COMNET III

Reports & Statistics

Release 2.5.1

COMPUWARE

Reports and Statistics
COMNET III™
Release 2.5.1

Part: CN-1700-0400-25

Compuware Corporation
31440 Northwestern Highway
Framington Hills, MI 48334-2564

Printed in the United States of America.

Contents

Section I

Statistics Monitors

1.1 Overview

COMNET III allows many more variables to be monitored during the course of a simulation than what is reported in the output reports. It is possible to collect more detailed information about the behavior of objects (nodes, links, traffic sources, etc.) in a model by turning on various statistics monitors. Most, but not all objects in COMNET III have statistics monitors. To turn on the statistics monitors for an object, it is necessary to edit the object and click on the **Statistics** button displayed on the various dialogs.

1.1.1 Selecting Statistics Requests

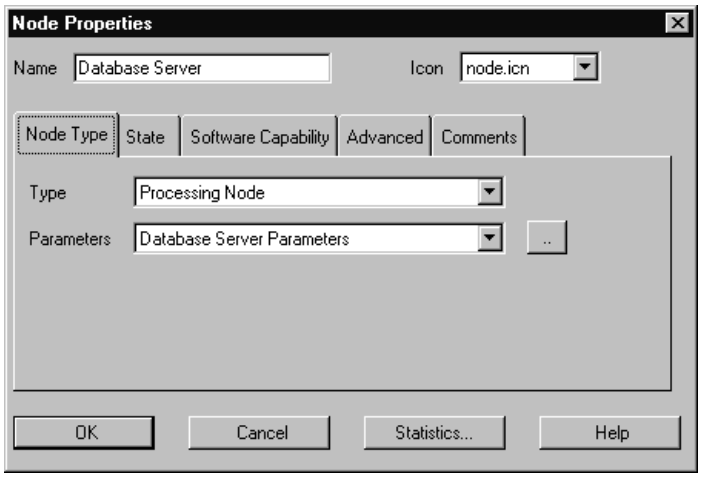

Clicking on the **Statistics** button opens a particular set of **Statistics Requests** that can be turned on to collect information during the simulation.

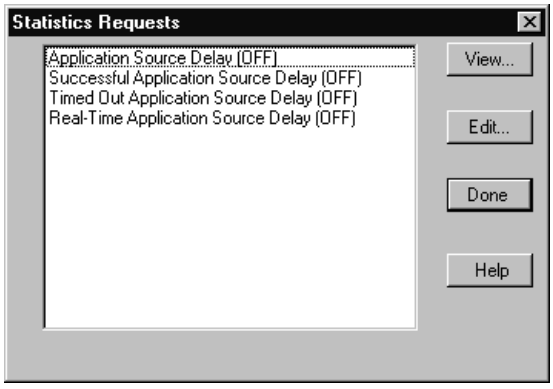

Turning on statistics requests enables the collection of raw observation data. This data can then be plotted, analyzed for percentiles or exported to other statistics tools, such as Excel. If you select a request and click on the **Edit** button, COMNET III presents the following dialog.

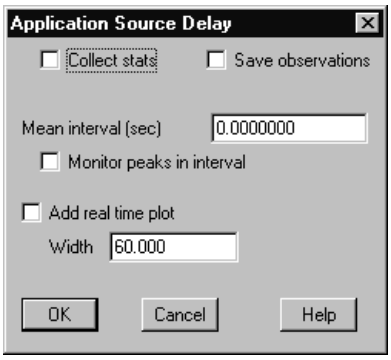

The fields of the dialog are described below.

Note: If the **Collect stats** checkbox is not marked, COMNET will not collect the indicated statistics regardless of the other settings.

Collect stats

Enables the collection of the data required to create the reports for the statistics selected on the **Statistics Requests** dialog. Normal statistics are available at the end of the simulation. Basic statistics such as mean and maximum are collected when the **Collect stats** button is turned on. Normal statistics are generally available for both the cases where the statistics are updated for each event, and for where the statistics are updated with the mean values over an interval.

Save observations

Saves to a file an observation for each monitored event. For example, if the monitor is of message delay, the file will capture the measured delay for each message. In the case of monitoring levels such as link utilization, each change in state is recorded to a file.

Note: This option provides valuable detailed data on each monitor, but use it judiciously. Unlike the **Collect stats** option, this option involves frequent file I/O that can slow down the simulation. There is also the possibility of creating very large and unmanageable observation files, if too many items are collecting observations.

Mean Interval

When this option is zero, the statistics are collected only on the raw measurements, and if observations are collected, the observations are of each individual event or state-change. When the option is nonzero, a separate set of statistics are collected by taking just one measurement each interval where the measurement is representative of the interval. If observations are saved, it is this single observation per interval that is saved.

The representative value for the interval may either be the mean value (the average over the interval) or the peak value (the maximum value over the interval).

This option has two benefits. First, it provides more informative statistics especially for the minimum, maximum, and percentiles, by collecting statistics of observations of short-term results rather than the raw observation. For example, it provides a mechanism to find the minimum or maximum of the short-term average link utilization or the short-term peaks of message delays or VC burst sizes.

The second benefit is that it reduces the size of saved-observation files by saving just one sample per interval rather than the potentially large number of events or state changes within that interval. This allows the simulation to run faster and lessens the size of the observation file.

Monitor peaks in interval

The **Mean Interval** observations recorded and summarized depend on the checkbox **Monitor peaks in interval**. If this checkbox is off, then the observations will be the average of the monitored values during the interval. This observation is especially useful for monitoring utilizations of links or processors.

If the checkbox is on, the observations will represent the peak value of the monitored quantity during the interval. This option is especially useful for monitoring buffer sizes and burst sizes. This option is useful for monitoring buffer sizes where the peak value (no matter how rare) is related to traffic loss.

Add real time plot

Add real time plot. This option presents a plot that comes up during the simulation. This plot is updated by each observation that is collected into the statistics. These observations are governed by the mean-interval and monitor-peak options of the statistics requests. This plot option does not require observations to be saved to a file, but it plots the same information that would be saved to the file.

1.1.2 Viewing Requested Statistics

After running the simulation, use the **Statistics** button, again, to **View** and then **Plot** the selected statistic. After bringing up the plot window by clicking on the **Plot** button, use **Plot** on the menubar to obtain histograms, percentiles, or plots of smoothed data.

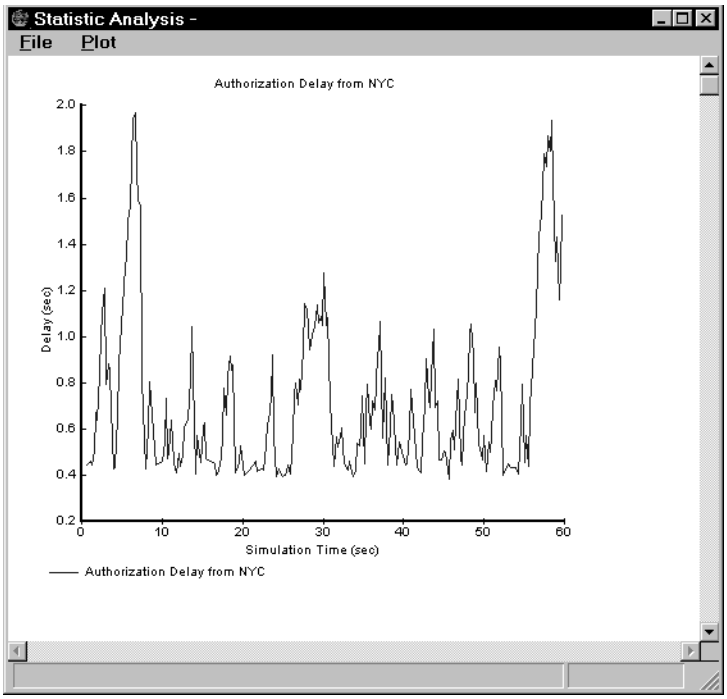

When plots of saved data are viewed after the simulation using the **Statistics** button, it is possible to obtain any desired percentile for the saved data. For example, you could determine the 95th percentile for message delays from a particular message source, provided **Save data** was requested for that source prior to the simulation.

Post-processed plots of saved data allow you to choose the number of points to be plotted, as well as the interval over which each plotted value is to be computed. For example, you could choose to plot 100 observations of utilization over an 80 second interval, with each observation of utilization computed as the time-weighted average utilization over the previous 2 second interval. If you request a plot of smoothed data for a delay type of statistic, you specify the time period [T1, T2] to be plotted and the number of points, N, to be plotted. The time period is divided into (T2 - T1) / N intervals and a single point is plotted for each interval. The point plotted for each interval is equal to the average of all of the delays observed during the interval. If no delays are observed, 0 is plotted.

If you request a plot of smoothed data for a level type of statistic (e.g., channel utilization), you specify the time period [T1, T2] to be plotted, the number of points, N, to be plotted, and the averaging interval, I, to be used in calculating the utilization at each plotted point. The point plotted at time T is the time-weighted average of the level statistic during the time interval [T-I, T]. If I > T, the time-weighted average is computed over the interval [0, T]. The averaging interval must be greater than or equal to (T2 - T1) / N. If the time period to be plotted is an entire replication and the averaging interval is set to the length of the replication, then each plotted point represents the time-weighted average of the level statistic from the beginning of the replication to the time of the plotted point.

To make the real time chart look like the one in the *COMNET III Sample Models Guide*, i.e. adjusted so that the entire time range fits in the chart without scrolling, right-mouse click on the chart. The following window will be displayed.

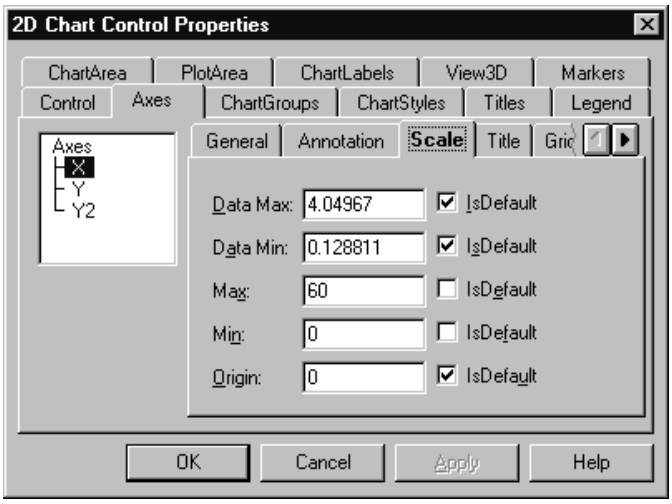

Select the **Axes** tab, select the **Scale** sub-tab, and check all the **Is Default** boxes. Then select the **Annotation** sub-tab, click the three-dot button next to **Annotation Method: Values**, and check all the **Is Default** boxes there.

See the 'Comnet III Chart Controls' help file for more information on manipulating the way charts are displayed.

1.1.3 Exporting Statistics Files

The statistics monitors in COMNET III collect raw observation data as well as data sampled over a specified interval. All of the simulation statistics collected during the run (raw observational data and interval sampled data) can be exported. Use the **Simulate/Run Parameters** option to display the **Run Parameters** tabs.

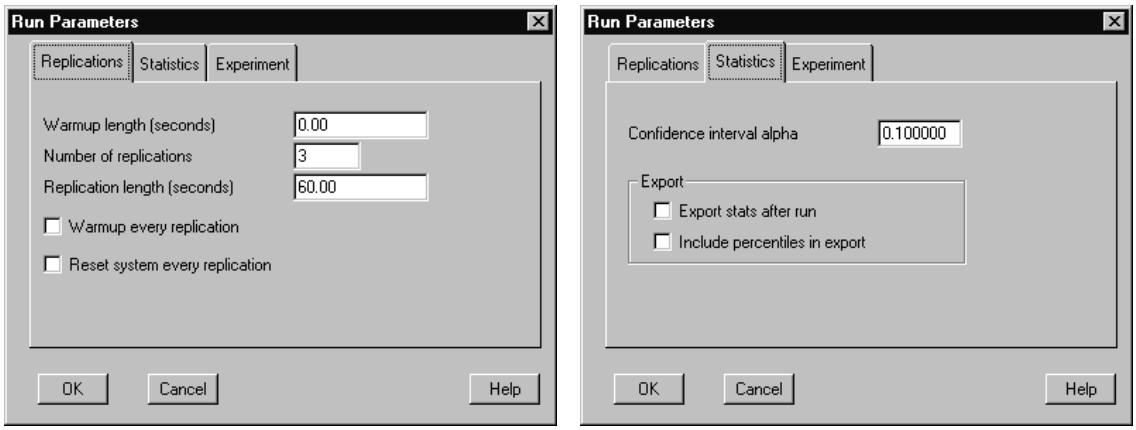

The **Replications** tab has a **Number of replications** field in which you can enter the number of times you wish your model to run. When multiple replications have been set, the desired

confidence interval alpha can be set under the **Statistics** tab. On the same tab, select the **Export stats after run** checkbox to automatically produce an export file of all statistics upon completing the simulation. Select **Include percentiles in export** if you wish percentiles to be included in the exported file. The export file can be imported directly into spreadsheet packages such as Microsoft Excel for further analysis.

Turning on export statistics does nothing to identify what gets exported. Instead it exports what has been requested elsewhere. It exports all monitors that have been turned on explicitly via the **Statistics** requests. It exports all data that is requested from the report requests. The export file contains all the collected statistics in the model except for the actual observations that may be available for plotting.

Once the simulation is completed, click on **File/Export/Simulation Statistics** and specify the file name to export to. The export file is an alternative to the report file and the report includes both monitor information and counter information. The file will have an extension of **`.xpt.`** The statistical information is exported in a tabular format. The first row of the file contains the model name and date of the simulation run. The second row contains the row titles for n following rows where n is the number of replications. These rows provide start and stop times for each replication. The interval sampled data is written to the fields that begin with **`Int.`,** while the raw data is written to the other fields. Each line in the file has the following tab-delimited fields:

Title

Units

Rep	The replication number. When there is more than one replication, replication -1 combines statistics for all replications as though the simulation consisted of one long replication. Replication 0 treats each replication as one observation and provides the statistics based on these summary observations.
Min	The smallest observed value or level.
Max	The largest observed value or level.
Mean	The mean observed value or level.
Std Dev	The standard deviation of the observed values or levels.
Sum	For delay types of performance measures, the sum of the delays. For level statistics (e.g., buffer usage), the levels are time-weighted. To compute the mean for a level statistic, divide the **Sum** by the length of a replication.
Sum of Squares	For delay types of performance measures, the sum of the observed values squared. For level statistics, the levels are time-weighted.
Count	The number of observations.

Lower Limit	The lower limit of the **confidence interval alpha** value you specify for the raw observational data. The **confidence interval alpha** is specified under the **Statistics** tab, which is found under the **Simulate/Run Parameters** menu option. **Lower Limit** values will only show up for simulations running more than one replication.
Upper Limit	The upper limit of the **confidence interval alpha** value you specify. The **confidence interval alpha** is specified under the **Statistics** tab, which is found under the **Simulate/Run Parameters** menu option. **Upper Limit** values will only show up for simulations running more than one replication.
Interval	The **confidence interval alpha** value entered by the user.
Int. Min	The smallest value of the samples from the intervals.
Int. Max	The largest value of the samples from the intervals.
Int. Mean	The mean value of the samples from the intervals.
Int. Std. Dev.	The standard deviation value of the samples from the intervals.
Int. Sum	The sum of values value of the samples from the intervals. The levels are time-weighted. To compute the mean for a level statistic, divide the **Sum** by the length of a replication.
Int. Sum of Squares	The sum-of-squared-values value of the samples from the intervals. The levels are time-weighted.
Int. Count	The count of intervals value of the samples from the intervals.
Int. Lower Limit	The lower limit of the **confidence interval alpha** value you specify for the interval sampled data. The **confidence interval alpha** is specified under the **Statistics** tab, found under the **Simulate/Run Parameters** menu option. **Lower Limit** values will only show up for simulations running more than one replication. The **Lower Limit** values show up on the 0^{th} replication labeled as **Cumulative** in the tables or replication 0 in the export file.
Int. Upper Limit	The upper limit of the **confidence interval alpha** value you specify for the interval sampled data. The **confidence interval alpha** is specified under the **Statistics** tab, which is found under the **Simulate/Run Parameters** menu option. **Upper Limit** values will only show up for simulations running more than one replication. The **Upper Limit** values show up on the 0^{th} replication labeled as **Cumulative** in the tables or replication 0 in the export file.

At the bottom of the monitor table are counter tables for message destination counters, link counters, node counters, other transmission counters, buffer packet counters and other counters.

1.2 Access Link Monitors

The simulation can monitor statistics for each access link's entry and exit utilization, and the statistics for the buffer size at the exit end. The buffer size is the number of frame-level bytes (including frame overhead) that reside in the buffer for both the frames waiting for transmission on the exit link and the frames currently transmitting on the exit link. The link utilization is the ratio of the number of circuits busy transmitting frames divided by the total number of circuits available. Generally there will be just a single circuit, and so the basic utilization will have either the value of zero or one.

The event statistics are labeled as **Exit Buffer Size**, **Entry Link Utilization**, and **Exit Link Utilization**. Observations of these events may be saved for plotting or post-run analysis. However, there is an observation collected for each frame entering and leaving an access link and, thus, there will generally be a very large number of observations to record.

Averaged statistics for the buffer size and link utilization are also available and they are labeled as **Mean Exit Buffer Size**, **Mean Entry Link Utilization**, and **Mean Exit Link Utilization**. These monitors will sample the averaged value at each sampling interval instead of each event, and, thus, there will be fewer observations to plot. However, the averaging interval can no longer be varied as part of the post-run analysis.

The average statistics may also be selected for real-time plots that will update at each update interval during the simulation. These real-time monitors are labeled **R-T Mean Exit Buffer Size**, **R-T Mean Entry Utilization**, and **R-T Mean Exit Utilization**, where **R-T** is an abbreviation for real-time.

Note: The term "real-time" refers to the plots that are available and updating while the simulation progresses. It does not refer to whether the simulation itself is running in real-time.

1.3 Virtual Circuit Monitors

For virtual circuits, the simulation can measure the frame delay and the leaky-bucket burst size.

The frame delay is the time the frame takes to traverse the cloud from the moment it is created at the input to the entry link to the time it is delivered at the output of the exit link. This frame delay counts only the frames that successfully go through the cloud, and, thus, it does not count frames that are dropped within the cloud. The frame delay monitor labeled **Frame Delay** collects statistics for the delay for each frame that is delivered. The **Mean Frame Delay** monitors the averaged frame delay over the mean-statistics interval, and the **Real-Time Mean Frame Delay** provides a real-time plot of the averaged delay for this virtual circuit.

The **Burst Size** is the current size of the leaky-bucket burst at the time a frame is accepted by the virtual circuit and, thus, the frame size is added to the burst. The **Burst Size** contains only the data bytes for the encapsulated packets and, thus, it does not include the frame

overhead or any padding necessary to meet minimum frame size. The monitor labeled **Burst Size** will collect the burst statistics for each frame accepted at the virtual circuit. The monitor labeled **Sampled Burst Size** collects burst size statistics based on a sample at the time of the mean-statistics update interval controlled by the cloud. **Real-Time Sampled Burst Size** provides a real-time plot of the sampled burst size.

Also, if the virtual circuit's **Show Burst Size** switch is on, the sampled burst size for that virtual circuit will be displayed over the virtual circuit icon during the simulation.

1.4 Processing Node Monitors

Used Bytes Level

Measure how much file space is used in the node's disk storage space. This monitor is useful for models that have Read and Write commands that change the size of files.

Busy Processors

Measure the number of busy processors. For a single processor node, the result is the utilization of the processor. Frequently the processor will be busy with short duration tasks and, thus, it is generally desirable to collect this utilization averaged over a convenient interval.

Input Buffer Level

Measure the input buffer level of the node. Since packets arrive and leave buffers very frequently, it is generally desirable to collect averaged buffer levels for plotting.

Output Buffer Level

Measure the output buffer level of the node. Since packets arrive and leave buffers very frequently, it is generally desirable to collect averaged buffer levels for plotting.

Call Bandwidth Level

Measure the call bandwidth allocated through this node. It is useful when modeling bandwidth-switched calls with call sources, and it gives the total bandwidth of all the calls using this node.

Calls In Progress

Measure the number of calls in progress established by the node.

Channels In Use

Measure the number of call channels in use by the node.

1.5 Network Device Monitors

Used Bytes Level

Measure how much file space is used in the network device's disk storage space. This monitor is useful for models that have Read and Write commands that change the size of files.

Busy Processors

Measure the number of busy processors at any time. For a single processor network device, the result is the utilization of the processor. Frequently the processor will be busy with short duration tasks and, thus, it is generally desirable to collect this utilization averaged over a convenient interval.

Input Buffer Level

Measure the input buffer level of the network device. Since packets arrive and leave buffers very frequently, it is generally desirable to collect averaged buffer levels for plotting.

Output Buffer Level

Measure the output buffer level of the network device. Since packets arrive and leave buffers very frequently, it is generally desirable to collect averaged buffer levels for plotting.

Call Bandwidth Level

Measure the call bandwidth allocated through the network device. It is useful when modeling bandwidth-switched calls with call sources, and it gives the total bandwidth of all the calls using the network device.

Calls In Progress

Measure the number of calls in progress established by the network device.

Channels In Use

Measure the number of call channels in use by the network device.

1.6 HOL Blocking Switch Monitors

Input Buffer Level

Measure the input buffer level of the HOL Blocking Switch. Since packets arrive and leave buffers very frequently, it is generally desirable to collect averaged buffer levels for plotting.

Call Bandwidth Level

Measure the call bandwidth allocated through the HOL Blocking Switch. It is useful when modeling bandwidth-switched calls with call sources, and it gives the total bandwidth of all the calls using the HOL Blocking Switch.

1.7 Aloha, CMSA, CMSA/CA, CMSA/CD and C.A.N. Link Monitors

The contention channel is the primary channel for data communications on the link.

Statistics for the control channel transmission delay are only gathered when the control channel has been enabled on the **Link Specifics** tab, and a **Control Node** has been specified. The control channel allows the non control nodes attached to the link to respond to control node messages along a contention free channel.

Contention Channel Utilization

Measure the percent of capacity at which the contention channel of the link is utilized.

Control Channel Utilization

Measure the percent of capacity at which the control channel of the link is utilized.

Contention Channel Transmission Delay

Measure the link layer frame delay of the contention channel.

Control Channel Transmission Delay

Measure the link layer frame delay of the control channel.

Deferral Delay

Measure the delay time of a link layer frame to transmit when the frame detects that the link is busy and *"defers"* transmission until the link is idle. This measure only applies to the contention channel.

Deferral Queue Size

Measure the queue size of link layer frames that have deferred transmission until a busy link becomes idle. This measure only applies to the contention channel.

Multiple Collision Episodes

Measure the number of times that two or more link layer frames collided.

Collisions Per Collided Frame

Measure the number of link layer frames that collided per collision episode.

Contention Channel Frame Size

Measure the size of the link layer frames being transmitted across the contention channel.

Control Channel Frame Size

This monitor measures the size of the link layer frames being transmitted across the control channel.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.8 DAMA Link Monitors

Busy Channels

Measure the number of channels of the link being used.

Transmission Delay

Measure the delay to transmit link layer frames across the link.

Frame Size

Measure the size of the link layer frames transmitted across the link.

Available Calls

Measure the number of call channels available.

Call Utilization

Measure the percent capacity at which the call channels are utilized on the link.

Calls In Progress

Measure the number of calls in progress on the link.

Channels In Use

Measure the number of channels in use on the link.

Session In-Progress

Measure the number of sessions in progress on the link.

1.9 Dialup Point-to-Point Link Monitors

The From X and From Y directions are determined by the order in which the nodes are connected to the Dialup Point-to-Point link. When editing the properties of the link, the X and Y nodes are listed across the top of the link dialog.

Busy Modem

Measure when the Dialup Point-to-Point link (modem) is busy. There is one modem and it can be busy or not busy (connected or disconnected). When a modem is busy, the channels may still be idle (no utilization) and, so the busy channels monitor is for the utilization on the channels when the modem is busy.

Busy Channels From X (To Pool)

Measure the number of channels from Node X to the Dialup Point-to-Point link being used.

Busy Channels From Y (To Pool)

Measure the number of channels from Node Y to the Dialup Point-to-Point link being used.

Time To Connect

Measure the time required to establish a connection before data can be transmitted.

Transmission Delay From X (To Pool)

Measure the delay to transmit link layer frames from Node X across the link.

Transmission Delay From Y (To Pool)

Measure the delay to transmit link layer frames from Node Y across the link.

Frame Size From X (To Pool)

Measure the size of the link layer frames from Node X transmitted across the link.

Frame Size From Y (To Pool)

Measure the size of the link layer frames from Node Y transmitted across the link.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.10 FDDI Basic Link Monitors

Node Active Duration

Measure the amount of time a node spends actively acquiring and using the token to empty its output buffer. Once the output buffer is empty, the node is no longer active. This monitor helps to identify if multiple token rotations are required to empty a node's buffer.

Active Node Count

Measures the number of nodes actively trying to capture the token. This is the number of nodes with non-empty buffers. The larger the number of active nodes, the slower the token rotation will be or the less time a node is allowed to transmit (for Priority Token Ring or FDDI).

Token Rotation Time

Measure the average time for the token to pass to all the nodes connected to the link.

Utilization

Measure the percent capacity at which the link is utilized.

Transmission Delay

Measure the delay to transmit link layer frames across the link.

Frame Size

Measure the size of the link layer frames transmitted across the link.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.11 FDM Link Monitors

Utilization Measure the percent capacity at which the link is utilized.

Transmission Delay

Measure the delay to transmit link layer frames across the link.

Frame Size Measure the size of the link layer frames transmitted across the link.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.12 FDMA Link Monitors

Busy Users

Measure the number of FDMA channels in use at any time.

Channel Used

Measure the number of channels used on the link.

Time To Connect

Measure the time required to establish a connection before data can be transmitted.

Node Activity Duration

Measure the time the nodes spend, both waiting for a channel to become available, and spent using the channel.

Active Node Count

Measure the number of nodes using channels, and the number of nodes waiting to use the channels.

Utilization Uplink

Measure the percent capacity at which the uplink is utilized.

Utilization Downlink

Measure the percent capacity at which the downlink is utilized.

Transmission Delay Uplink

Measure the delay to transmit link layer frames across the uplink.

Transmission Delay Downlink

Measure the delay to transmit link layer frames across the downlink.

Frame Size

Measure the size of the link layer frames being transmitted across the link.

Session In-Progress

Measure the number of sessions in progress on the link.

1.13 Link Group Link Monitors

Host Channel Frame Size

Measure the size of the link layer frames transmitted across the host channel.

Switch channel Frame Size

Measure the size of the link layer frames transmitted across the switch channel.

Sessions In-Progress

Measure the number of sessions in progress on the link.

Host Channel Utilization

Measure the percent capacity at which the host channel of the link is utilized.

Switch Channel Utilization

Measure the percent capacity at which the switch channel of the link is utilized.

Host Channel Transmission Delay

Measure the delay to transmit link layer frames across the host channel of the link.

Switch Channel Transmission Delay

Measure the delay to transmit link layer frames across the switch channel of the link.

Ring Node Activity Duration

Measure the amount of time a node spends actively acquiring and using the token to empty its output buffer. Once the output buffer is empty, the node is no longer active. This monitor helps to identify if multiple token rotations are required to empty a node's buffer. Applies only to token passing and FDDI link types.

Ring Active Node Count

Measure the number of nodes actively trying to capture the token. This is the number of nodes with non-empty buffers. Applies only to token passing and FDDI link types.

Ring Token Rotation Time

Measure the average time for the token to pass to all the nodes connected to the link. Applies only to token passing and FDDI link types.

CSMA/CD Deferral Delay

Measure the delay time of a link layer frame to transmit across CSMA/CD link group link, when the frame detects that the link is busy and *"defers"* transmission until the link is idle. This measure only applies to the contention channel.

CSMA/CD Deferral Queue Size

Measure the queue size of link layer frames that have deferred transmission until a busy CSMA/CD link group link becomes idle. This measure only applies to the contention channel.

CSMA/CD Multiple Collision Episodes

Measure the number of times more than two link layer frames collided on a CSMA/CD link group link.

CSMA/CD Collisions Per Collided Frame

Measure the number of link layer frames that collided per collision episode on a CSMA/CD link group link.

1.14 Modem Pool Monitors

Busy Modems

Measure the number of modems in the modem pool being used.

Busy Channels From X (To Pool)

Measure the number of channels from Node X to the modem pool link being used. The From X direction is determined by the order in which the nodes are connected to the modem pool link. When editing the properties of the link, the X and Y nodes are listed across the top of link dialog.

Busy Channels From Y (To Pool)

Measure the number of channels from Node Y to the modem pool link being used. The From Y direction is determined by the order in which the nodes are connected to the modem pool link. When editing the properties of the link, the X and Y nodes are listed across the top of link dialog.

Time To Connect

Measure the time required to establish a connection before data can be transmitted.

Transmission Delay From X (To Pool)

Measure the delay to transmit link layer frames from Node X to the modem pool link.

Transmission Delay From Y (To Pool)

Measure the delay to transmit link layer frames from Node Y to the modem pool link.

Frame Size From X (To Pool)

Measure the size of the link layer frames transmitted from Node X to the modem pool.

Frame Size From Y (To Pool)

Measure the size of the link layer frames transmitted from Node Y to the modem pool.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.15 Point-to-Point Link Monitors

The From 1st and 2nd Node directions are determined by the order in which the nodes are connected to the modem pool link. When editing the properties of the link, the 1st and 2nd nodes are listed across the top of the link dialog.

Busy Channels From 1st Node

Measure the number of channels of the link being used from the 1st Node.

Busy Channels From 2nd Node

Measure the number of channels of the link being used from the 2nd Node.

Transmission Delay From 1st Node

Measure the delay to transmit link layer frames across the link from the 1st Node.

Transmission Delay From 2nd Node

Measure the delay to transmit link layer frames across the link from the 2nd Node.

Frame Size From 1st Node

Measure the size of the link layer frames transmitted across the link from the 1st Node.

Frame Size From 2nd Node

Measure the size of the link layer frames transmitted across the link from the 2nd Node.

Available Calls

Measure the number of call channels available.

Call Utilization

Measure the percent capacity at which the call channels are utilized on the link.

Channels In Use

Measure the number of channels in use on the link.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.16 Polling Link Monitors

Statistics for the control channel utilization are only gathered when the control channel has been enabled on the **Link Specifics** tab, and a **Control Node** has been specified. The control channel allows the non control nodes attached to the link, to respond to control node messages along a contention free channel.

The contention channel is the primary channel for data communications on the link.

Contention Channel Utilization

Measure the percent of capacity at which the contention channel of the link is utilized.

Control Channel Utilization

Measure the percent of capacity at which the control channel of the link is utilized.

Contention Channel Transmission Delay

Measure the link layer frame delay of the contention channel.

Control Channel Transmission Delay

Measure the link layer frame delay of the control channel.

Contention Channel Frame Size

Measure the size of the link layer frames transmitted across the contention channel.

Control Channel Frame Size

Measure the size of the link layer frames transmitted across the control channel.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.17 Priority FDDI and Token Ring Link Monitors

Node Active Duration

Measure the amount of time a node spends actively acquiring and using the token to empty its output buffer. Once the output buffer is empty, the node is no longer active. This monitor helps to identify if multiple token rotations are required to empty a node's buffer.

Active Node Count

Measure the number of nodes actively trying to capture the token. This is the number of nodes with non-empty buffers.

Token Rotation Time

Measure the average time for the token to pass to all the nodes connected to the link.

Utilization

Measure the percent capacity at which the link is utilized.

Transmission Delay

Measure the delay to transmit link layer frames across the link.

Frame Size

Measure the size of the link layer frames transmitted across the link.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.18 STK Link Monitors

Transmission Delay

Measure the delay to transmit link layer frames across the link.

Frame Size

Measure the size of the link layer frames transmitted across the link.

Repeaters Assigned To User

Measure the number of repeaters/satellites assigned to a user node. A user node may not have a repeater assigned to it, either because the repeater is not visible or the visible repeaters have reached their limits for other users.

Repeaters Assigned To Net Access

Measure the number of repeaters/satellites assigned to net access nodes. A net access node may not have a repeater assigned to it, either because the repeater is not visible or the visible repeaters have reached their limits for other net accesses.

Users Assigned To Repeater

Measure the number of user nodes assigned to a repeater/satellite.

Net Accesses Assigned To Repeater

Measure the number of net accesses that are assigned to a repeater/ satellite.

Net Accesses Assigned to User

Measure the number of net accesses that are assigned to a user node.

Channels Used By User

Measure the number of voice or data channels used by a user node. A channel is used to support a voice call or a data call, but not both at the same time.

Channels Used By Net Access

Measure the number of voice or data channels used by a net access node. A channel is used to support a voice call or a data call, but not both at the same time.

Channels Used By Repeater

Measure the number of voice or data channels used by a repeater/satellite node. A channel is used to support a voice call or a data call, but not both at the same time.

Data Utilization From User

Measure the data transmissions from a user node on the channels when the channels are available for a user or a network access. Once a channel is allocated, the data utilization of the channel may vary.

Data Utilization From Net Access

Measure the data transmissions from a net access node on the channels when the channels are available for a user or a network access. Once a channel is allocated, the data utilization of the channel may vary.

User Call Statistics

Measure the counts of user node calls attempted, carried, blocked due to traffic, blocked due to visibility, disconnected, and preempted. This monitor just saves the total call counts.

Net Access Call Statistics

Measure the counts of net access node calls attempted, carried, blocked due to traffic, blocked due to visibility, disconnected, and preempted. This monitor just saves the total call counts.

Repeater Call Statistics

Measure the counts of repeater/satellite calls attempted, carried, blocked due to traffic, blocked due to visibility, disconnected, and preempted. This monitor just saves the total call counts.

User Data Statistics

Measure the counts of user node data transmissions attempted, carried, blocked due to traffic, blocked due to visibility, disconnected, and preempted. This monitor just saves the total call counts.

Net Access Data Statistics

Measure the counts of net access node data transmissions attempted, carried, blocked due to traffic, blocked due to visibility, disconnected, and preempted. This monitor just saves the total call counts.

Repeater Data Statistics

Measure the counts of repeater/satellite data transmissions attempted, carried, blocked due to traffic, blocked due to visibility, disconnected, and preempted. This monitor just saves the total call counts.

Hop Availability Duration

Measure the length of time when a user has a path through the STK link where a repeater and a network access node are assigned to that user. An available hop means that a path is possible, but the user may not use that available hop if it does not have traffic to send.

Time To Connect

Measure data connections where there is a connection delay and possibly retries to attempt to acquire a channel to use for the data connections. In order for a user to get a channel, a repeater and network access must be assigned to the user and each must have spare channel capacity to support the data call from the user.

Channels Used By Crosslink

Each network access node has a number of crosslinks available for connecting to other network access. This monitor measures the channel usage of each crosslink on each network access. The number of rows in the table is the number of network accesses times the number of possible crosslinks per network access.

Available Crosslink on Net Access

The crosslinks are allocated to mutually visible network accesses. A network access may have the capability of up to N crosslinks to other network accesses where N is the link parameter "'Net Accesses' Crosslinks: Number available". However due to visibility constraints, at any time less then a maximum number of crosslinks may be possible. This monitor measures the number of crosslinks that are available at any time. It is a level monitor measuring a number for each network access and that number changes when visibility to another network access changes. The number of rows in the table is the number of network accesses.

Crosslink Available

This monitor is for each possible crosslink on each network access. It is a level monitor that measures the time it is assigned to a particular network access and, thus, is available for transmitting data. The number of rows in the table is the number of network accesses times the number of crosslinks per network access.

Data utilization from crosslink

Measure the data transmissions from a crosslink node on the channels when the channels are available for a crosslink. Once a channel is allocated, the data utilization of the channel may vary.

Crosslink call statistics

Measure the counts of crosslink node calls attempted, carried, blocked due to traffic, blocked due to visibility, disconnected, and preempted. This monitor just saves the total call counts.

set of twelve statistics requests of the form "[User|Net Access|Crosslink]+[bytes|packets]+by+[Protocol|App type]".

This monitor provides link-breakout usage statistics for each combination, as listed. These monitors are counter monitors that are distinct from statistics monitors because they only track the count of occurrences rather than statistics for means or standard deviations. Each row is a user, a network access, or a network access crosslink. Each column is a different protocol or application type. The values in the table are either in bytes or in packets. The table may be plotted as multi-series bar charts.

[User|Net Access|Crosslink] frame delay.

The basic monitor for collecting statistics on the frame delay for each user, network access, or network access crosslink.

1.19 TDM Link Monitors

Node Activity Duration

Measure the time the nodes spend both waiting for a channel to become available, and the time it spends using the channel.

Active Node Count

Measure the number of nodes using channels, and the number of nodes waiting to use the channels.

Utilization

Measure the percent capacity at which the link is utilized.

Transmission Delay

Measure the delay to transmit link layer frames across the link.

Frame Size

Measure the size of the link layer frames transmitted across the link.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.20 TDMA Link Monitors

Busy Users

Measure the number of TDMA channels in use at any time.

Cycle Used

Measure the number of cycles used on the link.

Time To Connect

Measure the time required to establish a connection before data can be transmitted.

Node Activity Duration

Measure the time the nodes spend both waiting for a channel to become available, and the time it spends using the channel.

Active Node Count

Measure the number of nodes using channels, and the number of nodes waiting to use the channels.

Utilization Uplink

Measure the percent capacity at which the uplink is utilized.

Utilization Downlink

Measure the percent capacity at which the downlink is utilized.

Transmission Delay Uplink

Measure the delay to transmit link layer frames across the uplink.

Transmission Delay Downlink

Measure the delay to transmit link layer frames across the downlink.

Frame Size

Measure the size of the link layer frames being transmitted across the link.

Session In-Progress

Measure the number of sessions in progress on the link.

1.21 Token Passing Link Monitors

Node Active Duration

Measure the amount of time a node spends actively acquiring and using the token to empty its output buffer. Once the output buffer is empty, the node is no longer active. This monitor helps to identify if multiple token rotations are required to empty a node's buffer.

Active Node Count

Measure the number of nodes actively trying to capture the token. This is the number of nodes with non-empty buffers.

Token Rotation Time

Measure the average time for the token to pass to all the nodes connected to the link.

Utilization

Measure the percent capacity at which the link is utilized.

Transmission Delay

Measure the delay to transmit link layer frames across the link.

Frame Size

Measure the size of the link layer frames transmitted across the link.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.22 WAN Link Monitors

The From X and From Y directions are determined by the order in which the nodes are connected to the Dialup Point-to-Point link. When editing the properties of the link, the X and Y nodes are listed across the top of the link dialog.

Burst Size From X

Measure the burst size of data from Node X on a packet by packet basis.

Burst Size From Y

Measure the burst size of data from Node Y on a packet-by-packet basis.

Transmission Delay From X

Measure the delay to transmit link layer frames across the link from Node X.

Transmission Delay From Y

Measure the delay to transmit link layer frames across the link from the Node Y.

Frame Size From X

Measure the size of the link layer frames transmitted across the link from the Node X.

Frame Size From Y

Measure the size of the link layer frames transmitted across the link from the Node Y.

Sessions In-Progress

Measure the number of sessions in progress on the link.

1.23 WAN Cloud Access Link Monitors

Exit Buffer Size

Measure the size of the exit or egress buffer on the WAN Cloud access link.

Mean Exit Buffer Size

Measure the mean of the exit or egress buffer on the WAN Cloud access link.

Entry Link Utilization

Measure the capacity at which the WAN Cloud entry link is utilized.

Mean Entry Link Utilization

Measure the mean capacity at which the WAN Cloud entry link is utilized.

Exit Link Utilization

Measure the capacity at which the WAN Cloud exit link is utilized.

Mean Exit Link Utilization

Measure the mean capacity at which the WAN Cloud exit link is utilized.

R-T Mean Exit-Buffer Level

Measure the mean exit buffer level of the WAN Cloud Access link and display the information in a real time graph.

R-T Mean Entry Utilization

Measure the mean entry link utilization of the WAN Cloud and display the information in a real time graph.

R-T Mean Exit Utilization

Measure the mean exit link utilization of the WAN Cloud and display the information in a real time graph.

1.24 WAN Cloud Virtual Circuit Link Monitors

Burst Size

Measure the burst size of the WAN Cloud virtual circuit link.

Sampled Burst Size

Measure the sampled burst size of the WAN Cloud virtual circuit link.

Frame Delay

Measure the link layer frame delay across the WAN Cloud virtual circuit link.

Mean Frame Delay

Measure the link layer mean frame delay across the WAN Cloud virtual circuit link.

Real-Time Sampled Burst Size

Displays a graph in real time of the sampled burst size of the WAN Cloud virtual circuit link.

Real-Time Mean Frame Delay

Displays a graph in real time of the link layer mean frame delay across the WAN Cloud virtual circuit link.

1.25 Message Sources Monitors

Message Delay

Measure the message delay, from a message source, on the message after the message has been reassembled by the destination, and the sender is notified that the message has been assembled. This is the message delay that is perceived by the sender.

Real-Time Message Delay

Displays the message delay in a real time graph.

Msg Delivered Delay (sec)

Measure the message delivery delay, from a message source, before the message is reassembled by the destination. The message may still be worked on by the source after the message is delivered because of retransmissions, waiting for acks, or for modeling the close-sequence of the connection, but these values are not factored into the message delivered delay.

Real-Time Msg Delivered Delay

Displays the message delivery delay in a real time graph.

Messages In-Transit

Measure the number of messages in transit, from a message source.

Real-Time Messages In-Transit

Displays the number of messages in transit from a message source in a real time graph.

1.26 Response and Session Sources Monitors

Message Delay

This monitor measures the message delay, from a response source, on the message after the message has been reassembled by the destination, and the sender is notified that the message has been assembled. This is the message delay that is perceived by the sender.

Real-Time Message Delay

Displays the message delay in a real time graph.

Msg Delivered Delay (sec)

Measure the message delivery delay, from a response source, before the message is reassembled by the destination. The message may still be worked on by the source after the message is delivered because of retransmissions, waiting for acks, or for modeling the close-sequence of the connection, but these values are not factored into the message delivered delay.

Real-Time Msg Delivered Delay

Displays the message delivery delay in a real time graph.

Messages In-Transit

Measure the number of messages in transit, from a response source.

Real-Time Messages In-Transit

Displays the number of messages in transit in a real time graph.

1.27 Application Sources Monitors

Application Source Delay

Measure the time to execute the completed applications, counted from the point at which they were first scheduled until the point at which their last command completes. This includes time spent waiting in the application pending list of the node while other applications are executing.

Successful Application Source Delay

Measure the time to execute the successfully completed applications, counted from the point at which they were first scheduled until the point at which their last command completes. This includes time spent waiting in the application pending list of the node while other applications are executing.

Timed Out Application Source Delay

Measure the time for an application to time out, counted from the point at which they were first scheduled until the point at which they timed out. This includes time spent waiting in the application pending list of the node while other applications are executing.

Real-Time Application Source Delay

Displays the time to execute the completed applications in a real time graph.

1.28 Input Buffer Monitors

Port Buffer Level

Measure the amount of buffer space in use at the port input buffer.

Real-Time Port Buffer Level

Measure and display in a real time graph the amount of buffer space in use at the port input buffer.

Packet Delay In Port

Measure the packet delay incurred on the port input buffer.

Real-Time Port Packet Delay

Measure and display in a real time graph the packet delay incurred on the port input buffer.

Busy Processors

Measure the number of busy processors. For a single processor network device, the result is the utilization of the processor.

1.29 Output Buffer Monitors

Port Buffer Level

Measure the amount of buffer space in use at the port output buffer.

Real-Time Port Buffer Level

Measure and display in a real time graph the amount of buffer space in use at the port output buffer.

Packet Delay In Port

Measure the packet delay incurred on the port output buffer.

Real-Time Port Packet Delay

Displays the packet delay incurred on the port output buffer in a real time graph.

Busy Processors

Measure the number of busy processors. For a single processor network device, the result is the utilization of the processor.

Section II

Reports

COMNET III generates a number of textual reports at the end of each replication of the model. They can be selectively turned on by choosing the menu item **Report/Select Reports** and then selecting the appropriate report from **Reporters** dialog. The **Reporters** dialog box shows each individual model element (node, link, etc.). Here is an example of the **Reporters** dialog box.

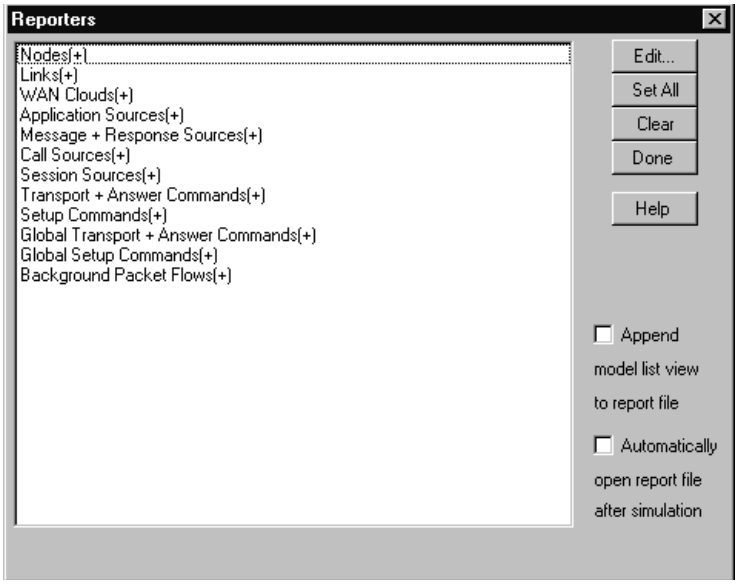

Reporters Dialog Box

Edit Use the **Edit** button to set individual reports.

Set All Select the **Set All** button if you want all reports turned on.

Append model list view to report file

Checking the box **Append model list view to report file** causes a listing of all characteristics of the model to be appended to the end of the report file. The listing includes everything that would be seen by selecting the **View** menu option and choosing the **View By List** option after selecting **All** and excluding items with default values. This feature is an excellent way to fully document a simulation run.

Automatically open report file after simulation

Checking the box **Automatically open report file after simulation** causes the report files to automatically open when the simulation completes.

Selecting the **Edit** button will cause a sub-dialog to open when there are various objects related to the individual model elements listed.

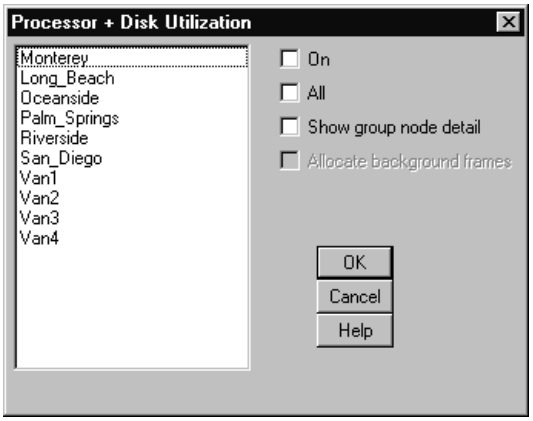

Selecting Reports on Processor and Disk Utilization

On
All

Select the objects for which there should be reports and check the **On** box, or if there should be reports on all objects shown, check the **All** box.

Show group node detail

In all of the reports data for Computer Group nodes is aggregated. Check to display the **Group Node Detail** dialog.

Allocate background frames

Checking the **Allocate background frames** causes background frames to be included in the listed protocol or application grouping in the **`Link Utilization by Protocol`** or **`Link Utilization by Application`** reports, respectively. The frames will be allocated to the listed groupings in proportion to the amount these groupings are observed in the non-background frames traversing the link.

If the checkbox is left unchecked, the background frames are listed in their own category. Background frames are frames generated by link loading. This checkbox does not apply to any other report.

Statistical monitoring during the simulation is active only for the statistics that need to be included in formatted reports (or viewed through **Statistics** buttons). By requesting fewer reports and including fewer objects in reports, you can speed up the simulation.

These reports are prepared as soon as the simulation run completes. They are in standard text format and are formatted to fit within 80 columns for easy viewing and printing. The **Report/Browse Reports** menu item may be used to view the reports once the simulation is complete, or any text editor may be used.

Wherever there is an object name in the report (i.e. node name, link name, source name, etc), only the first 20 characters will be displayed. The remaining characters will be truncated. This means that a node name of "Application Workstatio" is really probably "Application Workstation".

The reports are placed in a subdirectory which has the same name as the model. Within the subdirectory all the report replications are in the file **`Stat1.rpt`**. If you run a simulation again, this file will be overwritten with the new reports for the next run.

Also, when a report presents a mean and a standard deviation for some measurement, it should not be interpreted to mean that the measured value is normally distributed about the average—it probably is not.

2. Nodes

Nodes represent processing, switching, and bridging equipment in the network. Refer to Chapter 3 in the *COMNET III Reference Guide* for more information.

COMNET III produces the fifteen different Node reports described in the following paragraphs.

The first column in most of the node reports is titled:

NODE

or

NODE NAME

This column lists the name of the node.

2.1 Node Full Utilization Report

NODES: NODE FULL UTILIZATION

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODES: / PORTS:	BUSY CENTRAL PROCESSORS — BUSY INPUT PROCESSORS MEAN	MAX	UTIL%	BUSY BUSSES — BUSY OUTPUT PROCESSORS MEAN	MAX	UTIL%
Message Workstation	.0001033	1	.0103333	N/A	N/A	N/A
Ethernet Hub	0.0	0	0.0	0.0	0	0.0
Session Server	.0009833	1	.0983333	N/A	N/A	N/A
Ethernet Hub	0.0	0	0.0	0.0	0	0.0
Application Workstatio	.0005188	1	.0518800	N/A	N/A	N/A
Ethernet Hub	0.0	0	0.0	0.0	0	0.0
Router 1	0.0	1	0.0	.0000024	1	.0002385
WAN Cloud	0.0	0	0.0	0.0	0	0.0
Ethernet Hub	0.0	0	0.0	0.0	0	0.0
Router 2	0.0	1	0.0	.0000013	1	.0001318
WAN Cloud	0.0	0	0.0	0.0	0	0.0
Ethernet Hub	0.0	0	0.0	0.0	0	0.0

The Node Full Utilization Report presents a summary of the utilization level for the processors and ports on each node.

Busy Central Processors—Mean, Max

The mean and maximum of the number of processors busy.

Util% Busy Central Processors

The percent of capacity at which the node's central processors were utilized.

Busy Busses—Mean, Maximum

The mean and maximum number of busses that are busy.

Util % Busy Busses

The percent of capacity at which the node's busses were utilized.

Ports The name of the link or links that the node is attached to in the model.

Busy Input Processors—Mean, Maximum

The mean and maximum number of busy input port processors on a node.

Util% Busy Input Processors

The percent of capacity at which the node's busy input port processors were utilized.

Busy Output Processors—Mean, Max

The mean and maximum number of busy output port processors on a node.

Util% Busy Output Processors

The percent of capacity at which the node's busy output port processors were utilized.

2.2 Processor and Disk Utilization Report

NODES: PROCESSOR + DISK UTILIZATION

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE	DISK REQSTS GRNTED	DISK USAGE (KILOBYTES) AVERAGE	MAXIMUM	STD DEV	PROCESSOR % UTIL
Message Workstation	0	10.000	10.000	0.000	0.0103
Session Server	0	65.000	65.000	0.000	0.0983
Application Workstat	36	10.000	10.000	0.000	0.0519
Application Server	20	65.000	65.000	0.000	96.04
Call Node 4	0	0.000	0.000	0.000	0.0000
Call Node 1	0	0.000	0.000	0.000	0.0000
Call Node 2	0	0.000	0.000	0.000	0.0000
Call Node 3	0	0.000	0.000	0.000	0.0000
Response Server	0	65.000	65.000	0.000	0.0163
Session Workstation	0	10.000	10.000	0.000	0.0000
WAN Workstation 1	0	10.000	10.000	0.000	0.0100
WAN Workstation 2	0	10.000	10.000	0.000	0.0090
Router 1	0	0.000	0.000	0.000	0.0000
Router 2	0	0.000	0.000	0.000	0.0000

The Processor and Disk Utilization Report presents a summary of the utilization level for the processors and disk utilization on each node.

Disk Requests Granted

The number of disk requests granted. Disk requests are issued by read and write commands to access files stored on the local disk of a Processing node, a Computer Group node, or a Network Device node. An interrupted read or write command (because of time slicing or node failure) will count as multiple access.

Disk Usage (KB)—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of disk usage.

Processor Utilization

At any instant in time the node processor is either idle or busy. The processor is counted as busy when executing process commands, read and write commands are accessing files, or packets are being created by message, session or response generators. When a packet is created, the processor is made busy for the packetizing delay followed by the packet switching time. In addition, the processor is busy for the switching time when packets are routed across the node.

2.3 Received Message Count Report

NODES: RECEIVED MESSAGE COUNTS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

RECEIVER	COUNT	MESSAGE NAME
Message Workstation	31	Message Source
Session Server	295	Session Source
Session Server	14	WAN Message 2
Session Server	6	WAN Message 1
Application Workstatio	16	Request Response
Application Workstatio	4	DataDB Request
Application Server	16	DataDB Request
Application Server	73	Local Setup
Application Server	16	Global DB Request
Application Server	16	Global Setup
Response Server	7	WAN Message 2
Response Server	11	WAN Message 1
Response Server	31	Message Source
Session Workstation	295	Session Source
WAN Workstation 1	11	WAN Message 1
WAN Workstation 1	6	WAN Message 2
WAN Workstation 2	7	WAN Message 2
WAN Workstation 2	13	WAN Message 1

The Received Message Counts Report presents a count of received messages for each destination node listed by message name.

Receiver The name of the receiving node.

Count The number of messages received.

Message Name

The name of the received message. This is the message text of the message, session, answer, or transport command that sent the message. This is controlled by the "Text" tab of the sources or commands.

2.4 Disk Access Error Counts Report

NODES: DISK ACCESS ERROR COUNTS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE NAME	# TIMES STORAGE CAPACITY EXCEEDED	# TIMES NO READ FILE EXISTED	# TIMES READ BYTES IN FILE EXCEEDED
Message Workstation	0	0	0
Session Server	0	0	0
Application Workstatio	0	16	0
Application Server	0	4	0
Call Node 4	0	0	0
Call Node 1	0	0	0
Call Node 2	0	0	0
Call Node 3	0	0	0
Response Server	0	0	0
Session Workstation	0	0	0
WAN Workstation 1	0	0	0
WAN Workstation 2	0	0	0
Router 1	0	0	0
Router 2	0	0	0

When there are problems with file I/O at any of the Processing, Computer Group, or Network Device nodes, COMNET III does not stop the simulation. Instead, it keeps a count of these problems and presents them in this report.

Number Of Times Storage Capacity Exceeded

When write commands are executed on a Processing, Computer Group, or a Network Device node, it is possible that the local disk capacity modeled for the node will be exceeded. The host node and the disk are counted as busy for the length of time it takes to execute the requested transfer and the number of times this occurred is reported. At the end of the transfer the disk will be full, but its size will not be increased.

Number Of Times No Read File Existed

When read commands are executed on Processing, Computer Group, or a Network Device node, it is possible that the requested file does not exist at that instant in time. The host node and the disk will be counted as busy as if the transfer took place, but the file will still not exist at the end of the transfer. The number of times this occurred is reported.

Number Of Times Read Bytes In file Exceeded

When read commands are executed on Processing, Computer Group, or a Network Device node, it is possible that the requested file is not as big

as the requested transfer. The host node and the disk will be counted as busy as if the full transfer took place, but the file will not be incremented in size at the end of the transfer. The number of times this occurred is reported.

2.5 Session Level Report

NODES: SESSION LEVEL

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

NODE	SESSIONS IN PROGRESS AVERAGE	STD DEV	MAXIMUM
Workstation 1	25.25	12.49	42
Workstation 2	9.30	5.48	17
Workstation 3	14.69	6.84	25

The Session Level Report presents a summary of the number of sessions in progress on each node. A node may originate sessions, carry session packets in route to their destination, or be the destination of the session. In all cases, a session is set up through the node to carry the session traffic.

Sessions set up across subnetworks where a connectionless routing algorithm is in operation are not reported, even if some of the session packets are routed through the node. This is because the complete session is not associated with the particular node.

Sessions In Progress—Average, Standard Deviation, Maximum

The average, standard deviation and maximum number of concurrent sessions in progress on this node.

2.6 Call Counts Report

NODES: CALL COUNTS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE NAME	CALLS ATTEMPTD	CALLS BLOCK AVAIL	CALLS BLOCK TRAFF	BLOCK PROB	CALLS CARRIED	CALLS DISCON-NECTED	CALLS PRE-EMPTED
Message Workstation	0	0	0	0.00	0	0	0
Session Server	0	0	0	0.00	0	0	0
Application Workstation	0	0	0	0.00	0	0	0
Switch	68	0	0	0.00	68	0	0
Application Server	0	0	0	0.00	0	0	0
Call Node 4	47	0	0	0.00	47	0	0
Call Node 1	34	0	0	0.00	34	0	0
Call Node 2	44	0	0	0.00	44	0	0
Call Node 3	43	0	0	0.00	43	0	0
Response Server	0	0	0	0.00	0	0	0
Session Workstation	0	0	0	0.00	0	0	0
WAN Workstation 1	0	0	0	0.00	0	0	0
WAN Workstation 2	0	0	0	0.00	0	0	0
Router 1	0	0	0	0.00	0	0	0
Router 2	0	0	0	0.00	0	0	0

The Call Counts Report represents a summary of all the calls attempted on a node, and whether they were blocked, carried, disconnected, or preempted.

Calls Attempted

The number of calls that have been attempted in terms of routing new call instances to their destination from this node.

Calls Blocked - Availability

The number of calls that could not be routed over this node because this node had failed.

Calls Blocked - Traffic

The number of calls that could not be routed over this node because all the node capacity had been allocated to other calls.

Blocking Probability

Total Calls Blocked divided by Calls Attempted for calls at this node.

Calls Carried

The number of calls that have been successfully routed to their destination over this node.

Calls Disconnected

The number of calls disconnected due to this node failing.

Calls Preempted

The number of calls carried at this node and disconnected due to preemption by a higher priority call routing through this node. Preemption occurs when there is insufficient bandwidth available on the node to route the higher priority call. The preempting call does not have to be between the same origin and destination or following the same route as the preempted call. Even if there is only one common node or link between the calls then preemption may occur.

2.7 Call Level Report

NODES: CALL LEVEL

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE NAME	% AVAIL	NODE FAILS	BANDWIDTH USED (KBPS) AVERAGE	STD DEV	MAXIMUM	NODE UTIL %
Message Workstation	100.00	0	0	0	0	0.00
Session Server	100.00	0	0	0	0	0.00
Application Workstation	100.00	0	0	0	0	0.00
Switch	100.00	0	2302	966	3072	23.02
Application Server	100.00	0	0	0	0	0.00
Call Node 4	100.00	0	1187	498	1600	11.87
Call Node 1	100.00	0	1115	478	1600	11.15
Call Node 2	100.00	0	1115	478	1600	11.15
Call Node 3	100.00	0	1187	498	1600	11.87
Response Server	100.00	0	0	0	0	0.00
Session Workstation	100.00	0	0	0	0	0.00
WAN Workstation 1	100.00	0	0	0	0	0.00
WAN Workstation 2	100.00	0	0	0	0	0.00
Router 1	100.00	0	0	0	0	0.00
Router 2	100.00	0	0	0	0	0.00

The Call Level Report represents a summary of all the call bandwidths used and available on a node.

Node Name The name of the node. Each node has a bandwidth capacity which it may use to carry call traffic.

% Availability

A node is not available when it is in the failed state. The percentage of time available is reported. This is the Up Time divided by the Run Length.

Node Failures

The number of failures that occurred.

Bandwidth Used (Kbps)—Average, Standard Deviation, Maximum

The average, standard deviation and maximum peak bandwidth used on the node.

Node Utilization

Average Bandwidth Used/Total Available Bandwidth expressed as a percentage.

2.8 Input Buffer Totals

NODES: INPUT BUFFER TOTALS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE	PACKETS ACCEPTED	PACKETS BLOCKED	BUFFER USE (P=PKTS, B=BYTES) AVERAGE	STD DEV	MAXIMUM	P/B
Message Workstation	186	0	0	0	1040	B
Session Server	1470	0	0	0	1431	B
Application Workstation	967	0	0	19	4540	B
Switch	8149	0	0	0	295	B
Application Server	824	0	4938	4994	43740	B
Call Node 4	47	0	0	0	295	B
Call Node 1	13049	0	0	0	193	B
Call Node 2	1219	0	0	0	295	B
Call Node 3	1005	0	0	0	275	B
Response Server	899	0	0	0	1040	B
Session Workstation	905	0	0	0	1040	B
WAN Workstation 1	141	0	0	0	1040	B
WAN Workstation 2	141	0	0	0	1040	B
Router 1	168	0	0	0	1040	B
Router 2	168	0	0	0	1040	B

The Input Buffer Nodes Report represents a summary of each node's input buffer usage. The node is connected to links via an interface port. Each port has an input buffer, the size of which is defined on the port (edit the arc that connects the link to the node to see this). The amount of input buffer space on the node is the sum total of all the port input buffer spaces. In addition, the node has an upper maximum on the total amount of input buffer space that can be in use at one time across all ports. This is defined on the node parameters screen.

When a packet is received, the port input buffer is first checked to see if there is space. If so, then the node input limit is checked to see if the node in total has space. If both tests succeed, the packet is received into the input buffer.

Packets Accepted

The number of packets received into the node across all input ports on the node.

Packets Blocked

The number of packets blocked across all input ports on the node due to insufficient buffer space (either port space or total node space).

Buffer Use - Average, Std Dev, Max

The average, standard deviation and maximum amount of buffer space in use across all input ports on the node described in either packets or bytes.

2.9 Output Buffer Totals Report

NODES: OUTPUT BUFFER TOTALS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE	PACKETS ACCEPTED	PACKETS BLOCKED	BUFFER USE (P=PKTS, B=BYTES) AVERAGE	STD DEV	MAXIMUM	P/B
Message Workstation	186	0	0	10	1040	B
Session Server	1076	0	1	34	1431	B
Application Workstation	654	0	3	131	12520	B
Switch	5358	0	1	14	896	B
Application Server	673	0	2	109	10840	B
Call Node 4	25	0	0	0	90	B
Call Node 1	3561	0	1	9	232	B
Call Node 2	520	0	0	10	307	B
Call Node 3	679	0	0	1	112	B
Response Server	582	0	0	12	1040	B
Session Workstation	1200	0	1	28	1040	B
WAN Workstation 1	141	0	16	127	1080	B
WAN Workstation 2	141	0	14	121	1080	B
Router 1	168	0	0	11	1040	B
Router 2	168	0	0	19	1040	B

The Output Buffer Totals Report represents a summary of each node's output buffer usage. The node is connected to links via an interface port. Each port has an output buffer, the size of which is defined on the port (edit the arc that connects the link to the node to see this). The amount of output buffer space on the node is the sum total of all the port output buffer spaces. In addition, the node has an upper maximum on the total amount of output space that can be in use at one time across all ports. This is defined on the node parameters screen.

When a packet is routed across the node it must be placed into an output buffer port. First the port output buffer is checked to see if there is space. If so, then the node output limit is checked to see if the node in total has space. If both tests succeed then the packet is placed into the output buffer.

Packets Accepted

The number of packets placed by the node's output buffers across all output ports on the node.

Packets Blocked

The number of packets blocked across all output ports on the node due to insufficient buffer space (either port space or total node space).

Buffer Use - Average, Std Dev, Max

The average, standard deviation and maximum buffer space in use across all output ports on the node.

2.10 Input Node Buffer Policy Report

NODES: INPUT NODE BUFFER POLICY

REPLICATION 1 FROM 0.0 TO 3600.0 SECONDS

	PREEMPTION		THRESHOLD	
NODE	PACKETS PREEMPTING	PACKETS PREEMPTED	PACKETS EXCEEDED	PACKETS REJECTED
Source/Destination 1	36	36	0	0

The Input Node Buffer Policy Report summarizes the input buffer policy actions aggregated over the entire node. The preemption and threshold values reported are aggregated for all input ports on the node.

The input buffer on a node first determines that there is room for a packet by checking that the specific input buffer has buffer space available to fit the incoming packet. It next checks that the total space used for all similar input buffers on the node does not exceed the buffer constraint in the node's parameter set. The packet may be accepted in an input buffer if it fits in both the port-specific buffer and in the total buffer space available on the node.

There are two other policies available at the input port to refine the initial test for admitting the packet into an input buffer. These policies are a "Preemption Policy" and a "Threshold Policy". These policies determine which packets may be accepted at an input buffer.

The "Preemption Policy" checks to see if the packet will be accepted into the input buffer if the packet has passed the first check and does not initially fit in the input buffer. In this case, if the packet has sufficient priority, it can cause lower priority packets to be ejected from the buffer, and thus blocked, until there is enough space in the buffer for the incoming packet. The lower priority packets will be ejected only if there can be enough space made available for the incoming packet.

The "Threshold Policy" checks to see if the packet will be accepted into the input buffer if the packet has passed the first test. This is so that if a packet fits in the buffer it can be rejected if the buffer size is above a threshold, and the packet lacks sufficient priority to use the congested buffer. This mechanism is commonly used for rejecting packets that are marked as "discard eligible" when the buffers are above their threshold in frame-switching networks such as ATM or frame-relay.

Node Name The name of the node sending packets.

Packets Preempting

The number of packets that preempted other packets in the buffer, thus causing those other packets to be blocked. A higher priority packet that arrives at the input buffer that is full and preempts a lower priority packet.

Packets Preempted

The number of packets that were blocked due to preemption. A lower priority packet has been ejected from the input buffer by a higher priority packet. A large, high priority packet can eject many smaller, lower priority packets from the input buffer.

Packets Exceeded

Packets that have been accepted at the input buffer when the input buffer threshold has been exceeded.

Packets Rejected

Packets that have been rejected at the input buffer when the input buffer threshold has been exceeded.

2.11 Output Node Buffer Policy Report

NODES: OUTPUT NODE BUFFER POLICY

REPLICATION 1 FROM 0.0 TO 3600.0 SECONDS

	PREEMPTION		THRESHOLD	
NODE	PACKETS PREEMPTING	PACKETS PREEMPTED	PACKETS EXCEEDED	PACKETS REJECTED
Source/Destination 1	36	36	0	0

The Output Node Buffer Policy Report summarizes the output buffer policy actions aggregated over the entire node. The preemption and threshold values reported are aggregated for all output ports on the node.

The output buffer on a node first determines that there is room for a packet by checking that the specific output buffer has buffer space available to fit the incoming packet. It next checks that the total space used for all similar output buffers on the node do not exceed the buffer constraint in the node's parameter set. The packet may be accepted in an output buffer if it fits in both the port-specific buffer and in the total buffer space available on the node.

There are two other policies available at the output port to refine the initial test for admitting the packet into an output buffer. These policies are a "Preemption Policy" and a "Threshold Policy". These policies determine which packets may be accepted at an output buffer.

The "Preemption Policy" checks to see if the packet will be accepted into the output buffer if the packet has passed the first check and does not initially fit in the output buffer. In this case, if the packet has sufficient priority, it can cause lower priority packets to be ejected from the buffer, and, thus, be blocked until there is enough space in the buffer for the incoming packet. The lower priority packets will be ejected only if there can be enough space made available for the incoming packet.

The "Threshold Policy" checks to see if the packet will be accepted into the output buffer if the packet has passed the first test. This is so that if a packet fits in the buffer it can be rejected if the buffer size is above a threshold, and the packet lacks sufficient priority to use the congested buffer. This mechanism is commonly used for rejecting packets that are marked as "discard eligible" when the buffers are above their threshold in frame-switching networks such as ATM or frame-relay.

Node Name The name of the node sending packets.

Packets Preempting

The number of packets that preempted other packets in the buffer, thus causing those other packets to be blocked. A higher priority packet that arrives at the output buffer that is full and preempts a lower priority packet.

Packets Preempted

The number of packets that were blocked due to preemption. A lower priority packet has been ejected from the output buffer by a higher priority packet. A large, high priority packet can eject many smaller, lower priority packets from the output buffer.

Packets Exceeded

Packets that have been accepted at the output buffer when the output buffer threshold has been exceeded.

Packets Rejected

Packets that have been rejected at the output buffer when the output buffer threshold has been exceeded.

2.12 Input Buffer Use By Port Report

NODES: INPUT BUFFER USE BY PORT

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE:	PACKETS		BUFFER USE (P=PKTS, B=BYTES)			
CONNECTED LINKS	ACCEPTED	BLOCKED	AVERAGE	STD DEV	MAXIMUM	P/B
Message Workstation:						
Ethernet Hub	186	0	0	0	1040	B
Background Pkt Src	0	0	0	0	0	B
Session Server:						
Ethernet Hub	1359	0	0	0	1040	B
Background Pkt Src	111	0	0	0	1431	B
Application Workstation:						
Ethernet Hub	921	0	0	19	4540	B
Background Pkt Src	46	0	0	0	95	B
Application Server:						
Ethernet Hub	727	0	3976	4833	43740	B
Background Pkt Src	97	0	962	944	3696	B

The Input Buffer Use By Port Report represents a summary of each node's port input buffer usage during the simulation. The node is connected to each link via an interface port. Each port has an input buffer, the size of which is defined on the port (edit the arc that connects the link to the node to see this). The amount of input buffer space on the node is the sum total of all the port input buffer spaces. In addition, the node has an upper maximum on the total amount of input buffer space that can be in use at one time across all ports. This is defined on the node parameters screen.

When a packet is received, first the port input buffer is checked to see if there is space. If so, the node input limit is checked to see if the node in total has space. If both tests succeed, the packet is received into the input buffer.

Node Name The name of the node receiving packets.

Connected Links

The name of the link whose input port is being reported.

Packets Accepted

The number of packets received into the port input buffer.

Packets Blocked

The number of packets blocked at this port input buffer.

Buffer Use - Average, Std Dev, Max

The average, standard deviation and maximum amount of buffer space in use at this port input buffer described in either packets or bytes.

2.13 Output Buffer Use By Port Report

NODES: OUTPUT BUFFER USE BY PORT

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE:	PACKETS		BUFFER USE (P=PKTS, B=BYTES)			
CONNECTED LINKS	ACCEPTED	BLOCKED	AVERAGE	STD DEV	MAXIMUM	P/B
Message Workstation:						
Ethernet Hub	186	0	0	10	1040	B
Session Server:						
Ethernet Hub	1076	0	1	34	1431	B
Application Workstation:						
Ethernet Hub	654	0	3	131	12520	B
Switch:						
Call Link 4	22	0	0	2	530	B
Call Link 1	3536	0	1	9	232	B
Call Link 2	699	0	0	1	88	B
Call Link 3	326	0	0	10	368	B
Ethernet Hub	775	0	0	1	112	B
Application Server:						
Ethernet Hub	673	0	2	109	10840	B

The Output Buffer Use By Port Report represents a summary of each node's port output buffer usage during the simulation. The node is connected to links via an interface port. Each port has an output buffer, the size of which is defined on the port (edit the arc that connects the link to the node to see this). The amount of output buffer space on the node is the sum total of all the port output buffer spaces. In addition, the node has an upper maximum on the total amount of output space that can be is use at one time across all ports. This is defined on the node parameters screen.

When a packet is routed across the node it must be placed into an output buffer port. First the port output buffer is checked to see if there is space. If so, then the node output limit is checked to see if the node in total has space. If both tests succeed then the packet is placed into the output buffer.

Node Name The name of the node sending packets.

Connected Links

The name of the link whose output port is being reported.

Packets Accepted

The number of packets routed into the port output buffer.

Packets Blocked

The number of packets blocked at the port output buffer.

Buffer Use - Average, Std Dev, Max

The average, standard deviation and maximum amount of buffer space in use at the port output buffer in either packets or bytes.

2.14 Input Port Buffer Policy Report

NODES: INPUT PORT BUFFER POLICY

REPLICATION 1 FROM 0.0 TO 3600.0 SECONDS

NODE: CONNECTED LINKS	PREEMPTION PACKETS PREEMPTING	PREEMPTION PACKETS PREEMPTED	THRESHOLD PACKETS EXCEEDED	THRESHOLD PACKETS REJECTED
Source/Destination 1:				
Link1	36	36	0	0

The Input Port Buffer Policy Report summarizes the input buffer policy actions on a port-by-port basis. The preemption and threshold values reported are specific for each input port on the node.

The input buffer on a node first determines that there is room for a packet by checking that the specific input buffer has buffer space available to fit the incoming packet. It next checks that the total space used for all similar input buffers on the node do not exceed the buffer constraint in the node's parameter set. The packet may be accepted in an input buffer if it fits in both the port-specific buffer and in the total buffer space available on the node.

There are two other policies available at the input port to refine the initial test for admitting the packet into an input buffer. These policies are a "Preemption Policy" and a "Threshold Policy". These policies determine which packets may be accepted at an input buffer.

The "Preemption Policy" checks to see if the packet will be accepted into the input buffer if the packet has passed the first check and does not initially fit in the input buffer. In this case, if the packet has sufficient priority, it can cause lower priority packets to be ejected from the buffer, and, thus, be blocked, until there is enough space in the buffer for the incoming packet. The lower priority packets will be ejected only if there can be enough space made available for the incoming packet.

The "Threshold Policy" checks to see if the packet will be accepted into the input buffer if the packet has passed the first test. This is so that if a packet fits in the buffer it can be rejected if the buffer size is above a threshold, and the packet lacks sufficient priority to use the congested buffer. This mechanism is commonly used for rejecting packets that are marked as "discard eligible" when the buffers are above their threshold in frame-switching networks such as ATM or frame-relay.

Node Name The name of the node sending packets.

Packets Preempting

The number of packets that preempted other packets in the buffer, thus causing those other packets to be blocked. A higher priority packet arrives at the input buffer that is full and preempts a lower priority packet.

Packets Preempted

The number of packets that were blocked due to preemption. A lower priority packet has been ejected from the input buffer by a higher priority packet. A large, high priority packet can eject many smaller, lower priority packets from the input buffer.

Packets Exceeded

Packets that have been accepted at the input buffer when the input buffer threshold has been exceeded.

Packets Rejected

Packets that have been rejected at the input buffer when the input buffer threshold has been exceeded.

2.15 Output Port Buffer Policy Report

NODES: OUTPUT PORT BUFFER POLICY

REPLICATION 1 FROM 0.0 TO 3600.0 SECONDS

NODE: CONNECTED LINKS	PREEMPTION PACKETS PREEMPTING	PREEMPTION PACKETS PREEMPTED	THRESHOLD PACKETS EXCEEDED	THRESHOLD PACKETS REJECTED
Source/Destination 1:				
Link1	36	36	0	0

This report summarizes the output buffer policy actions on a port-by-port basis. The preemption and threshold values reported are specific for each output port on the node.

The output buffer on a node first determines that there is room for a packet by checking that the specific output buffer has buffer space available to fit the incoming packet. It next checks that the total space used for all similar output buffers on the node does not exceed the buffer constraint in the node's parameter set. The packet may be accepted in an output buffer if it fits in both the port-specific buffer and in the total buffer space available on the node.

There are two other policies available at the output port to refine the initial test for admitting the packet into an output buffer. These policies are a "Preemption Policy" and a "Threshold Policy". These policies determine which packets may be accepted at an output buffer.

The "Preemption Policy" checks to see if the packet will be accepted into the output buffer if the packet has passed the first check and does not initially fit in the output buffer. In this case, if the packet has sufficient priority, it can cause lower priority packets to be ejected from the buffer, and, thus, be blocked until there is enough space in the buffer for the incoming packet. The lower priority packets will be ejected only if there can be enough space made available for the incoming packet.

The "Threshold Policy" checks to see if the packet will be accepted into the output buffer if the packet has passed the first test. This is so that if a packet fits in the buffer it can be rejected if the buffer size is above a threshold, and the packet lacks sufficient priority to use the congested buffer. This mechanism is commonly used for rejecting packets that are marked as "discard eligible" when the buffers are above their threshold in frame-switching networks such as ATM or frame-relay.

Node Name The name of the node sending packets.

Packets Preempting

The number of packets that preempted other packets in the buffer, thus causing those other packets to be blocked. A higher priority packet arrives at the output buffer that is full and preempts a lower priority packet.

Packets Preempted

The number of packets that were blocked due to preemption. A lower priority packet has been ejected from the output buffer by a higher priority packet. A large, high priority packet can eject many smaller, lower priority packets from the output buffer.

Packets Exceeded

Packets that have been accepted at the output buffer when the output buffer threshold has been exceeded.

Packets Rejected

Packets that have been rejected at the output buffer when the output buffer threshold has been exceeded.

3. Links

Links represent communication channels between nodes. See Chapter 3 in the *COMNET III Reference Guide* for more information.

COMNET III produces the eleven different Link reports described in the following paragraphs.

The first column in most of the Link reports is titled:

LINK

or

LINK NAME

This column lists the name of the link.

3.1 Channel Utilization Report

LINKS: CHANNEL UTILIZATION

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK	FRAMES DELIVERED	FRAMES RST/ERR	TRANSMISSION DELAY (MS) AVERAGE	TRANSMISSION DELAY (MS) STD DEV	TRANSMISSION DELAY (MS) MAXIMUM	% UTIL
Ethernet Hub	5314	0	0.258	0.995	59.082	0.3707
Call Link 4						
FROM Switch	22	0	0.609	0.369	1.536	0.0045
FROM Call Node 4	25	0	0.391	0.023	0.469	0.0033
Call Link 1						
FROM Call Node 1	3561	0	0.557	0.099	0.604	0.6610
FROM Switch	3536	0	0.551	0.121	0.604	0.6496
Call Link 2						
FROM Call Node 2	520	0	0.907	0.595	1.536	0.1572
FROM Switch	699	0	0.244	0.032	0.313	0.0568

The Channel Utilization Report provides utilization rates for links used to carry connectionless and virtual circuit messages This report presents the number of link layer frames delivered and resent due to error, and transmission delays and link utilization.

Transmission delay is the time between when the frame (which may be part of a packet or contain several packets) is created at the input to the link and when the frame is delivered at the end of the link. It includes transmission, contention-resolution (for LANs), and propagation time.

From Node Name

The name of the transmitting node. For full duplex links transmissions can occur simultaneously and in both directions.

Frames Delivered

The number of frames removed from the output buffer at the transmitting node on the link and subsequently placed in the input buffer of the receiving node. Frames that are in transmission when the report is produced (because of transmission delay and propagation delay) are not reported.

Frames Resent/Error

The number of retransmitted frames. On a link, a framing error probability may be specified which causes statistically picked frames to be retransmitted as if they were in error. This feature is generally used to model noisy lines.

Transmission Delay—Average, Standard Deviation, Maximum

The average, standard deviation, and maximum delay observed for any packet across the link.

Link Utilization

The transmission time for a frame is calculated from its size divided by the link speed. The link is in use for this time. The additional propagation delay experienced by the packet does not utilize the link so that the link may transmit another packet while others are propagating. Utilization is the total usage time divided by the simulation run length.

3.2 Utilization By Application Report

LINKS: UTILIZATION BY APPLICATION

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK / APP TYPES	PACKETS DELIVERED	PKTS/ SEC	BYTES DELIVERED	KBPS DELIVERED	% BYTES	UTIL (%)
Ethernet Hub						
Other	4012	13.373	1134680	30.258	89.243	0.33
Novell NetWare	1009	3.363	73434	1.958	5.776	0.02
NNTP	30	0.100	42930	1.145	3.376	0.01
XNS-SPP	22	0.073	704	0.019	0.055	0.00
XNS-SMB	102	0.340	7729	0.206	0.608	0.00
XNS-NETBIOS	41	0.137	2549	0.068	0.200	0.00
NetBIOS Session	49	0.163	5224	0.139	0.411	0.00
Network Time	9	0.030	810	0.022	0.064	0.00
Domain	34	0.113	2820	0.075	0.222	0.00

The Utilization By Application Report provides a breakdown of link usage by application for connectionless and virtual circuit messages. For each category of packets flowing on a link, the report gives a count of the number of packets delivered, the average rate of delivery in kilobits per second (kbps), the percentage of bytes in the category, and the link utilization percentage produced by packets in the category.

A packet's application type is determined by the command (or source) that produced the packet.

In this report, if **Allocate background frames** is not checked, a separate line will be printed to indicate background frames. This formerly worked only with packet flows and Random link-loading. With Analytical link-loading, no frames are actually simulated, so in effect the "frames" are always allocated, no matter the state of the check box. This has been changed so that, if the allocation checkbox is unchecked, a line indicting effective background frames is printed.

App Types The name of the application utilizing the link.

Packets Delivered

The number of packets per application routed by the link.

Packets/Second

The number of packets per second, per application routed by the link.

Bytes Delivered

The total number of bytes delivered per application by the link.

KBPS Delivered

The number of kilobytes per second delivered by the link per application.

% Bytes The percentage of total bytes delivered by the link for each application using the link.

Utilization %

The percent utilization of the link per application.

3.3 Utilization By Protocol Report

LINKS: UTILIZATION BY PROTOCOL

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK / PROTOCOLS	PACKETS DELIVERED	PKTS/ SEC	BYTES DELIVERED	KBPS DELIVERED	% BYTES	UTIL (%)
Ethernet Hub						
Generic	295	0.983	295000	7.867	23.202	0.09
TCP/IP - Microso	3717	12.390	839680	22.391	66.041	0.24
IPX	1009	3.363	73434	1.958	5.776	0.02
IP	128	0.427	52360	1.396	4.118	0.02
XNS	165	0.550	10982	0.293	0.864	0.00
Call Link 4 FROM Switc						
IP	10	0.033	1515	0.040	58.926	0.00
XNS	12	0.040	1056	0.028	41.074	0.00
Call Link 4 FROM Call						
IP	2	0.007	180	0.005	9.590	0.00
XNS	23	0.077	1697	0.045	90.410	0.00
Call Link 1 FROM Call						
IPX	3561	11.870	380724	10.153	100.000	0.66
Call Link 1 FROM Switc						
IPX	3399	11.330	369790	9.861	98.828	0.64
XNS	137	0.457	4384	0.117	1.172	0.01

The Utilization By Protocol Report provides a breakdown of link usage by protocol. For each category of packets flowing on a link, the report gives a count of the number of packets delivered, the average rate of delivery in kilobits per second (kbps), the percentage of bytes in the category, and the link utilization percentage produced by packets in the category.

A packet's protocol is determined by the command (or source) that produced the packet.

Protocols The name of the protocol utilizing the link.

Packets Delivered

The number of packets per protocol routed by the link.

Packets/Second

The number of packets per second, per protocol routed by the link.

Bytes Delivered

The total number of bytes delivered per protocol by the link.

KBPS Delivered

The number of kilobytes per second, per protocol delivered by the link.

% Bytes The percent of total bytes delivered by the link for each protocol using the link.

Utilization %

The percent utilization of the link per protocol.

3.4 Frame Size Report

LINKS: FRAME SIZE

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK	COUNT	FRAME SIZES (BYTES) AVERAGE	STD DEV	MAXIMUM
Ethernet Hub	5314	261.567	406.866	1518.000
Call Link 4				
FROM Switch	22	116.864	70.790	295.000
FROM Call Node 4	25	75.080	4.507	90.000
Call Link 1				
FROM Call Node 1	3561	106.915	18.913	116.000
FROM Switch	3536	105.818	23.316	116.000
Call Link 2				
FROM Call Node 2	520	174.129	114.330	295.000
FROM Switch	699	46.770	6.053	60.000
Call Link 3				
FROM Call Node 3	679	49.353	15.158	112.000
FROM Switch	326	251.399	65.239	275.000

The Frame Size report provides a breakdown of link layer frames transmitted across the link. It presents the total number of link layer frames that traverse the link, along with their average, maximum, and standard deviation frame sizes. A packet's protocol is determined by the command (or source) that produced the packet.

From Node Name

The name of the transmitting node. For full duplex links transmissions can occur simultaneously and in both directions.

Count The number of link layer frames transmitted across the link.

Frame Sizes—Average, Standard Deviation, Maximum

The average, standard deviation and maximum size, in bytes, of a link layer frame transmitted across the link.

3.5 Collision Stats Report

```
                    LINKS: COLLISION STATS

              REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK NAME                  Ethernet
                             Hub

ACCESS PROTOCOL              CSMA/CD

COLLISION EPISODES              264

COLLIDED FRAMES                 529

NBR OF TRIES TO RESOLVE
  AVERAGE                      1.76
  STANDARD DEVIATION           1.40
  MAXIMUM                        10

NBR OF DEFERRALS                169

DEFERRAL DELAY  (MS)
  AVERAGE                      0.30
  STANDARD DEVIATION           0.42
  MAXIMUM                      1.19

DEFERRAL QUEUE SIZE  (FRAMES)
  AVERAGE                      0.00
  STANDARD DEVIATION           0.01
  MAXIMUM                         1

MULTIPLE COLLISION EPISODES
  NBR EPISODES                    1
  AVG PER EPISODE              3.00
  MAX PER EPISODE                 3
```

The Collision Stats Report provides a summary of collision statistics that occurred on each link. Only those links that can produce collision statistics will be reported.

Access Protocol

The protocol being used on the link.

Collision Episodes

How many times a collision occurred on the link; that is when two or more nodes try to transmit inside the same collision window.

Collided Frames

The total number of frames involved in collisions.

Number Of Retries To Resolve—Average, Standard Deviation, Maximum

The average, standard deviation, and maximum observed number of retries attempted before an initially collided frame was transmitted.

Number Of Deferrals

The number of transmission attempts deferred. If a node attempts to transmit a frame and sees the link busy, it defers its transmission until the link becomes idle (plus the contention interval).

Deferral Delay—Average, Standard Deviation, Maximum

The average, standard deviation and maximum deferral delay observed.

Deferral Queue Size(Frames)—Average, Standard Deviation, Maximum

The average, standard deviation and maximum deferral queue size observed.

Number Of Multiple Collision Episodes

The number of times that more than two frames are involved in a collision. When a collision occurs in the collision window, it may be between two or more frames arising from different nodes. The simple case is between just two frames, but in a heavily congested system it is possible that frames from more than two nodes are involved in the collision.

Collisions Per Multiple Collision Episode—Average, Maximum

The average and maximum number of frames colliding in one contention interval for multiple collision episodes.

3.6 Token Ring Stats Report

LINKS: TOKEN RING STATS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK	TOKEN ROTATION TIME (MS) AVERAGE	TOKEN ROTATION TIME (MS) STD DEV	NODES ACTIVE AVG	ACTIVE TIME (MS) AVG	FDDI RT/TTRT>2 COUNT
Token Ring	0.074	0.201	0.004	0.162	N/A

The Token Ring Stats Report provides summary statistics for token passing links on token activity and token rotation times.

Token Rotation Time (MS)—Average, Standard Deviation

The average and standard deviation of the time for the token to pass to all the nodes connected to the link.

Average Active Nodes

The number of active nodes holding and passing the token.

Average Active Time (MS)

The average time the token is held by a node attached to the link.

FDDI Rotation Time/TTRT >2 Count

The number of times that the token rotation time divided by the target token rotation time is greater than 2.

3.7 Connection Stats Report

LINKS: LINK CONNECTION STATS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

	BUSY CONNECTIONS			CONNECTION DELAY (MS)		
LINK	MAX	MEAN	SDEV	MEAN	STD DEV	MAXIMUM
FDMA Link	2	2.00	0.06	1000.000	0.000	1000.000

The Connection Stats Report provides a summary of connection and connection delay statistics that occurred on each link. This report will only record statistics for Dialup Point-to-Point, Modem Pool, FDMA, and TDMA links.

Busy Connections—Mean, Standard Deviation, Maximum

The mean, standard deviation and maximum of the number of busy link connections.

Connection Delay (MS)—Mean, Standard Deviation, Maximum

The mean, standard deviation and maximum link connection delay for link layer frames transmitted across the link

3.8 WAN Link Statistics Report

LINKS: WAN LINK STATISTICS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK	DE FRAMES CREATED	DE FRAMES DROP	BURST SIZE AVERAGE	BURST SIZE STD DEV	BURST SIZE MAXIMUM	BURST SIZE UNITS
WAN Link						
FROM WAN Workstation	0	0	4.298	5.208	19.744	kbits
FROM WAN Workstation	0	0	3.451	4.153	18.496	kbits

The WAN Link Statistics Report provides statistics on the number of discard eligible frames created and dropped on the WAN link, along with burst size statistics.

Discard Eligible Frames Created

The number of discard eligible frames created by the WAN link.

Discard Eligible Frames Dropped

The number of discard eligible frames dropped by the WAN link.

Burst Size—Average, Standard Deviation, Maximum

The average, standard deviation and maximum burst size of packets across the WAN link.

Burst Size Units

The units used to define the burst size.

3.9 Session Level Report

LINKS: SESSION LEVEL

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

	SESSIONS IN PROGRESS		
LINK	**AVERAGE**	**STD DEV**	**MAXIMUM**
LAN	**19.44**	**11.14**	**38**

The Session Level Report provides a summary of the number of sessions set up across a link. A link carries the packets that comprise a session. If a session is set up across a subnetwork that uses connection oriented routing then all packets of the session will follow the same route.

Sessions set up across subnetworks where a connectionless routing algorithm is in operation are not reported, even if some of the session packets are routed across the link. This is because the complete session is not associated with the particular link.

Sessions In Progress—Average, Standard Deviation, Maximum

The average, standard deviation, and maximum number of concurrent sessions in progress on this link.

3.10 Call Counts Report

LINKS: CALL COUNTS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK NAME	CALLS ATTEMPTD	CALLS BLOCK AVAIL	CALLS BLOCK TRAFF	BLOCK PROB	CALLS CARRIED	CALLS DISCON- NECTED	CALLS PRE- EMPTED
Call Link 4	47	0	15	0.32	32	0	0
Call Link 1	34	0	4	0.12	30	0	0
Call Link 2	44	0	11	0.25	33	0	0
Call Link 3	43	0	9	0.21	34	0	0

The Call Counts Report represents a summary of all the calls attempted to be transmitted across a link, and whether they were blocked, carried, disconnected, or preempted.

Calls Attempted

The number of calls that have been attempted in terms of routing call instances to their destination over this link. The origin and destination are not necessarily connected to this link—it may be an intermediate part of the route.

Calls Blocked - Availability

The number of calls that could not be routed over this link because this link had failed.

Calls Blocked - Traffic

The number of calls that could not be routed over this link because all its bandwidth had been allocated to other calls.

Blocking Probability

Total Calls Blocked divided by Calls Attempted for calls blocked and carried on this link.

Calls Carried

The number of calls that have been successfully routed to their destination over this link.

Calls Disconnected

The number of calls carried on this link and subsequently disconnected due to failure of this link.

Calls Preempted

The number of calls carried on this link and which were disconnected due to preemption by a higher priority call which requires bandwidth on this link. The preempting call does not have to be between the same origin and destination or following the same route as the preempted call. Even if there is only one common node or link between the calls then preemption may occur.

3.11 Call Level Report

LINKS: CALL LEVEL

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

LINK NAME	% AVAIL	LINK FAILS	BANDWIDTH USED (KBPS) AVERAGE	STD DEV	MAXIMUM	LINK UTIL %
Call Link 4	100.00	0	1187	498	1536	77.29
Call Link 1	100.00	0	1115	478	1536	72.59
Call Link 2	100.00	0	1115	478	1536	72.59
Call Link 3	100.00	0	1187	498	1536	77.29

The Call Level Report provides utilization rates and other statistics for each link's circuit switched call traffic.

Link Name The name of the link. Each link has a bandwidth capacity which it may use to carry call traffic. This is not shared with packet switched traffic.

Availability %

A link is not available when it is in the failed state. The percentage of time available is reported. This is the Up Time divided by the Run Length.

Link Failures

The number of link failures that occurred in the simulation.

Bandwidth Used (Kbps)—Average, Standard Deviation, Maximum

The average, standard deviation and maximum peak bandwidth used on the link during the simulation.

Link Utilization

Average Bandwidth Used/Total Available Bandwidth expressed as a percentage.

4. WAN Cloud

A WAN Cloud represents a set of communication channels, the fine details of which are unknown or unimportant to the simulation, typically of a commercially available service. Refer to Chapter 3 in the *COMNET III Reference Guide* for more information.

COMNET III produces the four WAN cloud reports described in the following paragraphs.

4.1 Frame Delay By VC Report

WAN CLOUDS: FRAME DELAY BY VC

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

CLOUD: VC	FRAME DELAY (MS) AVG	STD	MAX	BURST SIZE (kb) AVG	MAX
WAN Cloud					
VC41	0	0	0	0	0
VC42	47	59	166	6	10
VC43	0	0	0	0	0
VC44	31	42	166	3	10
VC48	0	0	0	0	0
VC49	54	64	166	6	10
VC50	0	0	0	0	0
VC51	37	50	166	5	10

The Frame Delay by VC Report presents statistics for the frame delay and burst size for each virtual circuit in a WAN cloud. The frame delay statistics are for all the frames that the WAN cloud successfully delivers. The burst size is monitored for each frame that the virtual circuit accepts, and thus it will capture instantaneous peak values.

Cloud The name of the WAN cloud.

Frame Delay (MS)—Average, Standard Deviation, Maximum

The average, standard deviation and maximum frame delay, in milliseconds, incurred by a link layer frame transiting the WAN cloud.

Burst Size (kb)—Average, Maximum

The average and maximum burst size, in kilobits, of packets transmitted across the virtual circuit.

4.2 Frame Counts By VC Report

WAN CLOUDS: FRAME COUNTS BY VC

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

CLOUD:		FRAMES / KILOBITS			
VC: FRAMES		ACCEPTED		DROPPED	
KILOBITS		NORMAL	DE	NORMAL	DE
WAN Cloud	(TOTAL KILOBITS TRANSMITTED =	744)			
VC41	Frm	0	0	0	0
	kb	0	0	0	0
VC42	Frm	84	0	0	0
	kb	163	0	0	0
VC43	Frm	0	0	0	0
	kb	0	0	0	0
VC44	Frm	84	0	0	0
	kb	83	0	0	0
VC48	Frm	0	0	0	0
	kb	0	0	0	0
VC49	Frm	84	0	0	0
	kb	195	0	0	0
VC50	Frm	0	0	0	0
	kb	0	0	0	0
VC51	Frm	84	0	0	0
	kb	115	0	0	0

The Cloud Throughput Report (Frame Counts by VC) provides throughput measures for each virtual circuit both in terms of the number of frames that the virtual circuit accepted or dropped and the number of kilobits accepted or dropped. These numbers are further divided between normal and discard eligible (DE) frames.

Cloud The name of the WAN cloud. The first line of data in the report identifies the WAN cloud that includes the following virtual circuits, and it lists the total number of kilobits that were successfully transmitted through the cloud.

Virtual Circuit (VC)

The name of the virtual circuit defined inside of the WAN cloud. For each virtual circuit, there are two rows of data. The first row lists the frames and the second row lists the kilobits. For each row there are separate counts for normal and DE frames and for frames accepted and dropped. In these counts, frames that are dropped due to the excess burst size being exceeded will be counted as DE frames. Also, the accepted frames may be dropped later when they arrive at a filled exit buffer.

4.3 Access Link Stats Report

WAN CLOUDS: ACCESS LINK STATS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

CLOUD: ACCESS LINK	(ENTRY) (EXIT)	FRAMES ACCEPTED	FRAMES DROPPED	BUFFER (BYTES) MAX	BUFFER (BYTES) AVG	BUFFER (BYTES) STD	% UTIL
WAN Cloud							
Router 2 Link	Entry	168	0	N/A	N/A	N/A	0.05
	Exit	0	0	0	0	0	0.00
Workstaion 1 Lin	Entry	141	0	N/A	N/A	N/A	1.78
	Exit	141	0	1080	9	96	1.16
Workstation 2 Li	Entry	141	0	N/A	N/A	N/A	1.63
	Exit	141	0	1080	11	104	1.30
Router 1 Link	Entry	0	0	N/A	N/A	N/A	0.00
	Exit	168	0	1040	1	27	0.08

The Access Link Stats Report presents statistics for each access link in a WAN cloud. For each access link, two rows are printed: the first is for the entry link and the second is for the exit link. No buffer is modeled for the entry link (the buffering here is in the port of the connected node) and thus those values are blank.

Cloud: The name of the WAN cloud.

Access Link The name of the access link in the WAN cloud. For each access link, two rows are printed: the first is for the entry link and the second is for the exit link.

Frames Accepted

The number of frames accepted by the entry and exit access link of the WAN cloud. For entry links, this count reflects the frames that found a path to the required destination. If a frame arrives that must go to a destination for which no virtual circuit is defined and the cloud does not allow transmission through non-VCs, then that frame will count as being dropped by the entry access link even though that frame's transmission will have utilized the link.

Frames Dropped

The number of frames dropped by the entry and exit access links of the WAN cloud. For exit links, the count reflects the number of frames accepted or blocked by the exit buffer. In cases where a low priority frame is initially accepted by the buffer and then preempted by a higher priority frame, then that low priority frame will count as being dropped and not accepted.

Buffer (Bytes)—Maximum, Average, Standard Deviation

The maximum, average and standard deviation of the buffer size for the entry and exit access link.

Utilization %

The entry and exit access link utilization in terms of percentage. The utilization is time-averaged over the replication, where at each event, the utilization is the number of busy circuits divided by the total number of circuits.

4.4 Cloud Access Buffer Policy Report

```
                WAN CLOUDS: CLOUD ACCESS BUFFER POLICY

               REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

CLOUD:                              FRAMES
  ACCESS LINK                  EPD          PPD
_____________________________  ___________  _________

WAN Cloud
  Router 2 Link                      0           0
  Workstaion 1 Lin                   0           0
  Workstation 2 Li                   0           0
  Router 1 Link                      0           0
```

The Cloud Access Buffer Policy Report is similar to the Buffer Policy Report on node buffers but it applies to the buffer policy on the cloud access exit (or egress) buffer.

Cloud: The name of the WAN cloud.

Access Link The name of the access link in the WAN cloud. For each access link, two rows are printed: the first is for the entry link and the second is for the exit link.

Early Packet Discard Frames

The number of frames that were flagged as early packet discard for each access link in the WAN Cloud.

Partial Packet Discard Frames

The number of frames that were flagged as partial packet discard for each access link in the WAN Cloud.

5. Application Sources

Application Sources are a flexible means of creating arbitrary network and processor workload. See Chapter 3 in the *COMNET III Reference Guide* for more information. COMNET III produces the three different Application Sources reports described in the following paragraphs.

The first column in most of the Application Sources reports is titled:

NODE: APPLICATION LIST

This column lists the name of the node and the names of the applications that run on the node.

5.1 Application Run Length Report

APPLICATION SOURCES: APPLICATION RUN LENGTH

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE: APPLICATION LIST	NUMBER COMPL	AVERAGE	APPLICATION DELAY MINIMUM	MAXIMUM	STD DEV
Application Workstati:					
Application Request	4	137.3808 S	67.1246 S	197.3535 S	47.7800 S
Application Server:					
Application Source	3	112.6697 S	67.4112 S	135.6193 S	32.0037 S

The Application Run Length Report presents a summary of delays experienced by applications executing on nodes. For each node, the list of applications running on that node is summarized. Application Delay is the average time required to execute the completed applications, counted from the point at which they were first scheduled until the point at which their last command completes. This includes time spent waiting in the application pending list of the node while other applications are executing.

Number Complete

The number of instances of the application that have completed execution of the last command in their command list. It is possible that, at the instant the report was produced, applications were still scheduled for the node. These applications will not be included on the report as they have not completed.

Application Delay—Average, Minimum, Maximum, Standard Deviation

The average, minimum, maximum, and standard deviation of the application delay for completed instances of the application.

5.2 Successfully Completed Applications Report

```
            APPLICATION SOURCES: SUCCESSFULLY COMPLETED APPLICATIONS

                    REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE:                    NUMBER              DELAY SUCCESSFULLY COMPLETING
  APPLICATION LIST       COMPL    AVERAGE    MINIMUM     MAXIMUM     STD DEV
_________________________ _____ __________ __________ __________ __________

Application Workstati:
  Application Request        4  137.3808 S   67.1246 S  197.3535 S   47.7800 S
Application Server:
  Application Source         3  112.6697 S   67.4112 S  135.6193 S   32.0037 S
```

The Successfully Completed Applications Report presents a summary of the instances of an application that completed successfully. For each node, the list of applications running on that node is summarized.

Application Delay is the time it required to execute the completed applications, counted from the point at which they were first scheduled until the point at which their last command completes. This includes time spent waiting in the application pending list of the node while other applications are executing.

Number Complete

The number of instances of the application that have successfully completed execution of the last command in their command list. It is possible that, at the instant the report was produced, applications were still scheduled for the node. These applications will not be included on the report as they have not completed.

Application Delay—Average, Minimum, Maximum, Standard Deviation

The average, minimum, maximum, and standard deviation of the application delay for completed instances of the application.

5.3 Applications Terminated By Time-Out Report

APPLICATION SOURCES: APPLICATIONS TERMINATED BY TIME-OUT

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

NODE: APPLICATION LIST	NUMBER TERMED	DELAY TIMING OUT AVERAGE	MINIMUM	MAXIMUM	STD DEV
Application Workstati:					
Application Request	0	0.000 MS	0.000 MS	0.000 MS	0.000 MS
Application Server:					
Application Source	0	0.000 MS	0.000 MS	0.000 MS	0.000 MS

The Applications Terminated By Time-Out Report presents a summary of the number of instances of an application that timed out before completion. For each node, the list of applications running on that node is summarized.

Application Delay is the time required to execute the timed-out applications, counted from the point at which they were first scheduled until the point at which they timed-out. This includes time spent waiting in the application pending list of the node while other applications are executing.

Number Terminated

The number of instances of the application that timed out before successfully completing execution. This means that they did not complete the last command in their command list. It is possible that, at the instant the report was produced, applications were still scheduled for the node. These applications will not be included on the report as they have not completed.

Application Delay—Average, Minimum, Maximum, Standard Deviation

The average, minimum, maximum and standard deviation of the application delay for timed out instances of the application.

6. Message and Response Sources

Message and Response Sources create packet-switched workload. Response Sources execute in response to a received message, typically from a Message Source, and automatically direct their traffic to the originator of that message. See Chapter 3 in the *COMNET III Reference Guide* for more information.

The eleven reports associated with Message and Response Sources are presented in the following paragraphs.

For each node in the model sending a Response Source message, the report lists **ECHO** as the destination list destination. **ECHO** indicates that the delay messages assembled and delay statistics are for all message source messages that triggered a response message. Thus, the **ECHO** line statistics display an aggregate number of messages assembled along with average delays. In order to see a breakdown of the messages assembled and the message delay statistics for each response source message being sent to each destination, it is necessary to deselect the **ECHO** checkbox on the **Destinations** tab window for the response source parameter set. If the checkbox is off, then the specific destinations will be listed for the response source, but the **ECHO** destination will remain in the report to catch unexpected destinations.

The first column in the Message and Response Sources reports is titled:

ORIGIN/MSG SRC NAME: DESTINATION LIST

or

ORIGIN : DESTINATION LIST

In the first instance the column lists the name of the node and the name of the Message and/or Response Source for each node that has Message and/or Response sources attached to it.

In the second instance the column lists the name of the node and the name of the source for each node that has Message and/or Response sources attached to it.

In both instances, the column also lists the destination of the Message Source. All destinations in the destination list of the source are listed. For Random Neighbor destinations, all nodes connected via one link to the originating node are listed.

6.1 Message Delay Report

MESSAGE + RESPONSE SOURCES: MESSAGE DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / MSG SRC NAME: DESTINATION LIST	MESSAGES ASSEMBLED	MESSAGE DELAY AVERAGE	STD DEV	MAXIMUM
Message Workstation / src Message Source:				
Response Server	31	202.150 MS	0.102 MS	202.709 MS
Response Server / src Message Response:				
ECHO	49	371.396 MS	219.525 MS	659.637 MS
Session Workstation / src Session Response:				
ECHO	295	2.799 MS	0.145 MS	4.472 MS
WAN Workstation 1 / src WAN Message 1:				
Response Server	11	456.388 MS	0.062 MS	456.584 MS
Session Server	6	464.434 MS	18.037 MS	504.766 MS
WAN Workstation 2	13	653.454 MS	36.681 MS	743.857 MS
WAN Workstation 2 / src WAN Message 2:				
Response Server	7	456.368 MS	0.000 MS	456.368 MS
WAN Workstation 1	6	701.309 MS	107.529 MS	926.997 MS
Session Server	14	474.131 MS	42.362 MS	578.573 MS

This Message Delay Report presents message delay statistics as perceived by the sender. Message Delay is the time between creating the first packet of the message on the originating node and the time that the originating node is notified that the message has been assembled by the destination. For each originating node in the model the report lists delays to each destination.

Messages Assembled

For each destination, the number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination.

Message Delay—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of the message delay.

6.2 Message Delivered Report

MESSAGE + RESPONSE SOURCES: MESSAGE DELIVERED

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / MSG SRC NAME: DESTINATION LIST	MESSAGES ASSEMBLED	MESSAGE DELAY AVERAGE	STD DEV	MAXIMUM
Message Workstation / src Message Source:				
Response Server	31	1.977 MS	0.102 MS	2.536 MS
Response Server / src Message Response:				
ECHO	49	151.719 MS	193.928 MS	406.218 MS
Session Workstation / src Session Response:				
ECHO	295	2.799 MS	0.145 MS	4.472 MS
WAN Workstation 1 / src WAN Message 1:				
Response Server	11	203.109 MS	0.000 MS	203.109 MS
Session Server	6	203.109 MS	0.000 MS	203.109 MS
WAN Workstation 2	13	361.646 MS	1.249 MS	365.973 MS
WAN Workstation 2 / src WAN Message 2:				
Response Server	7	203.109 MS	0.000 MS	203.109 MS
WAN Workstation 1	6	409.809 MS	108.501 MS	652.425 MS
Session Server	14	203.521 MS	1.484 MS	208.873 MS

The Message Delivery Report presents statistics on the delay before the message is reassembled by the destination; that is the time between creating the first packet of the message on the originating node and the time of receiving the last packet on the destination node. The message may still be worked on by the source after the message is delivered because of retransmissions, waiting for acks, or for modeling the close-sequence of the connection, but these values are not factored into the message delivered delay.

Messages Assembled

For each destination, the number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination.

Message Delay—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of message delay.

6.3 Packet Delay Report

MESSAGE + RESPONSE SOURCES: PACKET DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN:	NUMBER OF PACKETS				PACKET DELAY (MS)	
DESTINATION LIST	CREATED	DELIVERED	RESENT	DROPPED	AVERAGE	MAXIMUM
Session Workstation / src Session Response:						
ECHO	295	295	0	0	0.826	1.035
WAN Workstation 1 / src WAN Message 1:						
Response Server	33	11	0	0	67.366	166.600
Session Server	18	6	0	0	69.722	166.600
WAN Workstation 2	39	13	0	0	120.695	315.259

The Packet Delay Report presents a summary of the number of packets created, delivered, resent, or dropped for each message and response source. It also provides the average and maximum packet delay. Packet Delay is the time between creating a packet on the originating node and the time of receiving the packet at the destination node.

Packets Created

The number of packets created to send to the listed destination. Not all packets need to be delivered before they appear on the report.

Packets Delivered

The number of packets received at the destination. This may differ from the number of packets created by the number of packets that are in transit at the instant the report is written.

Packets Retransmitted

The number of packets retransmitted. Packets may be retransmitted from the origin because they are blocked at some point in route to the destination. Blocking can occur because input or output buffers are full, or because a node or link on the route fails.

Packets Dropped

The number of packets dropped. When a node or link fails, you can specify whether the transmission should be reattempted. If so, the retransmissions will be counted in the **Packets Retransmitted** field. Otherwise, the packet is dropped.

Packet Delay (millisecs)—Average, Maximum

The average and maximum packet delay.

6.4 Window Stats Report

MESSAGE + RESPONSE SOURCES: WINDOW STATS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN: DESTINATION LIST	WINDOW AVG	WINDOW MAX	NO. OF RESETS	CONG AVOID	PKT CREATION INTERVAL (MS) AVG	PKT CREATION INTERVAL (MS) MAX	PKT CREATION INTERVAL (MS) STD DEV
Message Workstation / src Message Source:							
Response Server	2.00	2	0	0	1	1	0
Response Server / src Message Response:							
ECHO	2.00	2	0	0	14	36	17
Session Workstation / src Session Response:							
ECHO	0.00	0	0	0	0	0	0
WAN Workstation 1 / src WAN Message 1:							
Response Server	2.00	2	0	0	36	36	0
Session Server	2.00	2	0	0	36	36	0
WAN Workstation 2	2.00	2	0	0	50	50	0
WAN Workstation 2 / src WAN Message 2:							
Response Server	2.00	2	0	0	36	36	0
WAN Workstation 1	2.00	2	0	0	99	341	108
Session Server	2.00	2	0	0	36	36	0

The Window Stats Report presents the statistics on the TCP/IP congestion window that varies as the connection starts up and when congestion is detected. It also presents the statistics for the interval between when packets are created. This statistic is applicable to all protocols.

Window Size—Average, Maximum

The average and maximum size of the TCP/IP congestion window, in packets, that varies as the connection starts up and when congestion is detected.

Number Of Congestion Resets

The number of times when the TCP congestion window is reset to "1" due to congestion (multiple dropped packets).

Number Of Congestion Avoidances

The number of times the TCP congestion window enters congestion avoidance where the window grows linearly instead of by the slow-start algorithm used initially.

Packet Creation Interval (MS)—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of the packet creation interval.

6.5 Retransmissions Report

MESSAGE + RESPONSE SOURCES: RETRANSMISSIONS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN: DESTINATION LIST	NO. OF EVENTS	NO. PACKETS AVG	NO. PACKETS MAX	PKT RETRIES AVG	PKT RETRIES MAX	BLCKD AFTER T-O PKTS	BLCKD AFTER T-O ACKS
Message Workstation / src Message Source:							
Response Server	0	0.00	0	0.00	0	0	0
Response Server / src Message Response:							
ECHO	0	0.00	0	0.00	0	0	0
Session Workstation / src Session Response:							
ECHO	0	0.00	0	0.00	0	0	0
WAN Workstation 1 / src WAN Message 1:							
Response Server	0	0.00	0	0.00	0	0	0
Session Server	0	0.00	0	0.00	0	0	0
WAN Workstation 2	0	0.00	0	0.00	0	0	0
WAN Workstation 2 / src WAN Message 2:							
Response Server	0	0.00	0	0.00	0	0	0
WAN Workstation 1	0	0.00	0	0.00	0	0	0
Session Server	0	0.00	0	0.00	0	0	0

The Retransmissions Report presents retransmission statistics when blocked packets are retransmitted. The retransmission timer accounts for the age of the packet before it is blocked. It schedules a packet retransmission only after the packet is actually blocked somewhere in the network. The retransmission time is adjusted so that the packet is retransmitted the specified delay after the packet was originally transmitted. It reports on quantities such as the number of times a packet has to be retransmitted, and the number of packets that are retransmitted when a blocked packet occurs.

Number Of Events

The number of times a blockage necessitates the retransmission of a packet or packets.

Number Packets—Average, Maximum

The average and maximum number of packets retransmitted.

Packet Retries—Average, Maximum

The average and maximum number of times a packet was resent.

Packets Blocked After T-O

The number of packets blocked after timing out.

Acknowledgments Blocked After T-O

The number of acknowledgments blocked after timing out.

6.6 Timeout Report

MESSAGE + RESPONSE SOURCES: TIMEOUT

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN: DESTINATION LIST	RETRANSMIT TIMEOUT (MS) AVG	RETRANSMIT TIMEOUT (MS) MAX	ROUND TRIP TIME (MS) AVG	ROUND TRIP TIME (MS) MAX
Message Workstation / src Message Source:				
Response Server	0.000	0.000	200.898	200.901
Response Server / src Message Response:				
ECHO	0.000	0.000	268.291	384.510
Session Workstation / src Session Response:				
ECHO	0.000	0.000	0.000	0.000
WAN Workstation 1 / src WAN Message 1:				
Response Server	0.000	0.000	384.369	384.565
Session Server	0.000	0.000	384.349	384.349
WAN Workstation 2	0.000	0.000	543.229	632.149
WAN Workstation 2 / src WAN Message 2:				
Response Server	0.000	0.000	384.349	384.349
WAN Workstation 1	0.000	0.000	535.429	535.429
Session Server	0.000	0.000	384.761	390.113

The Timeout Report presents time-out timer statistics for message and response sources. This report is most relevant for enhanced sliding window or TCP/IP window protocols that adapt their time-out timers based on measured round-trip delays. Because the sliding window algorithms acknowledge each packet, they have an opportunity to measure round-trip packet delay and use that delay to adjust the retransmission time out timer. The timer is only sampled when blocked packets are retransmitted.

Retransmit Timeout (MS)—Average, Maximum

The average and maximum timeout before a retransmission of the packet.

Round Trip Time (MS)—Average, Maximum

The average and maximum round trip delay time for a packet to get to a destination. This time is measured from when a packet is transmitted from the source node to when its corresponding ack is received back at the source node.

6.7 Ack Delay Report

MESSAGE + RESPONSE SOURCES: ACK DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN: DESTINATION LIST	NUMBER OF ACKS CREATED	DROPPED	ACK DELAY (MS) AVERAGE	MAXIMUM	STD DEV
Message Workstation / src Message Source:					
Response Server	93	0	200.898	200.901	0.001
Response Server / src Message Response:					
ECHO	147	0	268.291	384.510	88.443
Session Workstation / src Session Response:					
ECHO	0	0	0.000	0.000	0.000
WAN Workstation 1 / src WAN Message 1:					
Response Server	33	0	384.369	384.565	0.062
Session Server	18	0	384.349	384.349	0.000
WAN Workstation 2	39	0	543.229	632.149	25.699
WAN Workstation 2 / src WAN Message 2:					
Response Server	21	0	384.349	384.349	0.000
WAN Workstation 1	18	0	535.429	535.429	0.000
Session Server	42	0	384.761	390.113	1.484

The Ack Delay Report complements the Packet Delay Report to give statistics on the acks created and dropped as well as the delay for the ack. The ack delay is a round trip delay from when the packet started to when the ack returned. The time starts with the last packet that results in sending the ack.

Number of Acks Created

The number of acks created in response to the source message.

Number of Acks Dropped

The number of acks dropped in response to the source message.

Ack Delay—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of the ack delay.

6.8 Packet Size Report

MESSAGE + RESPONSE SOURCES: PKT SIZE

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN: DESTINATION LIST	MAX PKT	MAX WINDOW	PACKET SIZE MIN	AVG	MAX
Message Workstation / src Message Source:					
Response Server	1460	3	1000	1000	1000
Response Server / src Message Response:					
ECHO	1460	3	1000	1000	1000
Session Workstation / src Session Response:					
ECHO	65535	1	1000	1000	1000
WAN Workstation 1 / src WAN Message 1:					
Response Server	1460	3	1000	1000	1000
Session Server	1460	3	1000	1000	1000
WAN Workstation 2	1460	3	1000	1000	1000
WAN Workstation 2 / src WAN Message 2:					
Response Server	1460	3	1000	1000	1000
WAN Workstation 1	1460	3	1000	1000	1000
Session Server	1460	3	1000	1000	1000

The Packet Size Report provides two functions:

- The first function is to report on the maximum packet and window size resulting from the socket constraints at either end. The size is set by the transport protocol plus any sockets that may be at either end of the connection. This is important because the sockets' packet or window size constraint can reduce the maximum protocol window or packet size.
- The second function is to report on the packet sizes that were actually created. This part is most useful for monitoring the sizes for messages that are smaller than the packet sizes or for monitoring the sizes from external traffic sources.

Packet Maximum

The maximum size of a packet allowed by the connection.

Window Maximum

The maximum window size, in packets, allowed by the connection.

Packet Size—Minimum, Average, Max

The minimum, maximum and average size of the packet.

6.9 Burst Size Report

MESSAGE + RESPONSE SOURCES: BURST SIZE

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN: DESTINATION LIST	BURST UNITS	TRAF TYPE	BURST SIZE MIN	MAX	AVG	STD DEV
Workstation 1 / src Client 1:						
Workstation 2	kBits	FRLB	6	56	31	0
Workstation 3	kBits	FRLB	4	56	31	0
Workstation 1 / src Client 4:						
Workstation 2	kBits	FRLB	2	56	32	0
Workstation 3	kBits	FRLB	8	56	30	0

The Burst Size Report provides statistics on the burst size measured by the traffic policy on the protocol, when the traffic policy is present. It is the burst measurement that is used to determine the DE status of the packet, or for the ATM traffic policy, it is the burst measurement used to determine the conformance of the packets. Thus, a determination is made whether the packets should be immediately dropped because they exceed their traffic contract. There are three burst measurement options implemented for modeling the traffic policing. The three measurements are Sliding Window, Jumping Window, and Leaky Bucket. The burst algorithms are set in the transport protocol under traffic policing.

The Sliding Window algorithm measures the burst as the total size of the packets that arrived during a fixed interval before the current time. This interval slides according to the time, and the measurement always includes all the packets that have arrived for the previous full interval.

The Jumping Window algorithm measures the burst based on accumulating a burst size with each packet, but then clearing that accumulation at fixed intervals.

The Leaky Bucket algorithm (or the Generalized Cell Rate Algorithm, GCRA for ATM) is similar to the jumping window in terms of allowing excess traffic, except that it periodically subtracts off an amount equal to the rate times the burst interval. This maintains some memory of the previous burst interval, especially if that burst exceeded the value of the burst rate times the burst interval.

Burst Units The type of units set for the burst size under the traffic policing option of the transport protocol.

Traffic Type The following types of traffic can be represented in this report:

FRLB - Frame Relay Leaky Bucket

FRSW - Frame Relay Sliding Window

FRJW - Frame Relay Jumping Window

ATM - ATM

NONE - (reserved)

Burst Size—Minimum, Maximum, Average and Standard Deviation

The minimum, maximum, average and standard deviation of the burst size.

6.10 Assembly Interval Report

MESSAGE + RESPONSE SOURCES: ASSY INTVL

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN: DESTINATION LIST	PKTS ASS'D	INTER-ASSY INTERVAL MIN	MAX	AVG	STD DEV	UNIT
Workstation 1/ src Client 1:						
Workstation 2	189	114	130666	13813	16346	uS
Workstation 3	270	101	75953	12990	12101	uS
Workstation 2 / src Client 2:						
Workstation 1	270	96	118073	14475	16223	uS
Workstation 3	252	32	152623	13376	18070	uS
Workstation 3 / src Client 3:						
Workstation 1	216	110	58625	13981	11790	uS
Workstation 2	270	45	146084	13426	15964	uS

The Assembly Interval Report provides statistics on the delay between the time when different packets from the same message are assembled at the destination. The assembly interval can result from delay variation through the network (due to varying buffer delays) or it can result from congestion on the destination node. In cell-based networks, this measure is the cell-delay variation, a quality of service measure for these services.

Packets Assembled

The number of packets assembled at the destination.

Inter-Assembly Interval—Minimum, Maximum, Average and Standard Deviation

The minimum, maximum, average and standard deviation for the delay between the packets when they are assembled at the destination.

Inter-Assembly Units

The time unit used to measure the delay.

6.11 Discard Eligible Packets Report

MESSAGE + RESPONSE SOURCES: DE PACKETS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN: DESTINATION LIST	DE 0 PACKETS CREATED	DE 0 PACKETS DROPPED	DE 1 PACKETS CREATED	DE 1 PACKETS DROPPED	CONGEST PACKETS
Message Workstation / src Message Source:					
Response Server	93	0	0	0	0
Response Server / src Message Response:					
ECHO	147	0	0	0	0
Session Workstation / src Session Response:					
ECHO	295	0	0	0	0
WAN Workstation 1 / src WAN Message 1:					
Response Server	33	0	0	0	0
Session Server	18	0	0	0	0
WAN Workstation 2	39	0	0	0	0
WAN Workstation 2 / src WAN Message 2:					
Response Server	21	0	0	0	0
WAN Workstation 1	18	0	0	0	0
Session Server	42	0	0	0	0

The Discard Eligible Packets Report provides information about the flags set on the packet: the DE (discard eligibility) flag from the protocol's traffic policing algorithm, and the congestion flags set by the buffers when they are set up for FECN or BECN.

Discard Eligible 0 Packets Created

The number of created packets flagged with a discard eligibility of 0.

Discard Eligible 0 Packets Dropped

The number of dropped packets flagged with a discard eligibility of 0.

Discard Eligible 1 Packets Created

The number of created packets flagged with a discard eligibility of 1.

Discard Eligible 1 Packets Dropped

The number of dropped packets flagged with a discard eligibility of 1.

Congestion Packets

The number of packets that had a congestion flag set by buffers. This occurs when the buffers are set up for FECN or BECN.

7. Call Sources

Call Sources create circuit-switched workload. See Chapter 3 in the *COMNET III Reference Guide* for more information.

COMNET III produces the three different Call Sources reports described in the following paragraphs.

The first column in the Call Source reports is titled:

ORIGIN/CALL NAME: DESTINATION LIST

This column lists the name of the node and the name of the source for each node that has Call Sources attached to it and also lists the destination of the call. A Call Source will establish circuit switched routes across the network to the destination and hold the bandwidth on the route for the call holding time. All destinations in the destination list of the source are listed. For Random Neighbor destinations, all nodes connected via one link to the originating node are listed.

Note: A call source may not be multicast.

7.1 Blocked Call Counts Report

CALL SOURCES: BLOCKED CALL COUNTS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / CALL NAME: DESTINATION LIST	CALLS ATTEMPTD	CALLS RETRY	BLOCK PROB	HOPS AVG	HOPS MAX
Call Node 4 / call Call 4:					
Call Node 3	35	0	0.457	2.0	2
SUBTOTAL	35	0	0.457	2.0	2
Call Node 4 (TOTAL)	35	0	0.457	2.0	2
Call Node 1 / call Call 1:					
Call Node 2	21	0	0.190	2.0	2
SUBTOTAL	21	0	0.190	2.0	2
SUBTOTAL	24	0	0.500	2.0	2
Call Node 3 (TOTAL)	24	0	0.500	2.0	2
** T O T A L S **	107	0	0.430	2.0	2

The Blocked Call Counts Report gives details on calls that were blocked and the average and maximum hops needed to complete calls. For each origin and call name, the report presents information for each connected destination.

Calls Attempted

The number of calls that have been attempted in terms of routing new calls to their destination.

Calls Retried

The number of calls retried. A call routing attempt may fail because no route can be found where all links or all nodes have sufficient remaining, operational capacity. If the call cannot be routed, you can specify whether a call should be retried later, or dropped.

Calls Blocked

The number of calls blocked due to insufficient routing capacity.

Average Hops

The average number of hops (or links) used across all routes. Different calls may use different routes to reach the destination.

Maximum Hops

The longest route used in hops (or links).

7.2 Disconnected Call Counts Report

CALL SOURCES: DISCONNECTED CALL COUNTS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / CALL NAME: DESTINATION LIST	PRI	CALLS ATTEMPTD	CALLS CARRIED	CALLS DISCON	CALLS REROUT
Call Node 4 / call Call 4:	1				
Call Node 3		35	19	0	0
SUBTOTAL		35	19	0	0
Call Node 2 (TOTAL)		27	13	0	0
Call Node 3 / call Call 3:	1				
Call Node 4		24	12	0	0
SUBTOTAL		24	12	0	0
Call Node 3 (TOTAL)		24	12	0	0
** T O T A L S **		107	61	0	0

The Disconnected Call Counts Report gives details on calls that were disconnected and rerouted for each origin, call name, and connected destination.

Priority The call source priority. When preemptive operation has been specified, and a high priority call is routed and no route is available, the call will attempt to preempt a lower priority call on the full link.

Calls Attempted

The number of calls that have been attempted in terms of routing new calls to their destination.

Calls Carried

The number of calls successfully routed to their destination.

Calls Disconnected

The number of calls disconnected due to link or node failure.

Calls Rerouted

The number of disconnected calls successfully rerouted to their destination.

7.3 Preempted Call Counts Report

CALL SOURCES: PREEMPTED CALL COUNTS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / CALL NAME: DESTINATION LIST	PRI	CALLS ATTEMPTED	CALLS CARRIED	CALLS PREEMPTED
Call Node 4 / call Call 4:	1			
Call Node 3		35	19	0
SUBTOTAL		35	19	0
Call Node 2 (TOTAL)		27	13	0
Call Node 3 / call Call 3:	1			
Call Node 4		24	12	0
SUBTOTAL		24	12	0
Call Node 3 (TOTAL)		24	12	0
** T O T A L S **		107	61	0

The Preempted Call Counts Report gives details on preempted calls for each origin, call name, and connected destination. The preempting call does not have to be between the same origin and destination or following the same route as the preempted call. Even if there is only one common node or link between the calls, then preemption may occur.

Priority The call source priority. When preemptive operation has been specified, and a high priority call is routed and no route is available, the call will attempt to preempt a lower priority call on the full link.

Calls Attempted

The number of calls attempted in terms of routing new calls to their destination.

Calls Carried

The number of calls successfully routed to their destination.

Calls Disconnected

The number of calls disconnected due to link or node failure on some part of their route.

Calls Preempted

The number of calls disconnected due to preemption by a higher priority call which requires bandwidth on one (or more) of the nodes or links used by the preempted call.

8. Session Sources

Session Sources establish a pattern of packet switched traffic, involving optional handshaking messages, and including any number of content messages, between the same source and destination. See Chapter 3 in the *COMNET III Reference Guide* for more information.

COMNET III produces the fourteen different reports described in the following paragraphs.

The first column in most of the Session Sources reports is titled:

ORIGIN/SESSION SRC: DESTINATION LIST

This column lists the name of the node and the name of the source for each node that has Session Sources attached to it. Each source will send messages to the destinations listed in its destination list. All destinations in the destination list of the source are listed. For Random Neighbor destinations, all nodes connected via one link to the originating node are listed.

8.1 Message Delay Report

SESSION SOURCES: MESSAGE DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	MESSAGES ASSEMBLED	MESSAGE DELAY AVERAGE	STD DEV	MAXIMUM
Session Server / src Session Source:				
Session Workstation	295	202.150 MS	0.161 MS	203.821 MS

The Message Delay Report presents message delay statistics for session sources. Message delay is the time between creating the first packet of the message on the originating node, and the time that the originating node is notified that the message has been assembled by the destination. This is the message delay that is perceived by the sender. For each originating node in the model it lists delays to each destination.

Messages Assembled

For each destination, the number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of the observed message delay.

8.2 Message Delivered Report

SESSION SOURCES: MESSAGE DELIVERED

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	MESSAGES ASSEMBLED	MESSAGE DELAY AVERAGE	STD DEV	MAXIMUM
Session Server / src Session Source:				
Session Workstation	295	1.973 MS	0.144 MS	3.648 MS

The Message Delivered Report presents statistics for session sources on the message delay before the message is reassembled by the destination. Thus, message delay is the time between creating the first packet of the message on the originating node and the time of receiving the last packet of the message on the destination node. The message may still be worked on by the source after the message is delivered because of retransmissions, waiting for acks, or for modeling the close-sequence of the connection, but these values are not factored into the message delivered delay.

Messages Assembled

The number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of the observed message delay.

8.3 Packet Delay Report

SESSION SOURCES: PACKET DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	NUMBER OF PACKETS CREATED	DELIVERED	RESENT	DROPPED	PACKET DELAY (MS) AVERAGE	MAXIMUM
Session Server / src Session Source:						
Session Workstation	**885**	**295**	**0**	**0**	**0.324**	**1.525**

The Packet Delay Report presents a summary for session sources of the number of packets that were created, delivered, resent, or dropped. An average and maximum packet delay is also displayed. Packet Delay is the time between creating a packet on the originating node and the time of receiving the packet at the destination node.

Packets Created

The number of packets created to send to the listed destination. Not all packets need to be delivered before they appear on the report.

Packets Delivered

The number of packets received. This may differ from the number of packets created by the number of packets that are in transit at the instant the report is written.

Packets Resent

The number of packets resent. Packets may be retransmitted from the origin because they are blocked at some point in route to the destination. Blocking can occur because input or output buffers are full, or because a node or link on the route fails.

Packets Dropped

The number of packets dropped. When a node or link fails, you can specify whether the transmission should be reattempted. If so the retransmissions will be counted in the **Packets Resent** field. If no retransmission is specified, then the packet will be dropped.

Packet Delay—Average, Maximum

The average and maximum packet delay.

8.4 Setup Delay Report

SESSION SOURCES: SETUP DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	SESSIONS SETUP	SETUP DELAY (MILLISECONDS) AVERAGE	STD DEV	MAXIMUM
Session Server / src Session Source:				
Session Workstation	20	0.112 MS	0.000 MS	0.112 MS

The Setup Delay Report presents summary statistics for session sources on the time required to set up a session between an originating node and a destination node. The number of sessions that were set up and delay statistics are reported. Setup Delay is the time difference between creating the session setup packet and receiving back the session connect packet.

Sessions Setup

The number of sessions set up. When a session is started a session setup packet is sent to the destination and a session connect packet is returned. The session is then counted as set up.

Setup Delay—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of setup delay.

8.5 Session Length Report

SESSION SOURCES: SESSION LENGTH

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN	SESSION SRC NAME	SESSIONS ENDED	SESSION LENGTH (SECONDS) AVERAGE	STD DEV	MAXIMUM
Session Server	Session Source	19	13.8784 S	4.4433 S	21.6407 S

The Session Length Report presents summary statistics for session sources on the duration of a session between an originating node and a destination node. It provides the number of sessions completed and session length statistics.

Session Length is the time between the session setup packet being created and the end of the session. A session is counted as complete when the last message of a session has been received at the destination, and all packets for response messages have been received at the session origination node, and all acks for all packets have been received, and all pending message notices created by the receipt of the session messages have been cleared. When these conditions are met the session is complete and cleared.

Origin The node name which originates the session by setup command.

Setup Command Name

The name of the setup command.

Sessions Ended

The number of sessions that have ended naturally is reported.

Session Length (seconds)—Average, Standard Deviation, Maximum

The average, maximum, and standard deviation of the length of completed sessions.

8.6 Setup Counts Report

```
                    SESSION SOURCES: SETUP COUNTS

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC NAME:                    NUMBER OF SESSIONS
  DESTINATION LIST                 TRIED    SETUP   RETRY  BLOCK  DISCON RERTD
_____________________________________ _______ _______ _____ _____ _____ _____

Session Server / src Session Source:
  Session Workstation                20       20      0      0      0      0
```

The Setup Counts Report presents summary statistics for session sources on the number of attempts to set up a session, the number of attempts that succeeded, the number of retries to set up a session based upon initial failures, the number of sessions attempted that were blocked, the number of sessions attempted that were disconnected, and the number of sessions that were rerouted due to being blocked.

Origin The node name which originates the session setup command.

Setup Command Name

The name of the session setup command.

Sessions Tried

The number of attempts to set up sessions. As the simulation executes the session traffic source, an attempt is made to set up a particular session. Whether this session succeeds or fails depends upon network conditions such as route availability, inside hop and session limits and buffer availability.

Sessions Setup

The number of successful session setup attempts.

Sessions Retried

The number of sessions that were retried. If a setup attempt fails, the session may be retried later depending on the settings you have entered.

Sessions Blocked

The number of blocked sessions. The session setup attempt may block due to insufficient buffer space on routing nodes, or no route being available inside the hop limit, or no route being available because at least 1 link on all routes is at its session limit, or nodes/links have failed and no route is available.

Sessions Disconnected

The number of disconnected sessions. A session in progress may be disconnected because a node or link through which it is routed fails. The session may optionally be rerouted if this happens.

Sessions Rerouted

The number of sessions rerouted following disconnection.

8.7 Window Stats Report

```
                    SESSION SOURCES: WINDOW STATS

                 REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC:     WINDOW     NO. OF  CONG     PKT CREATION INTERVAL (MS)
  DESTINATION LIST       AVG    MAX  RESETS  AVOID     AVG      MAX      STD DEV
_____________________  _______  ___  _______  _______  _______  _______  _______

Session Server / src Session Source:
  Session Workstation     2.00    2       0       0         1        3        0
```

The Window Stats Report presents the statistics for session sources on the TCP/IP congestion window that varies as the connection starts up and when congestion is detected. It also presents the statistics for the interval between packets being created.

Window Size—Average, Maximum

The average and maximum size, in packets, of the TCP/IP congestion window.

Number Of Congestion Resets

The number of times that the TCP/IP "congestion" window (dynamic window size) is reset to 1 packet. Thus, congestion window size is reset.

Number Of Congestion Avoidances

After a reset, the window size starts to grow more slowly (linearly) than the initial "slow start" window growth. This algorithm switch is counted as "Congestion Avoidance" and gives a measure of how often the network is recovering from congestion. If Congestion Avoidance is nearly the same as the Number of Congestion Resets, then congestion may be spurious or short-lived. If there are lots of Congestion Resets and few Congestion Avoidances, then either the network is very congested or the message sizes are not very large.

Packet Creation Interval (MS)—Average, Maximum, Standard Deviation

The average, maximum, and standard deviation of the time interval between the creation of packets.

8.8 Retransmissions Report

SESSION SOURCES: RETRANSMISSIONS

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	NO. OF EVENTS	NO. PACKETS AVG	NO. PACKETS MAX	PKT RETRIES AVG	PKT RETRIES MAX	BLCKD AFTER T-O PKTS	BLCKD AFTER T-O ACKS
Workstation 1 / src Sess17:							
Workstation 3	225	1.00	1	0.00	0	0	0
Workstation 2	108	1.00	1	5.00	5	0	1

The Retransmissions Report report presents retransmission statistics for session sources when blocked packets are retransmitted. The retransmission timer accounts for the age of the packet before it is blocked. It schedules a packet retransmission only after the packet is actually blocked somewhere in the network. The retransmission time is adjusted so that the packet is retransmitted the specified delay after the packet was originally transmitted. It reports on quantities such as the number of times a packet has to be retransmitted, and the number of packets that are retransmitted when a blocked packet occurs.

Number Of Events

The number of retransmissions.

Number Packets—Average, Maximum

The average and maximum number of packets retransmitted.

Packet Retries—Average, Maximum

The average and maximum number of times a packet was resent.

Packets Blocked After T-O

The number of packets blocked after timing out.

Acknowledgments Blocked After T-O

The number of acknowledgments blocked after timing out.

8.9 Timeout Report

SESSION SOURCES: TIMEOUT

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	RETRANSMIT TIMEOUT (MS) AVG	RETRANSMIT TIMEOUT (MS) MAX	ROUND TRIP TIME (MS) AVG	ROUND TRIP TIME (MS) MAX
Session Server / src Session Source:				
Session Workstation	0.000	0.000	200.899	201.327

The Timeout Report presents time-out timer statistics for session sources. This report is most relevant for enhanced sliding window or TCP/IP window protocols that adapt their time-out timers based on measured round-trip delays. Because the sliding window algorithms acknowledge each packet, they have an opportunity to measure round-trip packet delay and use that delay to adjust the retransmission time out timer. The timer is only sampled when blocked packets are retransmitted.

Retransmit Timeout (MS)—Average, Maximum

The average and maximum timeout before a retransmission of the packet.

Round Trip Time (MS)—Average, Maximum

The average and maximum round trip delay time for a packet to get to a destination.

8.10 Ack Delay Report

SESSION SOURCES: ACK DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	NUMBER OF ACKS CREATED	NUMBER OF ACKS DROPPED	ACK DELAY (MS) AVERAGE	ACK DELAY (MS) MAXIMUM	ACK DELAY (MS) STD DEV
Session Server / src Session Source:					
Session Workstation	885	0	200.899	201.327	0.025

The Ack Delay Report complements the Packet Delay Report for Session Sources, to give statistics on the acks created and dropped as well as the delay for the ack. The ack delay is a round trip delay from when the packet started to when the ack returned. The time starts with the last packet that results in sending the ack.

Number of Acks Created and Dropped

The number of acks created and dropped in response to the source message.

Ack Delay (MS)—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of the ack delay in response to the source message.

8.11 Packet Size Report

SESSION SOURCES: PKT SIZE

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	MAX PKT	MAX WINDOW	PACKET SIZE MIN	AVG	MAX
Session Server / src Session Source:					
Session Workstation	1460	3	1000	1000	1000

The Packet Size Report provides two functions:

- The first function is to report on the maximum packet and window size resulting from the socket constraints at either end for session sources. This is important because the sockets' packet or window size constraint can reduce the maximum protocol window or packet size. The packet size is set by the transport protocol.
- The second function is to report on the packet sizes that were actually created for session sources. This part is most useful for monitoring the sizes for messages that are smaller than the packet sizes or for monitoring the sizes from external traffic sources.

Packet Maximum

The maximum size of a packet that was transmitted to the destination.

Window Maximum

The maximum window size, in packets.

Packet Size—Minimum, Average, Maximum

The average, minimum and maximum size of packets transmitted to the destination.

8.12 Burst Size Report

SESSION SOURCES: BURST SIZE

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	BURST UNITS	TRAF TYPE	BURST SIZE MIN	MAX	AVG	STD DEV
Workstation 1 / src ftp session:						
Workstation 2	kBits	FRLB	1	56	31	0
Workstation 3	kBits	FRLB	4	56	28	0

The Burst Size Report provides statistics for session sources on the burst size measured by the traffic policy on the protocol, when the traffic policy is present. It is the burst measurement that is used to determine the DE status of the packet, or for the ATM traffic policy, it is the burst measurement used to determine the conformance of the packets. Thus a determination is made whether the packets should be immediately dropped because they exceed their traffic contract. There are three burst measurement options implemented for modeling the traffic policing. The three measurements are Sliding Window, Jumping Window, and Leaky Bucket. The burst algorithms are set in the transport protocol under traffic policing.

The Sliding Window algorithm measures the burst as the total size of the packets that arrived during a fixed interval before the current time. This interval slides according to the time, and the measurement always includes all the packets that have arrived for the previous full interval.

The Jumping Window algorithm measures the burst based on accumulating a burst size with each packet, but then clearing that accumulation at fixed intervals.

The Leaky Bucket algorithm (or the Generalized Cell Rate Algorithm, GCRA for ATM) is similar to the jumping window in terms of allowing excess traffic, except that it periodically subtracts off an amount equal to the rate times the burst interval. This maintains some memory of the previous burst interval, especially if that burst exceeded the value of the burst rate times the burst interval.

Burst Units The type of units set for the burst size under the traffic policing option of the transport protocol.

Traffic Type One of the following:

FRLB = Frame Relay Leaky Bucket
FRSW = Frame Relay Sliding Window
FRJW = Frame Relay Jumping Window
ATM
None

Burst Size—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation of the burst size observed.

8.13 Assembly Interval Report

SESSION SOURCES: ASSY INTVL

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	PKTS ASS'D	INTER-ASSY INTERVAL MIN	MAX	AVG	STD DEV	UNIT
Workstation 1 / src ftp session:						
Workstation 2	119	586	6911623	867676	1466433	uS
Workstation 3	145	21	5057755	187540	584611	uS

The Assembly Interval Report provides statistics. The Inter-Assembly Interval is the delay between packets of the same message, when they are assembled at the destination. The assembly interval can result from delay variation through the network (due to varying buffer delays) or it can result from congestion on the destination node. In cell-based networks, this measure is the cell-delay variation, a quality of service measure for these services.

Packets Assembled

The number of packets assembled at the destination.

Inter-Assembly Interval—Minimum, Maximum, Average

The minimum, maximum, average and Standard Deviation Inter-Assembly Interval.

Inter-Assembly Units

The time unit used to measure delay between the packets.

8.14 Discard Eligible Packets Report

SESSION SOURCES: DE PACKETS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SESSION SRC: DESTINATION LIST	DE 0 PACKETS CREATED	DE 0 PACKETS DROPPED	DE 1 PACKETS CREATED	DE 1 PACKETS DROPPED	CONGEST PACKETS
Session Server / src Session Source:					
Session Workstation	885	0	0	0	0

The Discard Eligible Packets Report provides information for session sources about the flags set on the packet: the DE (discard eligibility) flag from the protocol's traffic policing algorithm, and the congestion flags set by the buffers when they are set up for FECN or BECN.

Discard Eligible 0 Packets Created and Dropped

The number of created and dropped packets flagged with a discard eligibility of 0.

Discard Eligible 1 Packets Created and Dropped

The number of created and dropped packets flagged with a discard eligibility of 1.

Congestion Packets

The number of packets that had a congestion flag set by buffers. This occurs when the buffers are set up for FECN or BECN.

9. Transport and Answer Commands

Transport and Answer Commands, when executed as part of an Application Source's command sequence, create packet-switched workload. Answer Commands automatically select the originator of the most recently received message as the destination. See Chapter 5 in the *COMNET III Reference Guide* for more information. The eleven reports associated with Transport and Answer Commands are presented in the following paragraphs.

For each node in the model sending an answer command, the report lists **Echo** as the destination list destination. **Echo** indicates that the delay messages assembled and the delay statistics are for all transport commands that triggered an answer command. Thus, the **Echo** line statistics display an aggregate number of messages assembled along with average delays. In order to see a breakdown of the messages assembled and the message delay statistics for each answer command message being sent to each destination, it is necessary to deselect the **Echo** checkbox on the **Destinations** tab window for the answer command parameter set.

The first column in each of the Transport and Answer commands reports is titled:

ORIGIN/COMMAND NAME: DESTINATION LIST

This column lists the name of the node and the name of the command for each node that has Transport Commands or Answer Commands defined for it and also lists the destination of the Transport or Answer Command. All destinations in the destination list of the command are listed. For Random Neighbor destinations, all nodes connected via one link to the originating node are listed.

9.1 Message Delay Report

TRANSPORT + ANSWER COMMANDS: MESSAGE DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	MESSAGES ASSEMBLED	MESSAGE DELAY AVERAGE	STD DEV	MAXIMUM
Application Workstation / cmd DataDB Request:				
Application Server	16	68832.718 MS	42828.548 MS	129.74992 S
Application Server / cmd Request Response:				
ECHO	0	0.000 MS	0.000 MS	0.000 MS
Application Workstat	5	27686.602 MS	32695.022 MS	68450.191 MS

The Message Delay Report presents message delay statistics, for transport and answer commands on the message, after the message has been reassembled by the destination, and the sender is notified that the message has been assembled. This is the message delay that is perceived by the sender. For each originating node in the model the Message Delay Report lists delays for each command to each destination.

Message delay is the time between creating the first packet of the message on the originating node and the time that the originating node is notified that the message has been assembled by the destination.

Messages Assembled The number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, standard deviation and maximum observed message delay.

9.2 Message Delivered Report

TRANSPORT + ANSWER COMMANDS: MESSAGE DELIVERED

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	MESSAGES ASSEMBLED	MESSAGE DELAY AVERAGE	STD DEV	MAXIMUM
Application Workstation / cmd DataDB Request:				
Application Server	16	67981.375 MS	42554.772 MS	128.60042 S
Application Server / cmd Request Response:				
ECHO	0	0.000 MS	0.000 MS	0.000 MS
Application Workstat	16	880.260 MS	518.383 MS	1817.572 MS

The Message Delivery Report presents statistics for transport and answer commands on the delay before the message is reassembled by the destination. The message may still be worked on by the source after the message is delivered because of retransmissions, waiting for acks, or for modeling the close-sequence of the connection, but these values are not factored into the message delivered delay.

Message delay is the time between creating the first packet of the message at the originating node and the time of receiving the last packet on the destination node.

Messages Assembled

The number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, standard deviation and maximum observed message delay.

9.3 Packet Delay Report

TRANSPORT + ANSWER COMMANDS: PACKET DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	NUMBER OF PACKETS CREATED	DELIVERED	RESENT	DROPPED	PACKET DELAY (MS) AVERAGE	MAXIMUM
Application Workstation / cmd DataDB Request:						
Application Server	68	16	0	0	23726.978	66561.951
Application Server / cmd Request Response:						
ECHO	0	0	0	0	0.000	0.000
Application Workstat	160	128	0	0	3.302	9.328

The Packet Delay Report presents a summary of the number of packets that were created, delivered, resent, or dropped for each transport and answer command. An average and maximum packet delay is also displayed.

Packet delay is the time between creating a packet on the originating node and the time of receiving the packet at the destination node.

Packets Created The number of packets that have been created at the origin to send to the listed destination. Not all packets need to be delivered to appear on the report.

Packets Delivered The number of packets that have been received. This may differ from the number of packets created by the number of packets that are in transit at the instant the report is written.

Packets Retransmitted

The number of retransmitted packets. Packets may be retransmitted from the origin because they are blocked at some point in route to the destination. Blocking can occur because input or output buffers are full, or because a node or link on the route fails.

Packets Dropped The number of dropped packets. When a node or link fails, you can specify whether the transmission should be reattempted. If no retransmission is specified then the packet will be dropped.

Packet Delay—Average, Maximum

The average and maximum of the observed packet delays.

9.4 Window Stats Report

```
                    TRANSPORT + ANSWER COMMANDS: WINDOW STATS

                     REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:    WINDOW      NO. OF  CONG      PKT CREATION INTERVAL (MS)
  DESTINATION LIST       AVG    MAX  RESETS  AVOID      AVG       MAX      STD DEV
_________________________ _______ ___ _______ _______ ________ ________ ________

Application Workstation / cmd DataDB Request:
  Application Server      2.00    2       0       0     32228    64275    19793

Application Server / cmd Request Response:
  ECHO                    0.00    0       0       0         0        0        0
  Application Workstat    2.88    3       0       0       109     1001      227
```

The Window Stats Report presents the statistics on the TCP/IP congestion window in packets that varies as the connection starts up and when congestion is detected. It also presents the statistics for the interval between packets being created.

Window Size—Average, Maximum

The average and maximum size of the TCP/IP congestion window.

Number Of Congestion Resets

The number of times that the TCP/IP "congestion" window (dynamic window size) is reset to 1 packet. Thus, congestion window size is reset.

Congestion Avoidances

After a reset, the window size starts to grow more slowly (linearly) than the initial "slow start" window growth. This algorithm switch is counted as "Congestion Avoidance" and gives a measure of how often the network is recovering from congestion. If Congestion Avoidance is nearly the same as the Number of Congestion Resets, then congestion may be spurious or short-lived. If there are lots of Congestion Resets and few Congestion Avoidances, then either the network is very congested or the message sizes are not very large.

Packet Creation Interval (MS)—Average, Maximum, Standard Deviation

The average, maximum, and standard deviation of the time interval between the creation of packets.

9.5 Retransmissions Report

TRANSPORT + ANSWER COMMANDS: RETRANSMISSIONS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	NO. OF EVENTS	NO. PACKETS AVG	NO. PACKETS MAX	PKT RETRIES AVG	PKT RETRIES MAX	BLCKD AFTER T-O PKTS	BLCKD AFTER T-O ACKS
Application Workstation / cmd DataDB Request:							
Application Server	0	0.00	0	0.00	0	0	0
Application Server / cmd Request Response:							
ECHO	0	0.00	0	0.00	0	0	0
Application Workstat	0	0.00	0	0.00	0	0	0

The Retransmissions Report presents retransmission statistics for transport and answer commands when blocked packets are retransmitted. The retransmission timer accounts for the age of the packet before it is blocked. It schedules a packet retransmission only after the packet is actually blocked somewhere in the network. The retransmission time is adjusted so that the packet is retransmitted the specified delay after the packet was originally transmitted. It reports on quantities such as the number of times a packet has to be retransmitted, and the number of packets that are retransmitted when a blocked packet occurs.

Number Of Events The number of events.

Number Packets—Average, Maximum

The average and maximum number of packets that were retransmitted.

Packet Retries—Average, Maximum

The average and maximum number of times a packet was resent.

Packets Blocked After T-O

The number of packets that were blocked after timing out.

Acknowledgments Blocked After T-O

The number of acknowledgments that were blocked after timing out.

9.6 Timeout Report

```
                    TRANSPORT + ANSWER COMMANDS: TIMEOUT

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:  RETRANSMIT TIMEOUT (MS)       ROUND TRIP TIME (MS)
  DESTINATION LIST          AVG        MAX              AVG         MAX
_____________________  __________  __________      __________  __________

Application Workstation / cmd DataDB Request:
  Application Server         0.000       0.000        33912.299   66848.408

Application Server / cmd Request Response:
  ECHO                       0.000       0.000            0.000       0.000
  Application Workstat       0.000       0.000        10950.799   66765.428
```

The Timeout Report presents time-out timer statistics for transport and answer commands. This report is most relevant for enhanced sliding window or TCP/IP window protocols that adapt their time-out timers based on measured round-trip delays. Because the sliding window algorithms acknowledge each packet, they have an opportunity to measure round-trip packet delay and use that delay to adjust the retransmission time out timer. The timer is only sampled when blocked packets are retransmitted.

Retransmit Timeout (MS)—Average, Maximum

The average and maximum timeout before the retransmission of the packet.

Round Trip Time (MS)—Average, Maximum

The average and maximum round trip delay time for a packet to get to a destination.

9.7 Ack Delay Report

TRANSPORT + ANSWER COMMANDS: ACK DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	NUMBER OF ACKS CREATED	DROPPED	ACK DELAY (MS) AVERAGE	MAXIMUM	STD DEV
Application Workstation / cmd DataDB Request:					
Application Server	57	0	33912.299	66848.408	32935.165
Application Server / cmd Request Response:					
ECHO	0	0	0.000	0.000	0.000
Application Workstat	112	0	10950.799	66765.428	24557.067

The Ack Delay Report complements the packet delay report for transport and answer commands, to give statistics on the acks created and dropped as well as the delay for the ack. The ack delay is a round trip delay from when the packet started to when the ack returned. The time starts with the last packet that results in sending the ack.

Number of Acks Created and Dropped

The number of acks created and dropped in response to the source message.

Delay (MS)—Average, Maximum, Standard Deviation Ack Delay

The average, maximum and standard deviation of the delay for the ack in response to the source message.

9.8 Packet Size Report

```
                    TRANSPORT + ANSWER COMMANDS: PKT SIZE

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:     MAX       MAX          PACKET SIZE
  DESTINATION LIST         PKT     WINDOW     MIN       AVG       MAX
_____________________   _______   _______   _______   _______   _______

Application Workstation / cmd DataDB Request:
  Application Server       1460         3      100       100       100

Application Server / cmd Request Response:
  ECHO                     1460         3        0         0         0
  Application Workstat     1460         3       20      1280      1460
```

The Packet Size Report provides two functions:

- The first function is to report on the maximum packet and window size resulting from the socket constraints at either end for transport and answer commands. This is important because the sockets' packet or window size constraint can reduce the maximum protocol window or packet size.
- The second function is to report on the packet sizes that were actually created for transport and answer commands. This function is most useful for monitoring the sizes of messages that are smaller than the packet sizes or for monitoring the sizes of messages from external traffic sources.

The packet size is set by the transport protocol.

Packet Maximum The maximum size of a packet transmitted to the destination.

Window Maximum The maximum window size, in packets, defined by the transport protocol.

Packet Size—Minimum, Average, Max

The average, minimum and maximum size of a packet that was transmitted to the destination.

9.9 Burst Size Report

SETUP COMMANDS: BURST SIZE

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	BURST UNITS	TRAF TYPE	BURST SIZE MIN	MAX	AVG	STD DEV
Workstation 1 / src ftp:						
Workstation 1 / src ftp 2:						
cmd send data stream:						
Workstation 2	kBits	FRLB	2	56	31	0
Workstation 3	kBits	FRLB	3	56	29	0

The Burst Size Report provides statistics for transport and answer commands on the burst size measured by the traffic policy on the protocol, when the traffic policy is present. It is the burst measurement that is used to determine the DE status of the packet, or for the ATM traffic policy, it is the burst measurement used to determine the conformance of the packets. Thus, a determination is made whether the packets should be immediately dropped because they exceed their traffic contract. There are three burst measurement options implemented for modeling the traffic policing. The three measurements are Sliding Window, Jumping Window, and Leaky Bucket. The burst algorithms are set in the transport protocol under traffic policing.

The Sliding Window algorithm measures the burst as the total size of the packets that arrived during a fixed interval before the current time. This interval slides according to the time, and the measurement always includes all the packets that have arrived for the previous full interval.

The Jumping Window algorithm measures the burst based on accumulating a burst size with each packet, but then clearing that accumulation at fixed intervals.

The Leaky Bucket algorithm (or the Generalized Cell Rate Algorithm, GCRA for ATM) is similar to the jumping window algorithm in terms of allowing excess traffic, except that it periodically subtracts off an amount equal to the rate times the burst interval. This maintains some memory of the previous burst interval, especially if that burst exceeded the value of the burst rate times the burst interval.

Burst Units The type of units set for the burst size under the traffic policing option of the transport protocol.

Traffic Type One of the following:

FRLB = Frame Relay Leaky Bucket
FRSW = Frame Relay Sliding Window
FRJW = Frame Relay Jumping Window
ATM
None

Burst Size—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation burst size defined under the traffic policing option of the transport protocol.

9.10 Assembly Interval Report

```
                          SETUP COMMANDS: ASSY INTVL

                     REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN / COMMAND NAME:   PKTS            INTER-ASSY INTERVAL
  DESTINATION LIST       ASS'D      MIN         MAX         AVG      STD DEV
UNIT
_________________________  _______  _________  _________  _________  _________

Workstation 1 / src ftp:
Workstation 1 / src ftp 2:

  cmd send data stream:
  Workstation 2             165        395   31895643    2017873    4672067 uS
  Workstation 3             109         39   16970902     385931    1967765 uS
```

The Assembly Interval Report provides statistics for transport and answer commands on the delay time between when different packets from the same message are assembled at the destination. The assembly interval can result from delay variation through the network (due to varying buffer delays) or it can result from congestion on the destination node. In cell-based networks, this measure is the cell-delay variation, a quality of service measure for these services.

Packets Assembled The number of packets assembled at the destination.

Inter-Assembly Interval—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation for the delay between the packets.

Inter-Assembly Units

The time unit used to measure delay between the packets.

9.11 Discard Eligible Packets Report

TRANSPORT + ANSWER COMMANDS: DE PACKETS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	DE 0 PACKETS CREATED	DE 0 PACKETS DROPPED	DE 1 PACKETS CREATED	DE 1 PACKETS DROPPED	CONGEST PACKETS
Application Workstation / cmd DataDB Request:					
Application Server	68	0	0	0	0
Application Server / cmd Request Response:					
ECHO	0	0	0	0	0
Application Workstat	160	0	0	0	0

The Discard Eligible Packets Report provides information for transport and answer commands about the flags set on the packet: the DE (discard eligibility) flag from the protocol's traffic policing algorithm, and the congestion flags set by the buffers when they are set up for FECN or BECN.

Discard Eligible 0 Packets Created and Dropped

The number of created and dropped packets that were flagged with a discard eligibility of 0.

Discard Eligible 1 Packets Created and Dropped

The number of created and dropped packets flagged with a discard eligibility of 1.

Congestion Packets

The number of packets that had a congestion flag set by buffers.

10. Setup Commands

Setup Commands, when executed as part of an Application Source's command sequence, establish a pattern of packet switched traffic, involving optional handshaking messages, and including any number of content messages, between the same source and destination. COMNET III produces the fourteen different Setup Commands reports described in the following paragraphs.

The first column in most of the Setup Commands reports is titled:

ORIGIN/COMMAND NAME: DESTINATION LIST

This column lists the name of the node and the name of the command for each node that has Setup Commands defined for it and also lists the destination of the Setup command. All destinations in the destination list of the command are listed. For Random Neighbor destinations, all nodes connected via one link to the originating node are listed.

10.1 Message Delay Report

SETUP COMMANDS: MESSAGE DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	MESSAGES ASSEMBLED	MESSAGE DELAY AVERAGE	STD DEV	MAXIMUM
Application Workstation / cmd Local Setup:				
Application Server	73	54927.612 MS	36443.594 MS	133.69003 S

The Message Delay Report presents message delay statistics for setup commands on the message, after the message has been reassembled by the destination and the sender is notified that the message has been assembled. This is the message delay that is perceived by the sender. For each originating node in the model it lists delays to each destination.

Message delay is the time between creating the first packet of the message on the originating node and the time that the originating node is notified that the message has been assembled by the destination.

Messages Assembled

The number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, standard deviation and maximum observed message delay.

10.2 Message Delivered Report

```
                    SETUP COMMANDS: MESSAGE DELIVERED

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:  MESSAGES                MESSAGE DELAY
  DESTINATION LIST      ASSEMBLED      AVERAGE       STD DEV       MAXIMUM
_____________________   _________   ____________  ____________  ____________

Application Workstation / cmd Local Setup:
  Application Server            73   49790.712 MS  38478.748 MS  132.54053 S
```

The Message Delivered Report presents statistics for setup commands on the message delay before the message is reassembled by the destination. The message may still be worked on by the source after the message is delivered because of retransmissions, waiting for acks, or for modeling the close-sequence of the connection, but these values are not factored into the message delivered delay.

Message delay is the time between creating the first packet of the message on the originating node and the time of receiving the last packet of the message on the destination node.

Messages Assembled

The number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, standard deviation and maximum observed message delay.

10.3 Packet Delay Report

SETUP COMMANDS: PACKET DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	NUMBER OF PACKETS CREATED	DELIVERED	RESENT	DROPPED	PACKET DELAY (MS) AVERAGE	MAXIMUM
Application Workstation / cmd Local Setup:						
Application Server	278	73	0	0	18200.026	66567.951

The Packet Delay Report presents a summary for setup commands of the number of packets that were created, delivered, resent, or dropped. An average and maximum packet delay is also displayed.

Packet Delay the time between creating a packet on the originating node and the time of receiving the packet at the destination node.

Packets Created

The number of packets created at the origin to send to the listed destination. Not all packets need to be delivered before they appear on the report.

Packets Delivered

The number of packets that have been received at the destination. This may differ from the number of packets created by the number of packets that are in transit at the instant the report is written.

Packets Retransmitted

The number of packets retransmitted. Packets may be retransmitted from the origin because they are blocked at some point in route to the destination. Blocking can occur because input or output buffers are full, or because a node or link on the route fails.

Packets Dropped

The number of packets dropped. When a node or link fails, you can specify whether the transmission should be reattempted. If so, the retransmissions will be counted in the **Packets Retransmitted** field. If no retransmission is specified, then the packet will be dropped.

Packet Delay (MS)—Average, Maximum

The average and maximum observed packet delay.

10.4 Setup Delay Report

```
                    SETUP COMMANDS: SETUP DELAY

                REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SETUP CMD:      SESSIONS          SETUP DELAY (MILLISECONDS)
  DESTINATION LIST         SETUP       AVERAGE       STD DEV       MAXIMUM
_____________________  _________  ____________  ____________  ____________

Application Workstation / cmd Local Setup:
  Application Server           25   32227.077 MS  19793.936 MS  64274.337 MS
```

The Setup Delay Report presents summary statistics for setup commands on the time required to set up a session between an originating node and a destination node. The number of sessions that were set up along with delay statistics are reported.

Setup Delay is the time difference between creating the session setup packet and receiving back the session connect packet.

Sessions Setup

The number of sessions set up. When a session is started by a setup command a session setup packet is sent to the destination and a session connect packet is returned. The session is then counted as set up.

Setup Delay (milliseconds)—Average, Standard Deviation, Maximum

The average, standard deviation and maximum observed setup delay.

10.5 Session Length Report

SETUP COMMANDS: SESSION LENGTH

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN	SETUP CMD NAME	SESSIONS ENDED	SESSION LENGTH AVERAGE	STD DEV	MAXIMUM
Application Workst	Local Setup	13	109.8704 S	35.1256 S	193.7018 S

The Session Length Report presents summary statistics for setup commands on the duration of a session between an originating node and a destination node. It reports the number of sessions completed along with the session length.

The session length is the time between the session setup packet being created and the end of the session.

Origin The node name which originates the session by setup command.

Setup Command Name

The name of the setup command.

Sessions Ended

The number of complete and cleared sessions. As the simulation executes, sessions are setup and messages transmitted across them. A session is counted as complete when the last message of a session has been received at the destination, and all packets for response messages have been received at the session origination node, and all ACKS for all packets have been received, and all pending message notices created by the receipt of the session messages have been cleared. When these conditions are met the session is complete and cleared.

Session Length—Average, Standard Deviation, Maximum

The average, standard deviation and maximum observed session length.

10.6 Setup Counts Report

```
                    SETUP COMMANDS: SETUP COUNTS

                REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SETUP COMMAND NAME:                  NUMBER OF SESSIONS
  DESTINATION LIST                  TRIED    SETUP   RETRY  BLOCK  DISCON RERTD
_____________________________________ _______ _______ _____ _____ _____ _____

Application Workstation / cmd Local Setup:
  Application Server                    27       25      0      0      0      0
```

The Setup Counts Report presents summary statistics for setup commands on the number of attempts to set up a session, the number of attempts that succeeded, the number of retries to set up a session based upon initial failures, the number of sessions attempted that were blocked, the number of sessions attempted that were disconnected, and the number of sessions that were rerouted due to being blocked.

Sessions Tried

The number of attempts to set up a session by the respective setup command. As the simulation executes, applications which call the session setup command try to execute it. The attempt to set up a particular session will succeed or fail depending upon network conditions. For instance, route availability inside hop and session limits, buffer availability, etc.

Sessions Setup

The number of session setup attempts which succeeded.

Sessions Retried

The number of retried sessions. If a setup attempt fails, the session may be retried later depending on the settings you have entered.

Sessions Blocked

The number of blocked sessions. The session setup attempt may block due to insufficient buffer space on routing nodes, or no route being available inside the hop limit, or no route being available because at least 1 link on all routes is at its session limit, or nodes/links have failed and no route is available.

Sessions Disconnected

The number of disconnected sessions. A session in progress may be disconnected because a node or link through which it is routed fails. The session may optionally be rerouted if this happens.

Sessions Rerouted

The number of sessions rerouted following disconnection.

10.7 Window Stats Report

SETUP COMMANDS: WINDOW STATS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	WINDOW AVG	WINDOW MAX	NO. OF RESETS	CONG AVOID	PKT CREATION INTERVAL (MS) AVG	PKT CREATION INTERVAL (MS) MAX	PKT CREATION INTERVAL (MS) STD DEV
Application Workstation / cmd Local Setup:							
Application Server	2.00	2	0	0	45466	66564	29512

The Window Stats Report presents the statistics for setup commands on the TCP/IP congestion window that varies as the connection starts up and when congestion is detected. It also presents the statistics for the interval between packets being created.

Window Size—Average, Maximum

The average and maximum size of the TCP/IP congestion window, in packets, that varies as the connection starts up and when congestion is detected.

Number Of Congestion Resets

The number of times that the TCP/IP "congestion" window (dynamic window size) is reset to 1 packet. Thus, congestion window size is reset.

Number Of Congestion Avoidances

After a reset, the window size starts to grow more slowly (linearly) than the initial "slow start" window growth. This algorithm switch is counted as "Congestion Avoidance" and gives a measure of how often the network is recovering from congestion. If Congestion Avoidance is nearly the same as the Number of Congestion Resets, then congestion may be spurious or short-lived. If there are lots of Congestion Resets and few Congestion Avoidances, then either the network is very congested or the message sizes are not very large.

Packet Creation Interval (MS)—Average, Maximum, Standard Deviation

The average, standard deviation and maximum packet creation interval.

10.8 Retransmissions Report

SETUP COMMANDS: RETRANSMISSIONS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	NO. OF EVENTS	NO. PACKETS AVG	NO. PACKETS MAX	PKT RETRIES AVG	PKT RETRIES MAX	BLCKD AFTER T-O PKTS	BLCKD AFTER T-O ACKS
Application Workstation / cmd Local Setup:							
Application Server	0	0.00	0	0.00	0	0	0

The Retransmissons Report presents retransmission statistics for setup commands when blocked packets are retransmitted. The retransmission timer accounts for the age of the packet before it is blocked. It schedules a packet retransmission only after the packet is actually blocked somewhere in the network. The retransmission time is adjusted so that the packet is retransmitted the specified delay after the packet was originally transmitted. It reports on quantities such as the number of times a packet has to be retransmitted, and the number of packets that are retransmitted when a blocked packet occurs.

Number Of Events

The number of retransmissions.

Number Packets—Average, Maximum

The average and maximum number of packets retransmitted.

Packet Retries—Average, Maximum

The average and maximum number of times a packet was resent.

Packets Blocked After T-O

The number of packets blocked after timing out.

Acknowledgments Blocked After T-O

The number of acknowledgments blocked after timing out.

10.9 Timeout Report

SETUP COMMANDS: TIMEOUT

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	RETRANSMIT TIMEOUT (MS) AVG	RETRANSMIT TIMEOUT (MS) MAX	ROUND TRIP TIME (MS) AVG	ROUND TRIP TIME (MS) MAX
Application Workstation / cmd Local Setup:				
Application Server	0.000	0.000	11037.487	66853.652

The Timeout Report presents time-out timer statistics for setup commands. This report is most relevant for enhanced sliding window or TCP/IP window protocols that adapt their time-out timers based on measured round-trip delays. Because the sliding window algorithms acknowledge each packet, they have an opportunity to measure round-trip packet delay and use that delay to adjust the retransmission time out timer. The timer is only sampled when blocked packets are retransmitted.

Retransmit Timeout (MS)—Average, Maximum

The average and maximum timeout before the retransmission of the packet.

Round Trip Time (MS)—Average, Maximum

The average and maximum round trip delay time for a packet to get to a destination.

10.10 Ack Delay Report

SETUP COMMANDS: ACK DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	NUMBER OF ACKS CREATED	NUMBER OF ACKS DROPPED	ACK DELAY (MS) AVERAGE	ACK DELAY (MS) MAXIMUM	ACK DELAY (MS) STD DEV
Application Workstation / cmd Local Setup:					
Application Server	226	0	11037.487	66853.652	23481.046

The Ack Delay Report complements the Packet Delay Report for setup commands. This report gives statistics on the acks created and dropped as well as the delay for the ack.

Ack Delay is the round trip delay from when the packet started to when the ack returned. The time starts with the last packet that results in sending the acknowledgment.

Number of Acks Created and Dropped

The number of acks created and dropped.

Ack Delay (MS)—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of the delay for the ack.

10.11 Packet Size Report

SETUP COMMANDS: PKT SIZE

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	MAX PKT	MAX WINDOW	PACKET SIZE MIN	AVG	MAX
Application Workstation / cmd Local Setup:					
Application Server	1460	3	1000	1000	1000

The Packet Size Report provides two functions:

- The first function is to report on the maximum packet and window size resulting from the socket constraints at either end for setup commands. This is important because the sockets' packet or window size constraint can reduce the maximum protocol window or packet size. The packet size is set by the transport protocol.
- The second function is to report on the packet sizes that were actually created for setup commands. This function is most useful for monitoring the sizes for messages that are smaller than the packet sizes or for monitoring the sizes from external traffic sources.

Packet Maximum

The maximum size of a packet transmitted to the destination.

Window Maximum

The maximum window size, in packets, defined by the transport protocol.

Packet Size—Minimum, Average, Max

The minimum, average and maximum observed size of a packet transmitted to the destination.

10.12 Burst Size Reports

```
                    SETUP COMMANDS: BURST SIZE

                REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:   BURST   TRAF          BURST SIZE
  DESTINATION LIST       UNITS   TYPE    MIN       MAX       AVG     STD DEV
_____________________   _______  ____  ________  ________  ________  ________

Workstation 1 / src ftp:
Workstation 1 / src ftp 2:

  cmd send data stream:
  Workstation 2         kBits    FRLB        2        56        31         0
  Workstation 3         kBits    FRLB        3        56        29         0
```

The Burst Size Report provides statistics for setup commands on the burst size measured by the traffic policy on the protocol, when the traffic policy is present. It is the burst measurement that is used to determine the DE status of the packet, or for the ATM traffic policy, it is the burst measurement used to determine the conformance of the packets. Thus, a determination is made whether the packets should be immediately dropped because they exceed their traffic contract. There are three burst measurement options implemented for modeling the traffic policing. The three burst measurement options implemented for modeling the traffic policing are Sliding Window, Jumping Window, and Leaky Bucket. The burst algorithms are set in the transport protocol under traffic policing. Burst size is defined under the traffic policing option of the transport protocol.

The Sliding Window algorithm measures the burst as the total size of the packets that arrived during a fixed interval before the current time. This interval slides according to the time, and the measurement always includes all the packets that have arrived for the previous full interval.

The Jumping Window algorithm measures the burst based on accumulating a burst size with each packet, but then clearing that accumulation at fixed intervals.

The Leaky Bucket algorithm (or the Generalized Cell Rate Algorithm, GCRA for ATM) is similar to the Jumping Window algorithm in terms of allowing excess traffic, except that it periodically subtracts off an amount equal to the rate times the burst interval. This maintains some memory of the previous burst interval, especially if that burst exceeded the value of the burst rate times the burst interval.

`Burst Units` The type of units set for the burst size.

`Traffic Type` One of the following:

`FRLB` = Frame Relay Leaky Bucket
`FRSW` = Frame Relay Sliding Window
`FRJW` = Frame Relay Jumping Window
`ATM`
`None`

Burst Size—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation of the observed burst size.

10.13 Assembly Interval Report

```
                    SETUP COMMANDS: ASSY INTVL

                REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:   PKTS            INTER-ASSY INTERVAL
  DESTINATION LIST       ASS'D      MIN        MAX        AVG      STD DEV  UNIT
_____________________  ________  _________  _________  _________  _________  __

Workstation 1 / src ftp:
Workstation 1 / src ftp 2:

  cmd send data stream:
  Workstation 2             165        395   31895643    2017873    4672067   uS
  Workstation 3             109         39   16970902     385931    1967765   uS
```

The Assembly Interval Report provides statistics for setup commands on the delay time between when different packets from the same message are assembled at the destination. The assembly interval can result from delay variation through the network (due to varying buffer delays) or it can result from congestion on the destination node. In cell-based networks, this measure is the cell-delay variation, a quality of service measure for these services.

The Inter-Assembly Interval is the delay between the packets when they are assembled at the destination.

Packets Assembled

The number of packets assembled at the destination.

Inter-Assembly Interval—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation for the inter-assembly interval.

Inter-Assembly Units

The time unit used to measure delay between the packets.

10.14 Discard Eligible Packets Report

SETUP COMMANDS: DE PACKETS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	DE 0 PACKETS CREATED	DE 0 PACKETS DROPPED	DE 1 PACKETS CREATED	DE 1 PACKETS DROPPED	CONGEST PACKETS
Application Workstation / cmd Local Setup:					
Application Server	278	0	0	0	0

The Discard Eligible Packets Report provides information for setup commands about the flags set on the packet: the DE (discard eligibility) flag from the protocol's traffic policing algorithm, and the congestion flags set by the buffers when they are set up for FECN or BECN.

Discard Eligible 0 Packets Created and Dropped

The number of packets created and dropped that were flagged with a discard eligibility of 0.

Discard Eligible 1 Packets Created and Dropped

The number of packets created and dropped that were flagged with a discard eligibility of 1.

Congestion Packets

The number of packets that had a congestion flag set by buffers.

11. Global Transport and Answer Commands

Global Transport and Answer Commands, when executed as part of an Application Source's command sequence, create packet-switched workload. Answer Commands automatically select the originator of the most recently received message as the destination. See Chapter 5 in the *COMNET III Reference Guide* for more information.

COMNET III produces the eleven different Global Transport and Answer Commands reports described in the following paragraphs.

The first column in most of the Global Transport and Answer Commands reports is titled:

ORIGIN/COMMAND NAME: DESTINATION LIST

This column lists the name of the node and the name of the command for each node that has Transport and Answer commands defined for it and also lists the destination of the Transport and Answer command. All destinations in the destination list of the command are listed. For Random Neighbor destinations, all nodes connected via one link to the originating node are listed.

11.1 Message Delay Report

```
                GLOBAL TRANSPORT + ANSWER COMMANDS: MESSAGE DELAY

                    REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:  MESSAGES                MESSAGE DELAY
  DESTINATION LIST      ASSEMBLED      AVERAGE       STD DEV       MAXIMUM
_____________________   _________   ____________  ____________  ____________

Application Workstation / src Application Request:
  cmd Global DB Request:
  Application Server           16   66728.513 MS     67.894 MS  66779.767 MS

Application Server / src Application Source:
  cmd Global Request Response:
  ECHO                          3   140.31324 S   53369.856 MS  193.69588 S
```

The Message Delay Report presents message delay statistics for global transport and answer commands, after the message has been reassembled by the destination and the sender has been notified. This is the message delay that is perceived by the sender. For each originating node in the model the report lists delays for each command to each destination.

Message Delay is the time between creating the first packet of the message on the originating node and the time that the originating node is notified that the message has been assembled by the destination.

Messages Assembled

For each destination, the number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of the observed message delay.

11.2 Message Delivered Report

```
            GLOBAL TRANSPORT + ANSWER COMMANDS: MESSAGE DELIVERED

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:  MESSAGES                 MESSAGE DELAY
  DESTINATION LIST      ASSEMBLED      AVERAGE       STD DEV       MAXIMUM
_____________________   _________   ____________  ____________  ____________

Application Workstation / src Application Request:
  cmd Global DB Request:
  Application Server            16      10.547 MS      5.121 MS     17.869 MS

Application Server / src Application Source:
  cmd Global Request Response:
  ECHO                           4   137.38081 S   47779.985 MS  197.35349 S
```

The Message Delivered Report presents statistics for global transport and answer commands on the delay before the message is reassembled by the destination. The message may still be worked on by the source after the message is delivered because of retransmissions, waiting for acks, or for modeling the close-sequence of the connection, but these values are not factored into the message delivered delay.

Message delay is the time between creating the first packet of the message on the originating node and the time of receiving the last packet on the destination node.

Messages Assembled

For each destination, the number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of the observed message delay.

11.3 Packet Delay Report

```
                GLOBAL TRANSPORT + ANSWER COMMANDS: PACKET DELAY

                    REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:            NUMBER OF PACKETS           PACKET DELAY (MS)
  DESTINATION LIST       CREATED DELIVERED RESENT DROPPED   AVERAGE    MAXIMUM
_________________________ _______ _________ _____ _____ _________ __________

Application Workstation / src Application Request:
  cmd Global DB Request:
  Application Server          48        16      0      0      2.980      11.655

Application Server / src Application Source:
  cmd Global Request Response:
  ECHO                        11         4      0      0      4.005      17.232
```

The Packet Delay Report presents a summary of the number of packets that were created, delivered, resent, or dropped for each global transport and answer command. An average and maximum packet delay is also displayed.

Packet Delay is the time between creating a packet on the originating node and the time of receiving the packet at the destination node.

Packets Created

The number of packets created at the origin to send to the listed destination. Not all packets need to be delivered before they appear on the report.

Packets Delivered

The number of packets received at the destination. This may differ from the number of packets created by the number of packets that are in transit at the instant the report is written.

Packets Retransmitted

The number of packets retransmitted. Packets may be retransmitted from the origin because they are blocked at some point in route to the destination. Blocking occurs when input or output buffers are full, or when a node or link on the route fails.

Packets Dropped

The number of packets dropped. When a node or link fails, you can specify whether the transmission should be reattempted. If so, the retransmissions will be counted in the **Packets Retransmitted** field. If no retransmission is specified then the packet will be dropped.

Packet Delay (MS)—Average, Maximum

The average and maximum observed packet delay.

11.4 Window Stats Report

```
                GLOBAL TRANSPORT + ANSWER COMMANDS: WINDOW STATS

                   REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:    WINDOW     NO. OF  CONG     PKT CREATION INTERVAL (MS)
  DESTINATION LIST       AVG    MAX  RESETS  AVOID     AVG      MAX     STD DEV
_____________________  _______  ___  ______  ______  _______  _______  _______

Application Workstation / src Application Request:
  cmd Global DB Request:
  Application Server      2.00    2       0       0        9       15        4

Application Server / src Application Source:
  cmd Global Request Response:
  ECHO                    2.00    2       0       0      293     1143      490
```

The Window Stats Report presents the statistics on the TCP/IP congestion window that varies as the connection starts up and when congestion is detected. It also presents the statistics for the interval between packets being created.

Window Size—Average, Maximum

The average and maximum size of the TCP/IP congestion window, in packets.

Number Of Congestion Resets

The number of times that the TCP/IP "congestion" window (dynamic window size) is reset to 1 packet. Thus, congestion window size is reset.

Congestion Avoidances

After a reset, the window size starts to grow more slowly (linearly) than the initial "slow start" window growth. This algorithm switch is counted as "Congestion Avoidance" and gives a measure of how often the network is recovering from congestion. If Congestion Avoidance is nearly the same as the Number of Congestion Resets, then congestion may be spurious or short-lived. If there are lots of Congestion Resets and few Congestion Avoidances, then either the network is very congested or the message sizes are not very large.

Packet Creation Interval (MS)—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of the observed time interval between the creation of packets.

11.5 Retransmissions Delays Report

GLOBAL TRANSPORT + ANSWER COMMANDS: RETRANSMISSIONS

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	NO. OF EVENTS	NO. PACKETS AVG	NO. PACKETS MAX	PKT RETRIES AVG	PKT RETRIES MAX	BLCKD AFTER T-O PKTS	BLCKD AFTER T-O ACKS
Workstation 1 / src ftp:							
cmd send data:							
Workstation 1	111	1.00	1	0.00	0	0	0
Workstation 2	166	1.00	1	11.00	11	0	0
Workstation 3	127	1.00	1	0.00	0	0	0
Workstation 1 / src ftp 2:							
cmd send data:							
Workstation 1	165	1.00	1	0.00	0	0	0
Workstation 2	77	1.00	1	4.00	4	0	0
Workstation 3	60	1.00	1	0.00	0	0	0

The Retransmissions Delay Report presents retransmission statistics for global transport and answer commands when blocked packets are retransmitted. The retransmission timer accounts for the age of the packet before it is blocked. It schedules a packet retransmission only after the packet is actually blocked somewhere in the network. The retransmission time is adjusted so that the packet is retransmitted the specified delay after the packet was originally transmitted. It reports on quantities such as the number of times a packet has to be retransmitted, and the number of packets that are retransmitted when a blocked packet occurs.

Number Of Events

The number of retransmissions.

Packets—Average, Maximum Number

The maximum and average number of packets retransmitted.

Packet Retries—Average, Maximum

The average and maximum number of times a packet was resent.

Packets Blocked After T-O

The number of packets blocked after timing out.

Acknowledgments Blocked After T-O

The number of acknowledgments blocked after timing out.

11.6 Timeout Report

```
                GLOBAL TRANSPORT + ANSWER COMMANDS: TIMEOUT

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:  RETRANSMIT TIMEOUT (MS)       ROUND TRIP TIME (MS)
  DESTINATION LIST          AVG        MAX              AVG         MAX
_____________________  __________  __________       __________  __________

Application Workstation / src Application Request:
  cmd Global DB Request:
  Application Server         0.000       0.000        66698.525   66757.454

Application Server / src Application Source:
  cmd Global Request Response:
  ECHO                       0.000       0.000        22524.207   66560.000
```

The Timeout Report presents time-out timer statistics for global transport and answer commands. This report is most relevant for enhanced sliding window or TCP/IP window protocols that adapt their time-out timers based on measured round-trip delays. Because the sliding window algorithms acknowledge each packet, they have an opportunity to measure round-trip packet delay and use that delay to adjust the retransmission time out timer. The timer is only sampled when blocked packets are retransmitted.

Retransmit Timeout (MS)—Average, Maximum

The average and maximum timeout before the retransmission of the packet.

Round Trip Time (MS)—Average, Maximum

The average and maximum round trip delay time for a packet to get to a destination.

11.7 Ack Delay Report

```
                GLOBAL TRANSPORT + ANSWER COMMANDS: ACK DELAY

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:    NUMBER OF ACKS              ACK DELAY (MS)
  DESTINATION LIST      CREATED    DROPPED    AVERAGE    MAXIMUM     STD DEV
_____________________  ________  ________  __________  __________  __________

Application Workstation / src Application Request:
  cmd Global DB Request:
  Application Server          48         0   66698.525   66757.454      72.746

Application Server / src Application Source:
  cmd Global Request Response:
  ECHO                        11         0   22524.207   66560.000   31138.511
```

The Ack Delay Report complements the Packet Delay Report for global transport and answer commands, to give statistics on the acks created and dropped as well as the delay for the ack. The ack delay is a round trip delay from when the packet started to when the ack returned. The time starts with the last packet that results in sending the ack.

Number of Acks Created

The number of acks created.

Number of Acks Dropped

The number of acks dropped.

Ack Delay (MS)—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of the observed ack delay.

11.8 Packet Size Report

```
                GLOBAL TRANSPORT + ANSWER COMMANDS: PKT SIZE

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:     MAX      MAX          PACKET SIZE
  DESTINATION LIST         PKT     WINDOW    MIN      AVG      MAX
_____________________   _______  _______  _______  _______  _______

Application Workstation / src Application Request:
  cmd Global DB Request:
  Application Server       1460        3     1000     1000     1000

Application Server / src Application Source:
  cmd Global Request Response:
  ECHO                     1460        3     1000     1000     1000
```

The Packet Size Report provides two functions:

- The first function is to report on the maximum packet and window size resulting from the socket constraints at either end for global transport and answer commands. This is important because the sockets' packet or window size constraint can reduce the maximum protocol window or packet size. The packet size is set by the transport protocol.
- The second function is to report on the packet sizes that were actually created for global transport and answer commands. This part is most useful for monitoring the sizes for messages that are smaller than the packet sizes or for monitoring the sizes from external traffic sources.

Packet Maximum

The maximum size of a packet transmitted to the destination.

Window Maximum

The maximum window size, in packets, defined by the transport protocol.

Packet Size—Minimum, Average, Max

The average, minimum and maximum observed size of a packet transmitted to the destination.

11.9 Burst Size Report

GLOBAL TRANSPORT + ANSWER COMMANDS: BURST SIZE

REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	BURST UNITS	TRAF TYPE	BURST SIZE MIN	MAX	AVG	STD DEV
Workstation 1 / src ftp:						
cmd send data:						
Workstation 2	kBits	FRLB	3	3	3	0
Workstation 1 / src ftp 2:						
cmd send data:						
Workstation 2	kBits	FRLB	21	21	21	0

The Burst Size Report provides statistics for global transport and answer commands on the burst size measured by the traffic policy on the protocol, when the traffic policy is present. It is the burst measurement that is used to determine the DE status of the packet, or for the ATM traffic policy, it is the burst measurement used to determine the conformance of the packets. Thus, a determination is made whether the packets should be immediately dropped because they exceed their traffic contract. There are three burst measurement options implemented for modeling the traffic policing. The three measurements are Sliding Window, Jumping Window, and Leaky Bucket. The burst algorithms are set in the transport protocol under traffic policing. The burst size is defined under the traffic policing option of the transport protocol.

The Sliding Window algorithm measures the burst as the total size of the packets that arrived during a fixed interval before the current time. This interval slides according to the time, and the measurement always includes all the packets that have arrived for the previous full interval.

The Jumping Window algorithm measures the burst based on accumulating a burst size with each packet, but then clearing that accumulation at fixed intervals.

The Leaky Bucket algorithm (or the Generalized Cell Rate Algorithm, GCRA for ATM) is similar to the jumping window in terms of allowing excess traffic, except that it periodically subtracts off an amount equal to the rate times the burst interval. This maintains some memory of the previous burst interval, especially if that burst exceeded the value of the burst rate times the burst interval.

Burst Units The type of units set for the burst size.

Traffic Type One of the following:

FRLB = Frame Relay Leaky Bucket
FRSW = Frame Relay Sliding Window
FRJW = Frame Relay Jumping Window
ATM
None

Burst Size—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation of the observed burst size.

11.10 Assembly Interval Report

```
                GLOBAL TRANSPORT + ANSWER COMMANDS: ASSY INTVL

                   REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN / COMMAND NAME:   PKTS            INTER-ASSY INTERVAL
  DESTINATION LIST       ASS'D      MIN         MAX         AVG      STD DEV  UNIT
_____________________   _______  _________  _________  _________  _________  __

Workstation 1 / src ftp:

  cmd send data:
  Workstation 2             102        114   29441488    2808389    5697652   uS
  Workstation 3             151        109    8029041     369997    1440288   uS
```

The Assembly Interval Report provides statistics for global transport and answer commands on the delay time between when different packets from the same message are assembled at the destination. The assembly interval can result from delay variation through the network (due to varying buffer delays) or it can result from congestion on the destination node. In cell-based networks, this measure is the cell-delay variation, a quality of service measure for these services.

The Inter-Assembly Interval is the delay between the packets when they are assembled at the destination.

Packets Assembled

The number of packets assembled at the destination.

Inter-Assembly Interval—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation for the observed inter-assembly interval.

Inter-Assembly Units

The time unit used to measure delay between the packets.

11.11 Discard Eligible Packets Report

GLOBAL TRANSPORT + ANSWER COMMANDS: DE PACKETS

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	DE 0 PACKETS CREATED	DE 0 PACKETS DROPPED	DE 1 PACKETS CREATED	DE 1 PACKETS DROPPED	CONGEST PACKETS
Application Workstation / src Application Request:					
cmd Global DB Request:					
Application Server	**48**	**0**	**0**	**0**	**0**
Application Server / src Application Source:					
cmd Global Request Response:					
ECHO	**11**	**0**	**0**	**0**	**0**

The Discard Eligible Packets Report provides information for global transport and answer commands about the flags set on the packet: the DE (discard eligibility) flag from the protocol's traffic policing algorithm, and the congestion flags set by the buffers when they are set up for FECN or BECN.

Discard Eligible 0 Packets Created and Dropped

The number of created and dropped packets flagged with a discard eligibility of 0.

Discard Eligible 1 Packets Created and Dropped

The number of created and dropped packets flagged with a discard eligibility of 1.

Congestion Packets

The number of packets that had a congestion flag set by buffers. This occurs when the buffers are set up for FECN or BECN.

12. Global Setup Commands

Global Setup Commands, when executed as part of an Application Sources's command sequence, establish a pattern of packet switched traffic, involving optional handshaking messages, between the same source and destination. See Chapter 3 in the *COMNET III Reference Guide* for more information. COMNET III produces the fourteen different Global Setup Commands reports described in the following paragraphs.

The first column in most of the Global Setup Commands reports is titled:

ORIGIN/COMMAND NAME: DESTINATION LIST

This column lists the name of the node and the name of the command for each node that has Global Setup commands defined for it and also lists the destination of the Setup Command. All destinations in the destination list of the command are listed. For Random Neighbor destinations, all nodes connected via one link to the originating node are listed.

12.1 Message Delay Report

```
                    GLOBAL SETUP COMMANDS: MESSAGE DELAY

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:  MESSAGES                MESSAGE DELAY
  DESTINATION LIST      ASSEMBLED      AVERAGE       STD DEV       MAXIMUM
_____________________   _________   _____________  _____________  _____________

Application Workstation / src Application Request:

  cmd Global Setup:
  Application Server           16   66718.771 MS     67.026 MS  66772.377 MS
```

The Message Delay Report presents message delay statistics for global setup commands on the message. Message delay is the time between creating the first packet of the message on the originating node and the time that the originating node is notified that the message has been assembled by the destination. This is the message delay that is perceived by the sender. For each originating node in the model it lists delays to each destination.

Messages Assembled

The number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of the observed message delay.

12.2 Message Delivered Report

```
                    GLOBAL SETUP COMMANDS: MESSAGE DELIVERED

                    REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:  MESSAGES                MESSAGE DELAY
  DESTINATION LIST      ASSEMBLED      AVERAGE       STD DEV       MAXIMUM
_____________________   _________   ____________  ____________  ____________

Application Workstation / src Application Request:

  cmd Global Setup:
  Application Server           16       7.066 MS      4.327 MS     14.207 MS
```

The Message Delivered Report presents statistics for global setup commands on the message delay before the message is reassembled by the destination. The message may still be worked on by the source after the message is delivered because of retransmissions, waiting for acks, or for modeling the close-sequence of the connection, but these values are not factored into the message delivered delay.

Message delay is the time between creating the first packet of the message on the originating node and the time of receiving the last packet of the message on the destination node.

Messages Assembled

For each destination, the number of messages that have been completely assembled at the destination. Messages are broken into packets at the source node according to the transport protocol characteristics and then each packet is sent to the destination. Only messages where all packets have been received are reported.

Message Delay—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of the observed message delay.

12.3 Packet Delay Report

```
                    GLOBAL SETUP COMMANDS: PACKET DELAY

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:            NUMBER OF PACKETS          PACKET DELAY (MS)
  DESTINATION LIST       CREATED DELIVERED RESENT DROPPED   AVERAGE    MAXIMUM
_____________________  ________ ________ _____ _____ _________ __________

Application Workstation / src Application Request:

  cmd Global Setup:
  Application Server          48        16     0     0     2.798       9.998
```

The Packet Delay Report presents a summary for global setup commands of the number of packets that were created, delivered, resent, or dropped. An average and maximum packet delay is also displayed.

Packet Delay is the time between creating a packet on the originating node and the time of receiving the packet at the destination node.

Packets Created

The number of packets created to send to the listed destination. Not all packets need to be delivered before they appear on the report.

Packets Delivered

The number of packets received. This may differ from the number of packets created by the number of packets that are in transit at the instant the report is written.

Packets Retransmitted

The number of packets retransmitted. Packets may be retransmitted from the origin because they are blocked at some point in route to the destination. Blocking occurs when input or output buffers are full, or when a node or link on the route fails.

Packets Dropped

The number of packets dropped. When a node or link fails, you can specify whether the transmission should be reattempted. If so, the retransmissions will be counted in the packets retransmitted field. If no retransmission is specified then the packet will be dropped.

Packet Delay (MS)—Average, Maximum

The average and maximum observed packet delay.

12.4 Setup Delay Report

```
                    GLOBAL SETUP COMMANDS: SETUP DELAY

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SETUP CMD:      SESSIONS          SETUP DELAY (MILLISECONDS)
  DESTINATION LIST         SETUP       AVERAGE       STD DEV       MAXIMUM
_____________________   _________   ____________  ____________  ____________

Application Workstation / src Application Request:

  cmd Global Setup:
  Application Server           16       7.362 MS      4.383 MS     11.923 MS
```

The Setup Delay Report presents summary statistics for global setup commands on the time required to set up a session between an originating node and a destination node. The number of sessions that were set up during the simulation along with delay statistics are reported.

Setup Delay is the time difference between creating the session setup packet and receiving back the session connect packet.

Sessions Setup

The number of sessions set up. When a session is started by a setup command, a session setup packet is sent to the destination and a session connect packet is returned. The session is then counted as set up.

Setup Delay (milliseconds)—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of the observed setup delay.

12.5 Session Length Report

GLOBAL SETUP COMMANDS: SESSION LENGTH

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN	SETUP CMD NAME	SESSIONS ENDED	SESSION LENGTH AVERAGE	STD DEV	MAXIMUM
Application Workstation / src:					
Application Reques	Global Setup	16	66.7286 S	0.0678 S	66.7798 S

The Session Length Report presents summary statistics for global setup commands on the duration of a session between an originating node and a destination node. The number of sessions completed along with the session length are reported.

Session length is the time between the session setup packet being created and the end of the session.

Origin The node which originates the session by global setup command.

Setup Command Name

The name of the global setup command.

Sessions Ended

The number of sessions that have completed and cleared. As the simulation executes, sessions are set up and messages transmitted across them. A session is counted as complete when the last message of a session has been received at the destination, and all packets for response messages have been received at the session origination node, and all acks for all packets have been received, and all pending message notices created by the receipt of the session messages have been cleared. When these conditions are met the session is complete and cleared.

Session Length—Average, Standard Deviation, Maximum

The average, maximum and standard deviation of the observed session length of completed sessions.

12.6 Setup Counts Report

```
                    GLOBAL SETUP COMMANDS: SETUP COUNTS

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / SETUP COMMAND NAME:                  NUMBER OF SESSIONS
  DESTINATION LIST                  TRIED    SETUP   RETRY  BLOCK  DISCON RERTD
_____________________________________ _______ _______ _____ _____ _____ _____

Application Workstation / src:
Application Request / cmd Global Setup:
  Application Server                    16       16      0      0      0      0
```

The Setup Counts Report presents summary statistics for global setup commands on the number of attempts to set up a session, the number of attempts that succeeded, the number of retries to set up a session based upon initial failures, the number of sessions attempted that were blocked, the number of sessions attempted that were disconnected, and the number of sessions that were rerouted due to being blocked.

Sessions Tried

The number of attempts to set up a session by respective setup command. As the simulation executes, applications which call the global session setup command are scheduled and try to execute the command. The attempt to set up a particular session will succeed or fail depending upon network conditions such as route availability inside hop and session limits, buffer availability, etc.

Sessions Setup

The number of session setup attempts that succeeded.

Sessions Retried

The number of sessions retried. If a setup attempt fails, the session may be retried later depending on the settings you have entered.

Sessions Blocked

The number of sessions blocked. The session setup attempt may block due to insufficient buffer space on routing nodes, or no route being available inside the hop limit, or no route being available because at least 1 link on all routes is at its session limit, or nodes/links have failed and no route is available.

Sessions Disconnected

The number of disconnected sessions. A session in progress may be disconnected because a node or link through which it is routed fails. The session may optionally be rerouted, if this happens.

Sessions Rerouted

The number of sessions rerouted following disconnection.

12.7 Window Stats Report

```
                    GLOBAL SETUP COMMANDS: WINDOW STATS

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:    WINDOW     NO. OF  CONG     PKT CREATION INTERVAL (MS)
  DESTINATION LIST       AVG    MAX  RESETS  AVOID     AVG      MAX      STD DEV
_____________________  _______  ___  _______  _______  _______  _______  _______

Application Workstation / src Application Request:
  cmd Global Setup:
  Application Server      2.00    2       0       0         8       13        4
```

The Window Stats Report presents the statistics for global setup commands on the TCP/IP congestion window that varies as the connection starts up and when congestion is detected. It also presents the statistics for the interval between packets being created. Data for computer group nodes is aggregated.

Window Size—Average, Maximum

The average and maximum size of the TCP/IP congestion window, in packets, that varies as the connection starts up and congestion is detected.

Number Of Congestion Resets

The number of times that the TCP/IP "congestion" window (dynamic window size) is reset to 1 packet. Thus, congestion window size is reset.

Number Of Congestion Avoidances

After a reset, the window size starts to grow more slowly (linearly) than the initial "slow start" window growth. This algorithm switch is counted as "Congestion Avoidance" and gives a measure of how often the network is recovering from congestion. If Congestion Avoidance is nearly the same as the Number of Congestion Resets, then congestion may be spurious or short-lived. If there are lots of Congestion Resets and few Congestion Avoidances, then either the network is very congested or the message sizes are not very large.

Packet Creation Interval (MS)—Average, Maximum, Standard Deviation

The average, maximum and standard deviation of the observed time interval between the creation of packets.

12.8 Retransmissions Report

```
                    GLOBAL SETUP COMMANDS: RETRANSMISSIONS

                   REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:  NO. OF    NO. PACKETS    PKT RETRIES   BLCKD AFTER T-O
  DESTINATION LIST      EVENTS     AVG    MAX    AVG     MAX    PKTS     ACKS
_________________________ _______ _______ ____ _______ ____ _______ _______

Workstation 1 / src ftp:
Workstation 1 / src ftp 2:

  cmd send data stream:
  Workstation 3             979     1.00     1     0.00     0        0        0
  Workstation 2             256     1.00     1     0.00     0        0        0
```

The Retransmissions Report presents retransmission statistics for global setup commands when blocked packets are retransmitted. The retransmission timer accounts for the age of the packet before it is blocked. It schedules a packet retransmission only after the packet is actually blocked somewhere in the network. The retransmission time is adjusted so that the packet is retransmitted the specified delay after the packet was originally transmitted. It reports on quantities such as the number of times a packet has to be retransmitted, and the number of packets that are retransmitted when a blocked packet occurs.

Number Of Events

The number of retransmissions.

Number Packets—Average, Maximum

The average and maximum number of retransmitted packets.

Packet Retries—Average, Maximum

The average and maximum number of times a packet was resent.

Packets Blocked After T-O

The number of packets blocked after timing out.

Acknowledgments Blocked After T-O

The number of acknowledgments blocked after timing out.

12.9 Timeout Report

GLOBAL SETUP COMMANDS: TIMEOUT

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	RETRANSMIT TIMEOUT (MS) AVG	RETRANSMIT TIMEOUT (MS) MAX	ROUND TRIP TIME (MS) AVG	ROUND TRIP TIME (MS) MAX
Application Workstation / src Application Request:				
cmd Global Setup:				
Application Server	0.000	0.000	66692.633	66756.569

The Timeout Report report presents time-out timer statistics for global setup commands. This report is most relevant for enhanced sliding window or TCP/IP window protocols that adapt their time-out timers based on measured round-trip delays. Because the sliding window algorithms acknowledge each packet, they have an opportunity to measure round-trip packet delay and use that delay to adjust the retransmission time out timer. The timer is only sampled when blocked packets are retransmitted.

Retransmit Timeout (MS)—Average, Maximum

The average and maximum timeout before a retransmission of the packet.

Round Trip Time (MS)—Average, Maximum

The average and maximum round trip delay time for a packet to get to a destination.

12.10 Ack Delay Report

GLOBAL SETUP COMMANDS: ACK DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME: DESTINATION LIST	NUMBER OF ACKS CREATED	NUMBER OF ACKS DROPPED	ACK DELAY (MS) AVERAGE	ACK DELAY (MS) MAXIMUM	ACK DELAY (MS) STD DEV
Application Workstation / src Application Request:					
cmd Global Setup:					
Application Server	48	0	66692.633	66756.569	71.898

The Ack Delay Report complements the packet delay report for global setup commands, by presenting statistics on the acks created and dropped as well as the delay for the ack. The Ack Delay is a round trip delay from when the packet started to when the ack returned. The time starts with the last packet that results in sending the ack.

Number of Acks Created and Dropped

The number of acks created and dropped.

Ack Delay (MS)—Average, Maximum, Standard Deviation Ack Delay

The average, maximum and standard deviation of the observed ack delay.

12.11 Packet Size Report

```
                    GLOBAL SETUP COMMANDS: PKT SIZE

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:     MAX      MAX          PACKET SIZE
  DESTINATION LIST         PKT     WINDOW    MIN      AVG      MAX
_____________________   _______  _______  _______  _______  _______

Application Workstation / src Application Request:
  cmd Global Setup:
  Application Server       1460        3     1000     1000     1000
```

The Packet Size Report provides two functions:

- The first function is to report on the maximum packet and window size resulting from the socket constraints at either end for global setup commands. This is important because the sockets' packet or window size constraint can reduce the maximum protocol window or packet size. The packet size is set by the transport protocol.
- The second function is to report on the packet sizes that were actually created for global setup commands. This part is most useful for monitoring the sizes for messages that are smaller than the packet sizes or for monitoring the sizes from external traffic sources.

Packet Maximum

The maximum size of a packet transmitted to the destination.

Window Maximum

The maximum window size, in packets.

Packet Size—Minimum, Average, Max

The minimum, average and maximum observed size of a packet transmitted to the destination.

12.12 Burst Size Report

```
                    GLOBAL SETUP COMMANDS: BURST SIZE

                  REPLICATION 1 FROM 0.0 TO 60.0 SECONDS

ORIGIN / COMMAND NAME:   BURST   TRAF          BURST SIZE
  DESTINATION LIST       UNITS   TYPE    MIN       MAX      AVG    STD DEV
_____________________  _______  ____  _______  _______  _______  _______

Workstation 1 / src ftp:
Workstation 1 / src ftp 2:

  cmd send data stream:
  Workstation 2         kBits    FRLB        2       56       31        0
  Workstation 3         kBits    FRLB        3       56       29        0
```

The Burst Size Report provides statistics for global setup commands on the burst size measured by the traffic policy on the protocol, when the traffic policy is present. It is the burst measurement that is used to determine the DE status of the packet, or for the ATM traffic policy, it is the burst measurement used to determine the conformance of the packets. Thus a determination is made whether the packets should be immediately dropped because they exceed their traffic contract. There are three burst measurement options implemented for modeling the traffic policing. The three measurements are Sliding Window, Jumping Window, and Leaky Bucket. The burst algorithms are set in the transport protocol under traffic policing. Burst size is set under the traffic policing option of the transport protocol.

The Sliding Window algorithm measures the burst as the total size of the packets that arrived during a fixed interval before the current time. This interval slides according to the time, and the measurement always includes all the packets that have arrived for the previous full interval.

The Jumping Window algorithm measures the burst based on accumulating a burst size with each packet, but then clearing that accumulation at fixed intervals.

The Leaky Bucket algorithm (or the Generalized Cell Rate Algorithm, GCRA for ATM) is similar to the jumping window in terms of allowing excess traffic, except that it periodically subtracts off an amount equal to the rate times the burst interval. This maintains some memory of the previous burst interval, especially if that burst exceeded the value of the burst rate times the burst interval.

Burst Units The type of units for the burst size.

Traffic Type One of the following:

FRLB = Frame Relay Leaky Bucket
FRSW = Frame Relay Sliding Window
FRJW = Frame Relay Jumping Window
ATM
None

Burst Size—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation of the observed burst size.

12.13 Assembly Interval Report

```
                    GLOBAL SETUP COMMANDS: ASSY INTVL

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:   PKTS            INTER-ASSY INTERVAL
  DESTINATION LIST       ASS'D      MIN        MAX        AVG     STD DEV  UNIT
_____________________  _______  _________  _________  _________  _________  __

Workstation 1 / src ftp:
Workstation 1 / src ftp 2:

  cmd send data stream:
  Workstation 2             165        395   31895643    2017873    4672067  uS
  Workstation 3             109         39   16970902     385931    1967765  uS
```

The Assembly Interval Report provides statistics for global setup commands on Inter-Assembly Intervals. The Inter-Assembly Interval is the delay between the packets when they are assembled at the destination. The Assembly Interval can result from delay variation through the network (due to varying buffer delays) or it can result from congestion on the destination node. In cell-based networks, this measure is the cell-delay variation, a quality of service measure for these services.

Packets Assembled

The number of packets assembled at the destination.

Inter-Assembly Interval—Minimum, Maximum, Average, Standard Deviation

The minimum, maximum, average and standard deviation of the observed Inter-Assembly Interval.

Inter-Assembly Units

The time unit used to measure the Inter-Assembly Interval.

12.14 Discard Eligible Packets Report

```
                    GLOBAL SETUP COMMANDS: DE PACKETS

                  REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN / COMMAND NAME:      DE 0 PACKETS          DE 1 PACKETS       CONGEST
  DESTINATION LIST       CREATED   DROPPED     CREATED    DROPPED    PACKETS
_____________________  _________  _________  _________  _________  ________

Application Workstation / src Application Request:
  cmd Global Setup:
  Application Server           48          0          0          0         0
```

The Discard Eligible Packets Report provides information for global setup commands about the flags set on the packet: the DE (discard eligibility) flag from the protocol's traffic policing algorithm, and the congestion flags set by the buffers when they are set up for FECN or BECN.

Discard Eligible 0 Packets Created and Dropped

The number of created and dropped packets flagged with a discard eligibility of 0.

Discard Eligible 1 Packets Created and Dropped

The number of created and dropped packets flagged with a discard eligibility of 1.

Congestion Packets

The number of packets that had a congestion flag set by buffers. This occurs when the buffers are set up for FECN or BECN.

13. Miscellaneous Reports

13.1 Background Packet Flows: Packet Delay Report

BACKGROUND PACKET FLOWS: PACKET DELAY

REPLICATION 1 FROM 0.0 TO 300.0 SECONDS

ORIGIN, APP, PROTOCOL DESTINATION	NUMBER OF PACKETS CREATED	 DELIVERED	 RESENT	 DROPPED	PACKET DELAY (MS) AVERAGE	 MAXIMUM
Call Node 1, Novell NetWare, IPX						
Call Node 1	2996	2996	0	0	0.000	0.000

The Packet Delay Report for Background Packet Flows provides summary packet statistics for captured external traffic imported into COMNET III through COMNET Baseliner. The number of packets created, delivered, resent and dropped, along with the average and maximum packet delay are reported for the external traffic.

The external traffic is captured first by a traffic collector such as an RMON probe or a packet analyzer. For a complete list of supported traffic capture tools please consult the *COMNET Baseliner Users Guide.*

Packet Delay is the time between creating a packet on the originating node and the time of receiving the packet at the destination node.

Origin, Application, Protocol

For each background packet flow that is imported into COMNET III an origin node is either set by the user, or automatically assigned. An application name and protocol are reported only if that information is available from the external traffic file.

Destination List

The destination of the background packet flow.

Packets Created

The number of packets created at the origin to send to the listed destination. Not all packets need to be delivered before they appear on the report.

Packets Delivered

The number of packets received at the destination. This may differ from the number of packets created, by the number of packets that are in transit at the instant the report is written.

Packets Resent

The number of packets retransmitted. Packets may be retransmitted from the origin because they are blocked at some point in route to the

destination. Blocking can occur because input or output buffers are full, or because a node or link on the route fails.

Packets Dropped

The number of packets dropped. When a node or link fails, you can specify whether the transmission should be reattempted. If so, the retransmissions are counted in the **Packets Resent** field. Otherwise, the packet is dropped.

Packet Delay (MS)—Average, Maximum

The average and maximum packet delay.

13.2 Global Traffic Command Reports

All of the reports that are available for local transport, answer, and setup commands are also available for the global commands.

The reports work a little differently, however. The local commands may be used in several sources attached to the same node but the statistics are collected at one place (on the command) and, thus, this command aggregates the statistics across multiple application sources that use this source. This subtlety can affect message delay results by aggregating unrelated delays in the same report when the command is used on multiple sources and the command uses the "use original message" text option or the command is an answer command.

The global command, however, is cloned so that each source that uses a global command has its own copy of the global command. Thus, different sources on the same node using the same global command will have separate reports for these commands. However, if the source reuses the same global command, it will reuse the same cloned copy of the command and thus aggregate the statistics of these instances. This behavior is likely to be what is desired since the reused command is probably doing something similar in each instance so that the statistics are comparable and related.

13.3 Response And Answer Destinations

The response sources and answer commands have a **Use 'ECHO' in reports** checkbox . If the checkbox is off, then the statistics are reported separately for each destination responded to by the response or answer. This is useful because the destinations may have different levels of congestion or the destinations may have different paths so that the delays are not appropriately aggregated in the same statistic.

However, this detail comes at the cost of longer initialization times and more memory than the simple **ECHO**. When statistics are not required or when it is appropriate to aggregate the statistics over all possible destinations (when all such destinations are on the same LAN segment, for instance), it is advantageous to continue using the **ECHO** option to save memory and improve the speed of the simulation.

13.4 Snapshot Reports and Alarms

Snapshot reports convey statistical information about the state of a simulated model, very much like post-run reports. They are also turned on and off using a hierarchical list interface similar to that of post-run report selection. Snapshot reports differ from post-run reports in these key ways:

1. Snapshot reports may be turned on or off, not only before the simulation starts, but also while the simulation is running.
2. The information provided by a snapshot report consists only of the value of a measurement, such as link utilization, at an instant in time or over a brief interval of time. At the next instant or interval, new information is computed and presented, and the previous information is lost. You obtain this perishable information by watching the simulation, or by selecting the menu/toolbar option **Take Snapshot**. If you need the information to be collected and saved for later use, you should instead use a post-run report.
3. Alarms can be set on snapshot performance measures. Alarm conditions ARE recorded, and alarms have the potential to trigger model traffic. Alarms can also be configured to pause the simulation.

As mentioned above, snapshot reports may be observed in one of two ways:

1. For each class of layout objects, such as nodes, you may select one snapshot report to be displayed. The value of the snapshot will be displayed as a small number or set of numbers above the icons associated with the objects. Note that the default condition is that no snapshot reports are displayed for any objects.
2. At any point in the simulation, you may select a layout object and then select the menu/toolbar option **Take Snapshot**. A dialog will appear listing the current values of all the snapshot performance measures currently turned on for that object.

Some snapshot reports, such as session counts, are updated whenever the value changes. Most, however, are updated at one of two times: at the end of a user-specified interval, or when the user selects **Take Snapshot**. In the former case, the reported value is an average of the measurement over the interval. In the latter case, the reported value is a weighted combination of the current average (over only a partial interval) and the previously reported value. The measurement interval is specified separately for nodes and links on the backbone or subnet property dialog.

Some snapshot reports, such as link utilization, involve two numbers, and when these snapshots are selected for display, you might have difficulty determining which number has which meaning. The **Take Snapshot** option presents a more verbose report that is handy for associating the numbers with their meanings.

Each snapshot report may be set to *sound* an alarm when the value exceeds a specified threshold. When the alarm sounds, the icon of the offending object will change to a specified color, and a record of the alarm condition will be kept and later added to the post-run report file. When an alarm sounds on an object inside a subnet, transit network, or cloud, the subnet/

transit net/cloud icon will be colored in the 2-D view, and will retain that color until the alarm *clears*, even if subsequent alarms of different colors sound.

Note that nodes and links can be configured to trigger model traffic upon the sounding or clearing of an alarm. This a very powerful, and very dangerous, feature. It is important to remember that, in combination with triggers, alarms have the ability to change model behavior, in which case they are more than mere instrumentation.

If you have an alarm set in your model, but you do not want to watch the simulation waiting for a color change, you can configure the alarm to force the simulation into single step mode when it sounds. To permit the simulation to continue running without further interruption, typically you must first reconfigure the alarm so it no longer forces single step mode, then use the **Trace** option to turn off single step.

As with post-run reports, your simulation will run slower when you have more snapshot reports turned on. A few snapshot reports are on by default, but these snapshots will have no effect on speed of simulation. One is the Disk Error snapshot report. It counts the number of disk errors, such as attempting to read from a non-existent file. When the number exceeds 1, an alarm sounds. For most models, a disk error indicates a mistake in the model.

Turning on **Remember values** and **Export values at end** will result in a file with the suffix **.snp** being written to the model directory. The file is a text file, either tab or comma delimited, containing each snapshot value for all snapshots turned on in the model.

Index

O

P

R

S

T

U

V

W