined in a parameter called loop-delay which represents the minimum round-trip time for a notional message with zero bits (i.e. no transmission time) and in the absence of any other traffic (i.e. no congestion delay).

Protocol Delay — Message flow is sometimes delayed by TCP delayed-ack timeouts and retransmission timeouts.

Analysis

The traffic has been decoded with NetData, a tool that reconstructs, measures, characterises and records in a database every transaction, including those with incomplete responses and requests without responses. If any transaction is slow, or fails, it is displayed prominently in one of NetData's charts.

Besides characterising transactions, NetData measures the round-trip times for all TCP data packets and their acknowledgements, and charts the flow-control decisions made by servers and workstations. These charts display any evidence of network congestion or packet loss, and expose any inappropriate configuration parameters that affect data flow.

The analysis is presented here under *four* main headings: ... *transactions; and miscellaneous signs of unhealthy system behaviour.* The last topic has been approached as a standard Measure IT health check, leaving "no stone unturned" in an attempt to find every problem that might have a bearing on performance.

The main findings are summarised and discussed at the head of the report. The body of the report presents the supporting evidence, largely in the form of annotated NetData charts and tables. The evidence is introduced with a set of dialogue charts to identify the main servers and their services seen in the captured traffic. Notes on reading the various types of charts appear in Appendix A. Appendix B defines terminology for types of transactions and for the components of network latency.

For a high-level overview it may be sufficient to read only ...