TCP/IP NETWORK ROUTE PERFORMANCE ANALYSIS

Including a detailed look at how AES CLEVER *e*Route 3.0 analyzes network routes and identifies the root-causes of performance problems

By David Cheng, VP of Engineering, AES - DaveC@aesclever.com

Published and Distributed by:
AES, 149 Commonwealth Drive, Menlo Park, CA 94025
(650) 617-2400
www.aesclever.com

Acknowledgements: The author would like to thank Dr. Leo Lo of IBM for his valuable contribution and also Mark Nguyen of AES, who has responsibility for development of the CLEVER *e*Route product.

TCP/IP NETWORK ROUTE PERFORMANCE ANALYSIS

INTRODUCTION

Burgeoning mission-critical applications are heavily dependent on Internet technologies. Users demand more and more functionality, with near zero downtime. These ever-increasing demands make the management of enterprise networks extremely challenging. When there is a response time issue, the user needs to know where to look for information and then exactly what to look for in order to resolve the problem. Most networks have a limited set of shared resources and there are always potential performance bottlenecks elsewhere in the network. It is counter-productive to simply eliminate a performance bottleneck, only to have it move to some other resource.

Internet performance can cover a wide variety of components, each of which can affect overall network performance. Applications, Web servers, Intranets, Local Area Networks (LANs), protocols, operating systems, and Wide Area Networks (WANs) can have an impact on overall Internet performance. This paper discusses only the network-related performance concerns, such as TCP/IP routing analysis.

The Hyper Text Transfer Protocol (HTTP) is an application-layer protocol sitting on top of the TCP/IP to establish the communication path between a Web surfer and Web server. HTTP maps the server name to an IP address, establishes a communication path with the Web server, and then sends the request using the Uniform Resource Locator (URL). It also receives the response from the Web server (e.g., Hyper Text Markup Language or imaging document). When the communication is complete, HTTP closes the TCP/IP connection. One needs to understand the potential performance impact of the HTTP and TCP/IP. For example, the protocol latency delay may consist of path delay and processing delay. Path delay refers to the time required to establish a path between the Web surfer and the Web server. The process delay refers to the request processing time across the established path.

Analyses of TCP/IP routes help network administrators and performance analysts identify the root-causes of performance problems. For example, a problem that may appear to be caused by poor transaction response time may actually be caused by network route congestion, gateway failure, packet loss, or an unreachable destination, making performance data collection just as important as data analysis. Without a good set of data with statistical significance, the results of analyses may be misinterpreted and problems may not be resolved correctly. This could result in unnecessary tuning or capacity upgrades.

PERFORMANCE DATA COLLECTION ARCHITECTURE

Transaction failure, transaction performance, and network resource usage data can be useful for analyzing TCP/IP transactions. Analyses are performed at the TCP/IP session level. One possible implementation of data collection architecture is a three-layer hierarchy that consists of a centralized Master Collector & Processor, regional Collectors, and Remote Collector Agents. The diagram below shows the structure.

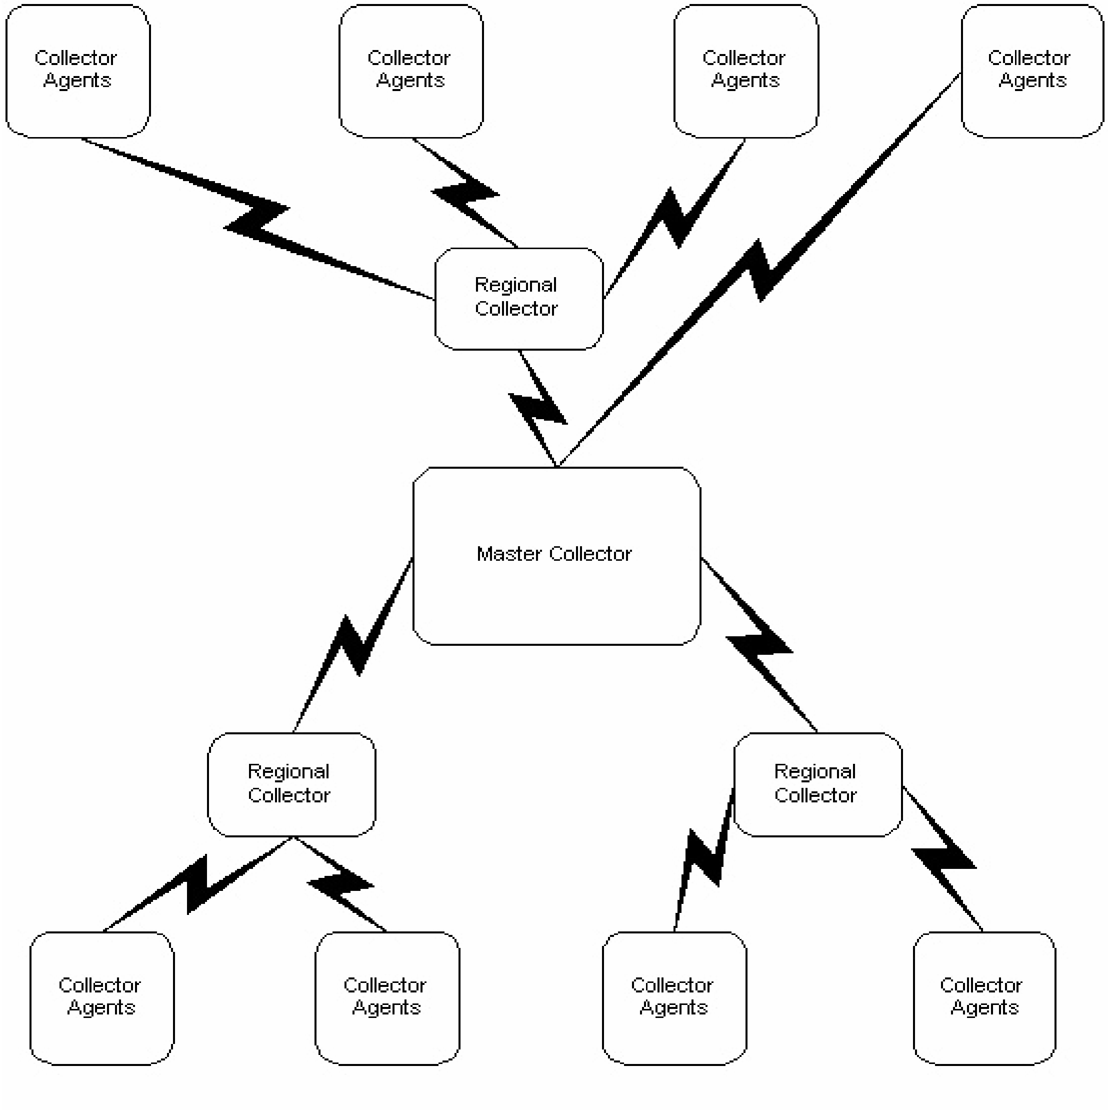

Each collector agent gathers TCP/IP transaction data using the TraceRoute command on a set of network nodes assigned under it. The Command transmits packets to each network node that constitutes a path or route between the data collector and the network node. The response time information on that path or route, along with the number of segments (also referred to as hops, the path between two network nodes), and the host node detected along the path will be collected and saved. The remote data collector agents transmit saved data at predetermined time intervals to their regional collectors. The regional collectors then transmit the gathered performance data to the Master Collector. The Master Collector can process the raw data and produce either scheduled or ad hoc reports.

3-Layer Data Collection Architecture

Three-layer architecture is somewhat complex for most implementations, however, and the collection model can be simplified and reduced to two layers in even the largest networks, whereby the Master Collector can directly control the remote data collection agents. A typical application of such architecture might be to install the Master Collector on a Windows system and remote data collection agents on Windows/LINUX systems, with the master collector also providing the analysis function. Multiple Master Collectors can be implemented in such a scenario, if necessary, making the two-tier structure ideal. CLEVER *e*Route, developed by AES, is an extremely good implementation of the two-layer architecture just described.

Three categories of data provide detailed insight, enabling network analysts to identify the cause(s) of the problem or the resource affected by a performance bottleneck:

- Transaction Failure data provides the times and types of failures, such as packet loss, looping, or connection failure, for each route or segment.
- Transaction performance, analyzed per segment or route, tracks the response time distribution over a selected time period, including peaks and valleys of transaction volumes.
- Network resource usage gives an indication of productive and non-productive use of network resources with detected peaks and valleys.

Network analysts can quickly isolate failed routes and segments by viewing these categories of performance data along with the overall network topology. Transaction volume statistics, such as patterns of shared resource usage dominance, can also be useful for better managing future network capacity growth.

One needs to be aware that TCP/IP route/path analysis only covers one aspect of overall TCP/IP performance analysis. It can be used in conjunction with other types of tools such as Web Performance analysis of selected URL pages or Web page usage trending analysis, to address specific areas of concern.

PERFORMANCE DATA COLLECTION

A solid data collection system design is vital to the success of subsequent performance analysis. There are four basic procedures for data collection and analysis:

- Configuration – the master data collector can configure all remote collector agents by issuing a Configure Command. Any remote data collector agent can be re-assigned to a different master data collector, if present.
- Collection – the master data collector can issue either real-time or periodic scheduled collections to data collection agents, each of which collects data locally on those network nodes where a collection module is initiated.
- Analysis – as the collection module is invoked, collected data are stored in the application's subdirectory for identification purposes. The collection module also imports, analyzes, and generates reports from the collected data.
- Identification – each physical network device can be identified via the device's Simple Network Management Protocol (SNMP) Management Information Base (MIB) information such that data can be correlated to the physical device.

Depending upon where and how the data collection mechanism is set up, users can select certain data collection parameters. The TraceRoute Command can pick up the Maximum Time to Live (TTL) or Number of Maximum Network Hops value for the analysis. The Timeout value in seconds specifies how long to wait for a response. The Interval value in minutes defines the time between collecting two consecutive samples. The Sample value determines the total numbers of samples that need to be collected. Users can select a particular time interval on each day to collect TCP/IP route data or can run the data collection continuously. There is no ideal data-sampling interval or number of samples since each networking environment has different characteristics. Data could be collected for example several times a minute, or every five minutes, and summarized every 15 minutes.

PERFORMANCE DATA REPORTING

(a) For route or Segment Analysis Reporting

Since a route failure may caused by several factors, analysis reports should cover, as a minimum, the following information, from which defective routes, segments and devices can be further identified:

- Session identifier assigned to the session during the collection period
- Route identifier assigned to the route during the analysis
- Target/Initiating Host as the starting point of the route including its IP address
- Receiving/ending Host of the router including its IP address
- Number of times the route was analyzed and reported
- Date on which the route failure occurred
- Time at which the route failure occurred
- Segment identifier assigned to the segment during the analysis & reporting period
- Response time in milliseconds – the time to takes to perform a Trace Route Command from the initiating host to the receiving host
- Data Collection time (time at which the route/segment information was collected)
- Data Collection date (date on which the route/segment information was collected)
- Peak/Valley times, including the minimum, average and maximum response times based on the total number of samplings collected for the route. The average response time is calculated by averaging the total response time of all hops preceding the current hop and subtracting that from the current hop’s response time.
- Number of successful segments in the route under analysis
- Number of failed segments in the route under analysis
- Percentage of segment availability (the ratio between successful and failed segments)
- Status (route/segment operating status for availability reporting)
- Failure device type, based on the highest failing, looping, or packet losing routes

(b) SNMP MIB Data Analysis Reporting

It is worthwhile to note that SNMP was designed to accommodate different network platforms, various protocols, and proprietary operating systems. SNMP exchanges network information through messages known as Protocol Data Units (PDUs). There are five useful PDUs for network monitoring: two deal with setting terminal data, two deal with reading terminal data, and one (called trap) is used for monitoring network events such as start-up and shut-down. There are two major types of MIBs: public and private. Public MIBs contain information that most network vendors support and private MIBs contain network vendor-specific information.

Reports that utilize public MIB information should contain at a minimum the following information:

Border Gateway Protocol (BGP)

Used for exchanging routing information between gateway hosts, each with its own route in a network of autonomous systems. It should contain the BGP version, local autonomous system number, and BGP identifier of the local system. Each BGP should also contain a table of its peer BGP information such as:

- Peer BGP ID and Status
- Local IP Address and Port
- Remote IP Address, Port, and Autonomous Systems
- Messages-In
- Messages-Out
- Connection Retry
- Peer BGP's Holding Time Received
- Keep-Alive Time
- Time-To-Live (TTL) Values

CSS Bridge

Contains a list of information for each port of a bridge. It should contain the following information:

- Port Number where the bridge management information is kept
- Interface Index for the interface corresponding to this port
- Port Circuit which is used to distinguish the same value used by another port on the same bridge
- Delay-Exceeded Discards (the number of frames discarded by this port due to excessive transit delay through the bridge)
- Size-Exceeded Discards (the number of frames discarded by this port due to an excessive size)

CSS Ethernet

Contains statistics for interface to an Ethernet-link medium. It should contain such information as the following:

Index — A unique identifier

Alignment Errors — The count on frames received without an integral number of octets in length

Frame Check Sequence (FCS) — Refers to the extra characters added to a frame for detecting and correcting errors

Errors — The count of frames due to FCS check

Single Collision Frames	A count of successfully transmitted frames due to a single collision
Multiple Collision Frames	A count of successfully transmitted frames that have more than one collision
Deferred Transmissions	A count of frames that experienced delay not due to collision
Excessive Collisions	A count of frames whose transmission failure was due to excessive collisions
Carrier Sense Errors	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame
Frame Too Long	A count of frames that exceed the maximum permitted frame size
Internal MAC Receive/Transmit Errors	A count of frames failed due to either MAC sub-layer receiver/transmitter error

CSS RMON Contains remote network monitoring statistics information. It should contain such information as the following:

- Data Source (any Ethernet interface on this device)
- Drop Events (the total number of events that probe packets are dropped due to lack of resources)
- Octets (total number of octets of data)
- Packets (total number of packets including bad, broadcast and multicast packets received)
- Undersize Packets (<64 octets)
- Oversize Packets (>1518 octets)
- Fragments (total number of packets received whose length is less than 64 octets with a bad FCS)
- Collisions (total number of estimated collisions on this Ethernet segment)
- Various packet sizes (counts on packets between 64 to 1518 sizes)
- Status

Internet Control Message Protocol (ICMP)

Uses 26 counters to keep track of the message types generated and received by the local ICMP entity including received, sent, received in error, or not sent due to error. The basic information should contain the following:

- Packet-In
- Packet-Out
- Errors
- Timing (time stamps, replies)
- ICMP Address Mask Messages (requested & received)

Interface

Contains generic information on the interface layers, such as:

- Interface Description
- Type
- Maximum Transmission Unit (MTU) size
- Transmission Rate (speed, number of INs & OUTs)
- Media Specific Address
- Administrative and Operational Status
- Discards (IN & OUT)
- Out Queue Length

IP MIB

Contains information on the IP subsystem of a managed node (e.g., if the device is acting as a router or a host, packets or datagram delivered, discarded, time-out values, etc.). It should contain such MIB information as:

- Forwarding Status
- Default Time-To-Live Value
- IN (received, delivered, discarded)
- OUT (requested, forwarded, discarded)
- Errors (header, address, unknown protocol, no routes)
- Assembly (required, success, failed, time-out)
- Fragment (success, failed, created)

IP Route

Contains the following information:

Destination address — Destination address, on this route and various routings, for primary routing, four alternative routes, Next Hop, etc.

Protocol — Learning various route protocols on possible four values: none, local, network management or obtained via ICMP

Age — Time in seconds since last route update

Mask — For comparison of destination address

Reference — Specific to the particular routing protocol responsible for this route

Open Shortest Path First (OSPF) Used with large autonomous networks in preference to the Routing Information Protocol (RIP). It should contain:

Router information	A 32-bit integer for the router in the autonomous system, and whether it is an area border router
OSPF Status	Version, administrative status
Link State Advertisements (LSA)	External LSA, Number of Originated & Received new LSAs, Exit Overflow Interval, etc.

Within OSPF, there are six sub-groups; each providing further information on OSPF performance:

OSPF Area

Contains a table with information on the configuration parameters and cumulative statistics of a router's attached areas. Information such as the following should be included:

- Authentication Types
- Imported Autonomous Systems External Link State
- Reachable Board Routers
- Area Status

OSPF Link State Data Base

Contains a table with information about the link state announcement from attached neighboring areas. Information such as the following should be included:

- Area IDs
- Type
- Originating Router ID
- Age
- Checksum
- LSA Content

OSPF Host

Contains a table with information on directly attached hosts to the router. Information such as the following should be included:

- Host IP address
- Type of Service
- Status
- Area ID to which the host is attached

OSPF Interface

Contains a table with information describing interfaces on this OSPF. Information pertinent to interface analysis includes the following:

- Interface IP address
- Area ID
- Type
- Status
- Priority
- polling interval
- authentication type
- key

OSPF neighbor

Contains a table with information on all neighbors in the locality of this router. Information important to the analysis includes the following:

- neighbor ID
- router IP address
- priority
- operating status

OSPF Area Advertisement Table

Replaces the OSPF Summary table containing the following Area Advertisement information:

- ID
- Type
- IP Addresses
- Subnet Mask
- Effect (advertise matching or not matching)

Printer Alert

Contains a list of critical and non-critical alerts currently active in the printer. A critical alert is one that stops the printer from printing immediately and prevents further printing. Non-critical alerts are those that do not stop printing, but may become critical at a future time. This table should contain information on severity, component and its detailed location, alert code, and a description of each currently active critical alert.

TCP The Transmission Control Protocol MIB information identifies the retransmission algorithm, maximum and minimum retransmission time-outs, and number of active/passive opens, resets, connects, etc. Information important to the analysis includes the following:

- Segment–In/Out
- Retransmission
- Numbers of Attempts
- Opens
- Resets
- Maximum Connections
- Time-Out Value (Min and Max)

UDP User Datagram Protocol (UDP) provides four counters and a table for datagram delivery, destined for unknown ports, discarded due to format errors, and send from UPD group. Information important to the analysis includes the number of total datagrams delivered/sent, and errors types.

APPLICATIONS & EXAMPLES

A two-layer architecture of the data collection system, CLEVER *e*Route, is used to demonstrate the application of TCP/IP route analysis. A Master Module can be installed on a workstation where the data collection commands and reports are generated. A number of data collection agents, or eCollectors, are installed on remote Linux or Windows machines. The eCollectors, in addition to performing route information collection, are also capable of auto-discovering local IP addresses, executing run-time scripts, launching the collection scheduler, and maintaining data logs.

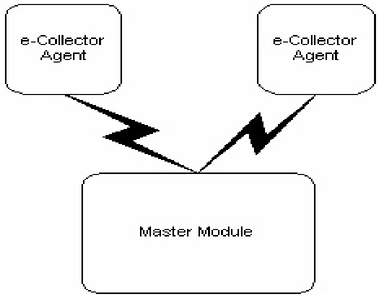

The 2-LayerArchitecture Employed by CLEVER eRoute

CLEVER *e*Route obtains its data from the *e*Route Host Collector (if present), the eCollector (local to the Master), and the remote eCollector Agents. Each eCollector Agent performs a TraceRoute command from the remote starting point to a set of network nodes (any IP-addressable device). The TraceRoute command transmits packets to each node along the path that constitutes the route and outputs a millisecond response time for each packet. Each response time associated with each network node is collected and saved in an output file that provides raw data for CLEVER *e*Route to analyze. The output file provides information on route performance, number of segments (the path between two network nodes, also called hops), and each host detected along the route. From this information CLEVER *e*Route analyzes the data, systematically organizing and presenting it to the user in an accessible and easily understandable format. Among its unique features, CLEVER eRoute has the ability to collect performance statistics *from* anywhere in the network *to* anywhere in the network. Its eCollectors can run outbound and inbound simultaneously, facilitating the analysis of return routes.

Analysis Reports

The following analysis reports are available for review:

- Most Failing Route
- Most Failing Segment
- Failures by Time of Day
- Packet-Loss by Time of Day
- Most Packet-Losing Route
- Most Packet-Losing Segment
- Looping by Time of Day
- Most Looping Route
- Most Looping Segment
- Route with Best Performance
- Segment with Best Performance
- Route with Worst Performance
- Segment with Worst Performance
- Performance by Time of Day
- Route Performance Distribution
- Segment Performance Distribution
- Most Dominant Route
- Most Shared Host(s)
- Most Shared Segment(s)

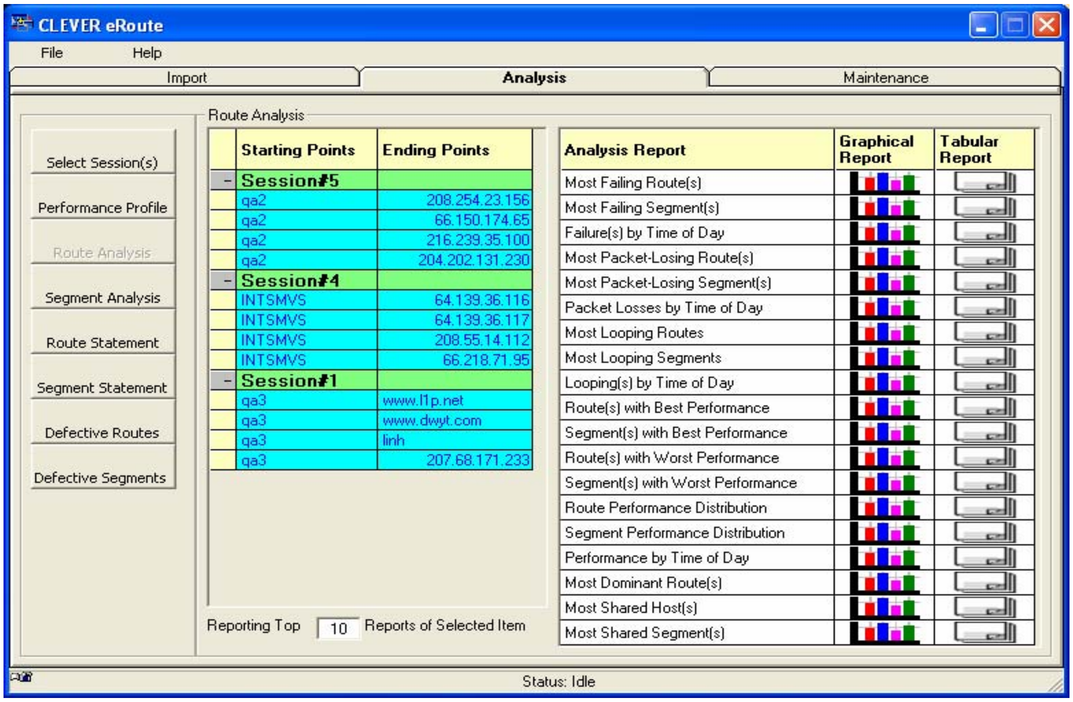

Route Analysis

	Starting Points	Ending Points
-	Session#5	
	qa2	208.254.23.156
	qa2	66.150.174.65
	qa2	216.239.35.100
	qa2	204.202.131.230
-	Session#4	
	INTSMVS	64.139.36.116
	INTSMVS	64.139.36.117
	INTSMVS	208.55.14.112
	INTSMVS	66.218.71.95
-	Session#1	
	qa3	www.l1p.net
	qa3	www.dwyt.com
	qa3	linh
	qa3	207.68.171.233

Reporting Top 10 Reports of Selected Item

Analysis Report	Graphical Report	Tabular Report
Most Failing Route(s)	📊	▭
Most Failing Segment(s)	📊	▭
Failure(s) by Time of Day	📊	▭
Most Packet-Losing Route(s)	📊	▭
Most Packet-Losing Segment(s)	📊	▭
Packet Losses by Time of Day	📊	▭
Most Looping Routes	📊	▭
Most Looping Segments	📊	▭
Looping(s) by Time of Day	📊	▭
Route(s) with Best Performance	📊	▭
Segment(s) with Best Performance	📊	▭
Route(s) with Worst Performance	📊	▭
Segment(s) with Worst Performance	📊	▭
Route Performance Distribution	📊	▭
Segment Performance Distribution	📊	▭
Performance by Time of Day	📊	▭
Most Dominant Route(s)	📊	▭
Most Shared Host(s)	📊	▭
Most Shared Segment(s)	📊	▭

Status: Idle

CLEVER eRoute Network Route Analysis

The Route Analysis menu shows two sections of information. The first one provides the selection options for display (e.g., performance profile of the selected session, route segments, segment statement, defective routes and defective segments). The starting and ending points for the selected sessions are listed for selection. The second one, the right side table, lists all the available reports with options for graphical or tabular formats. An example of a graphic format report is shown below:

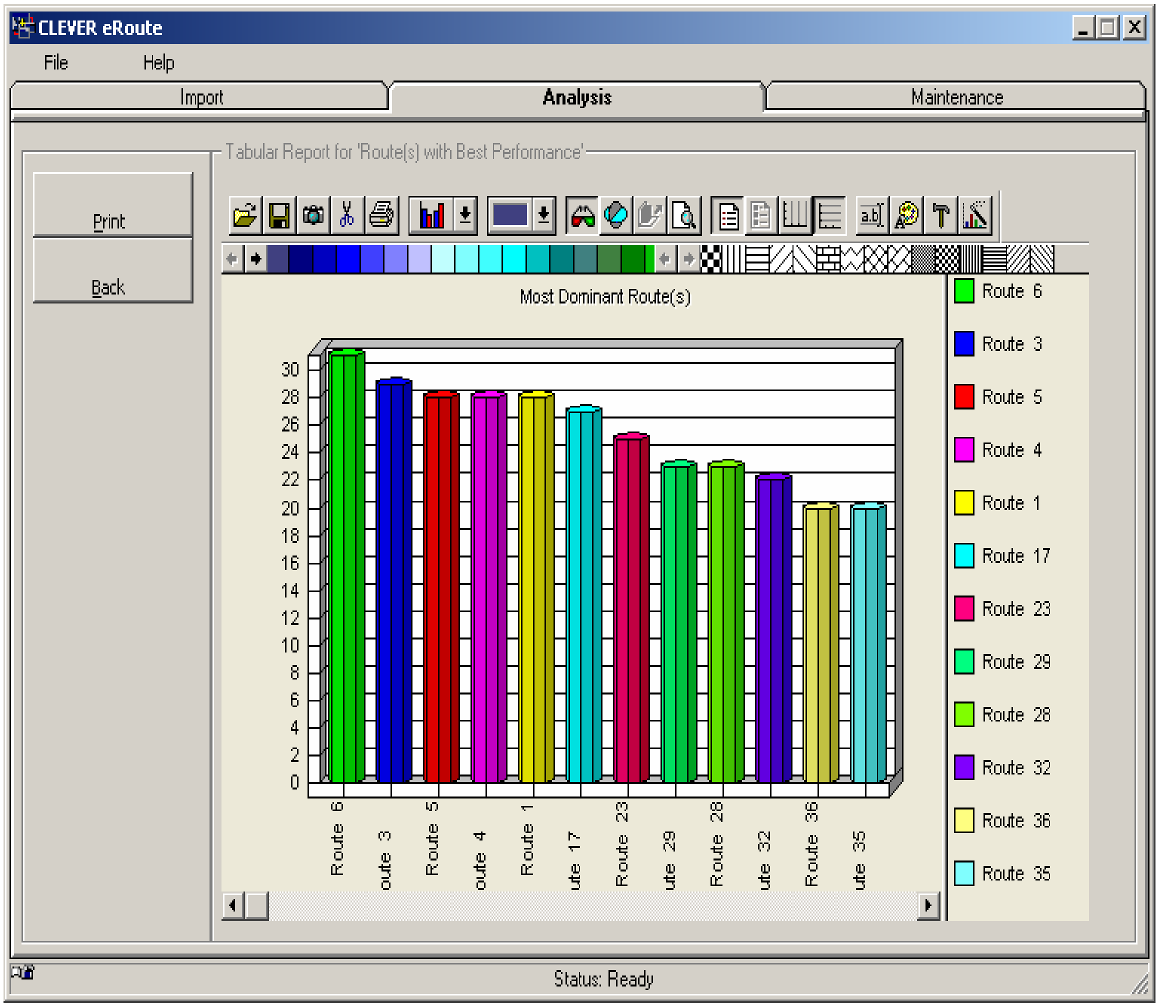

An Example of CLEVER eRoute Graphic Reports

The appropriate tabs can be selected to view the best and worst routes and segments. The report contains the following information:

- Session – the identifier assigned to the session during the data collection period
- Route ID – the identifier assigned to the route during the data collection period
- Segment ID – the identifier assigned to this segment during data collection period
- Response time – the time it takes to perform a TraceRoute Command from the starting point to the ending point, in milliseconds
- Frequency – the number of times the route was reported
- Collect Time – the time at which the route information was collected
- Collect Date – the date on which the route information was collected
- Starting Host – starting point for the route
- Ending Host – ending point for the route

The report can show the minimum, maximum, and average response time to illustrate Performance by Time of Day:

- Minimum RT (ms) – the minimum response time for the segment based upon the total number of data scans performed for the route
- Maximum RT (ms) – the maximum response time for the segment based upon on the number of data scans performed for the route
- Averaged RT (ms) – the averaged response time for the segment, which is calculated by averaging the total response times of all network hops preceding the current one and subtracting that from the current hop's response time

A Segment Analysis report for a given segment can be displayed to provide detailed performance data. Again, the display can be in either graphical or tabular format.

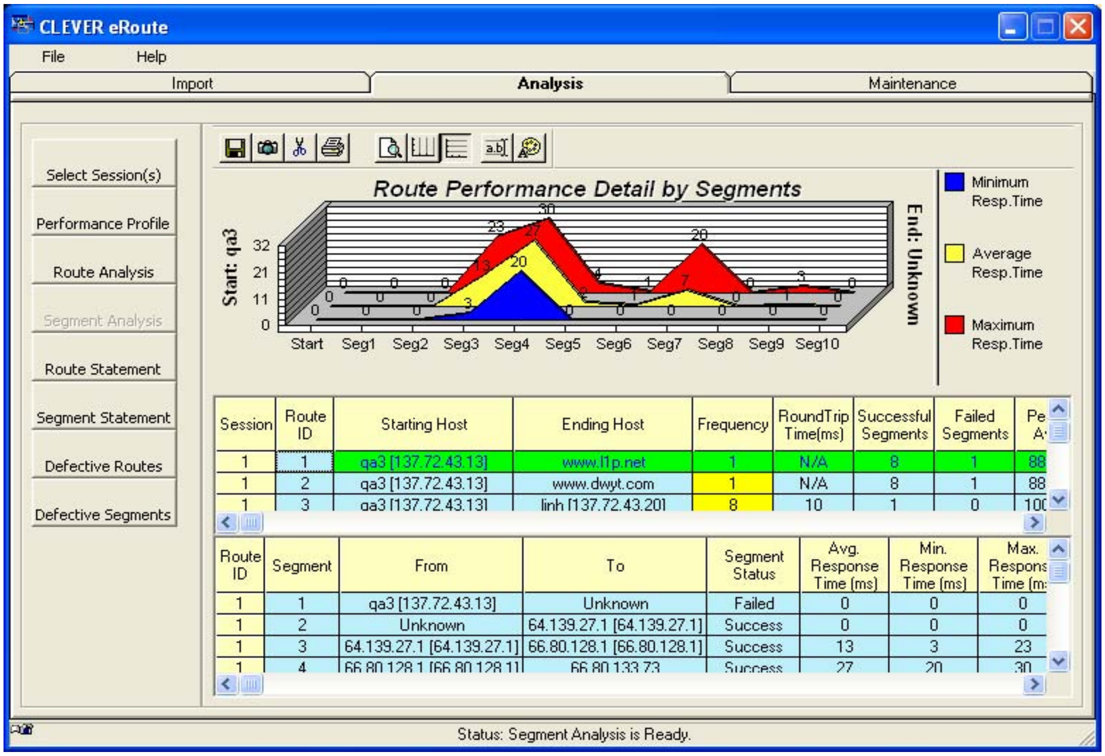

Session	Route ID	Starting Host	Ending Host	Frequency	RoundTrip Time(ms)	Successful Segments	Failed Segments	Pe A
1	1	qa3 [137.72.43.13]	www.l1p.net	1	N/A	8	1	88
1	2	qa3 [137.72.43.13]	www.dwyt.com	1	N/A	8	1	88
1	3	qa3 [137.72.43.13]	linh [137.72.43.20]	8	10	1	0	100

Route ID	Segment	From	To	Segment Status	Avg. Response Time (ms)	Min. Response Time (ms)	Max. Respons Time (m
1	1	qa3 [137.72.43.13]	Unknown	Failed	0	0	0
1	2	Unknown	64.139.27.1 [64.139.27.1]	Success	0	0	0
1	3	64.139.27.1 [64.139.27.1]	66.80.128.1 [66.80.128.1]	Success	13	3	23
1	4	66.80.128.1 [66.80.128.1]	66.80.133.73	Success	27	20	30

Status: Segment Analysis is Ready.

CLEVER eRoute Route Performance Detail by Segment

The sample graphic display report shows two sections. The upper section is a summary of the route information for the selected session with minimum, averaged and maximum response times. The lower section shows the following detailed segment performance data:

- Session ID
- Route ID
- Staring Host
- Ending Host
- Frequency
- Round Trip Time (ms) – the amount of time it took for the packet to reach its Ending Host and return to its Starting Host (if the route fails, the display will indicate N/A)
- Successful Segments – the number of segments that succeeded in the route
- Failed Segments – the number of segments that failed due to unavailability of the next IP address in the amount of time specified for the data collection

- Percent Availability - the ratio between successful and failed segments
- Status – displays complete or partial, depending upon the successful completion or failure of the route

In the MIB data analysis, there are several very useful reports for route performance analysis.

CSS-Ethernet Report

Some of the data fields in the CSS-Ethernet report (see example below) can provide insight into Ethernet interface performance. The Single Collision Frames column provides the count of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision. The Late Collision field shows the number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet. The Excessive Collisions field points towards a particular interface failing due to excessive collisions. The Frames Too Long data field shows the count of frames received on a particular interface that exceed the maximum permitted frame size.

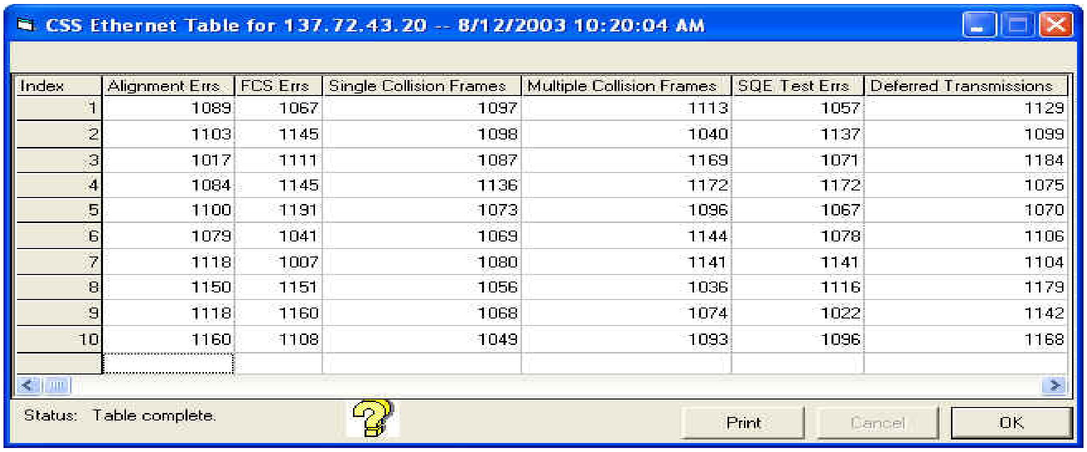
CSS Ethernet Table for 137.72.43.20 -- 8/12/2003 10:20:04 AM

Index	Alignment Errs	FCS Errs	Single Collision Frames	Multiple Collision Frames	SQE Test Errs	Deferred Transmissions
1	1089	1067	1097	1113	1057	1129
2	1103	1145	1098	1040	1137	1099
3	1017	1111	1087	1169	1071	1184
4	1084	1145	1136	1172	1172	1075
5	1100	1191	1073	1096	1067	1070
6	1079	1041	1069	1144	1078	1106
7	1118	1007	1080	1141	1141	1104
8	1150	1151	1056	1036	1116	1179
9	1118	1160	1068	1074	1022	1142
10	1160	1108	1049	1093	1096	1168

Status: Table complete.

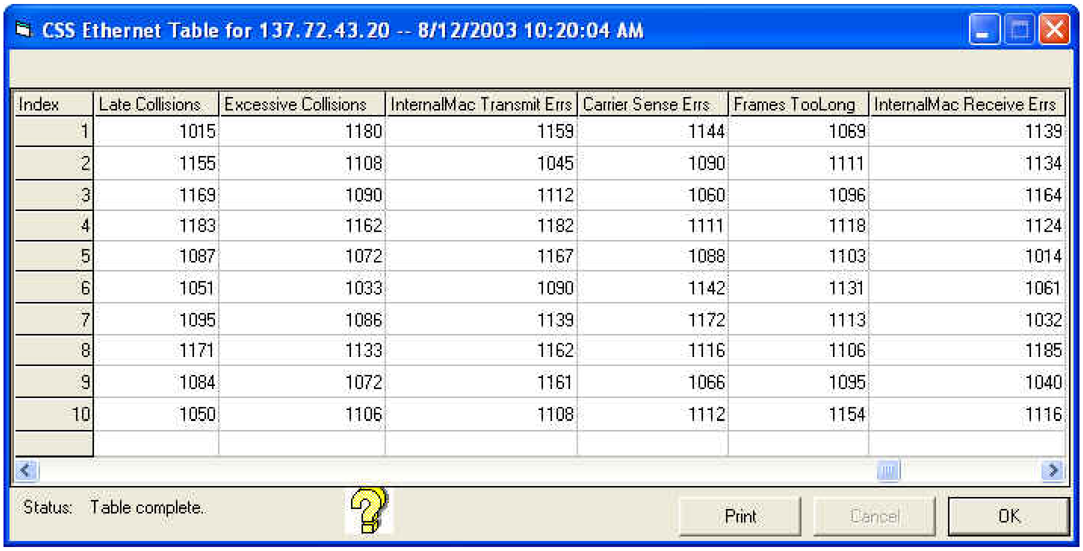
CSS Ethernet Table for 137.72.43.20 -- 8/12/2003 10:20:04 AM

Index	Late Collisions	Excessive Collisions	InternalMac Transmit Errs	Carrier Sense Errs	Frames TooLong	InternalMac Receive Errs
1	1015	1180	1159	1144	1069	1139
2	1155	1108	1045	1090	1111	1134
3	1169	1090	1112	1060	1096	1164
4	1183	1162	1182	1111	1118	1124
5	1087	1072	1167	1088	1103	1014
6	1051	1033	1090	1142	1131	1061
7	1095	1086	1139	1172	1113	1032
8	1171	1133	1162	1116	1106	1185
9	1084	1072	1161	1066	1095	1040
10	1050	1106	1108	1112	1154	1116

Status: Table complete.

CLEVER eRoute CSS Ethernet Reporting

CSS – RMON Report

The CSS RMON report provides remote network monitoring statistics, as seen in these examples:

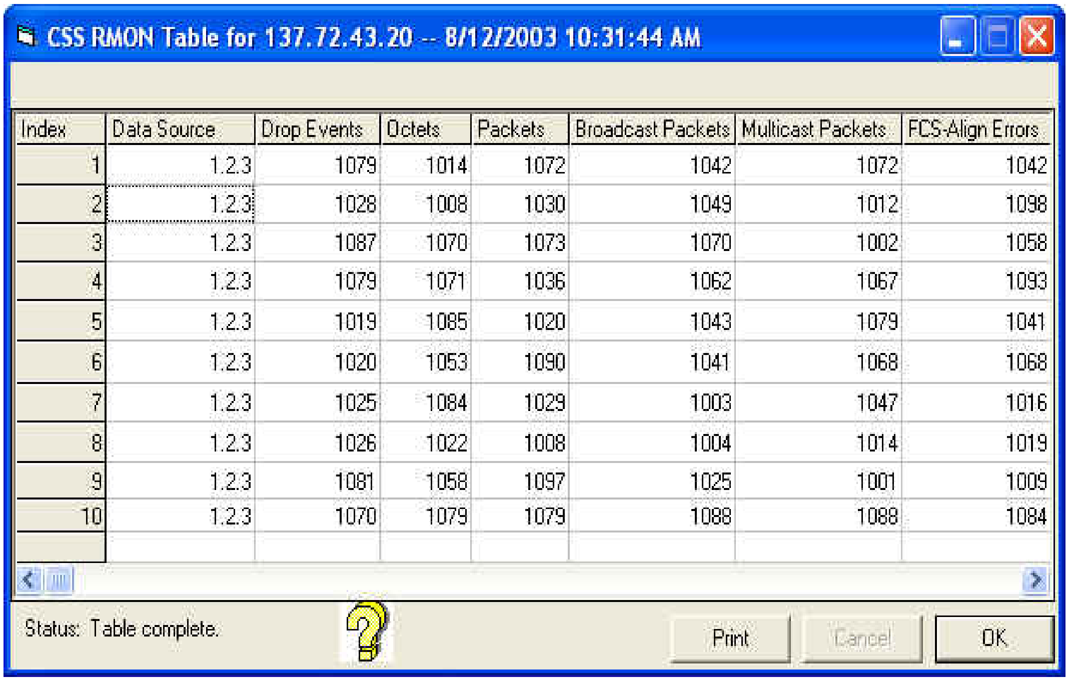

CSS RMON Table for 137.72.43.20 -- 8/12/2003 10:31:44 AM

Index	Data Source	Drop Events	Octets	Packets	Broadcast Packets	Multicast Packets	FCS-Align Errors
1	1.2.3	1079	1014	1072	1042	1072	1042
2	1.2.3	1028	1008	1030	1049	1012	1098
3	1.2.3	1087	1070	1073	1070	1002	1058
4	1.2.3	1079	1071	1036	1062	1067	1093
5	1.2.3	1019	1085	1020	1043	1079	1041
6	1.2.3	1020	1053	1090	1041	1068	1068
7	1.2.3	1025	1084	1029	1003	1047	1016
8	1.2.3	1026	1022	1008	1004	1014	1019
9	1.2.3	1081	1058	1097	1025	1001	1009
10	1.2.3	1070	1079	1079	1088	1088	1084

Status: Table complete.

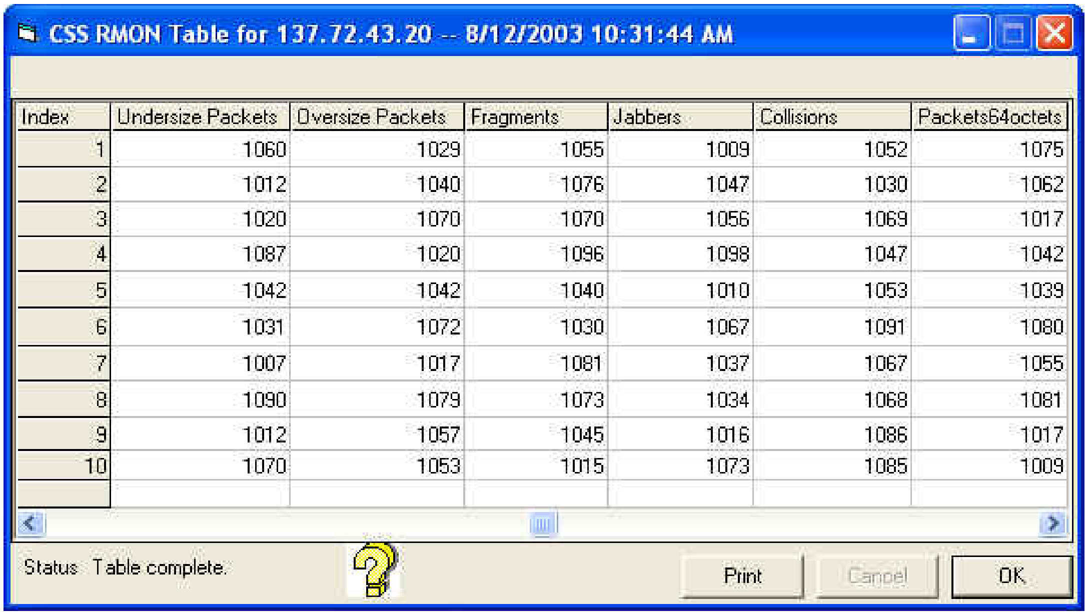

CSS RMON Table for 137.72.43.20 -- 8/12/2003 10:31:44 AM

Index	Undersize Packets	Oversize Packets	Fragments	Jabbers	Collisions	Packets64octets
1	1060	1029	1055	1009	1052	1075
2	1012	1040	1076	1047	1030	1062
3	1020	1070	1070	1056	1069	1017
4	1087	1020	1096	1098	1047	1042
5	1042	1042	1040	1010	1053	1039
6	1031	1072	1030	1067	1091	1080
7	1007	1017	1081	1037	1067	1055
8	1090	1079	1073	1034	1068	1081
9	1012	1057	1045	1016	1086	1017
10	1070	1053	1015	1073	1085	1009

Status: Table complete.

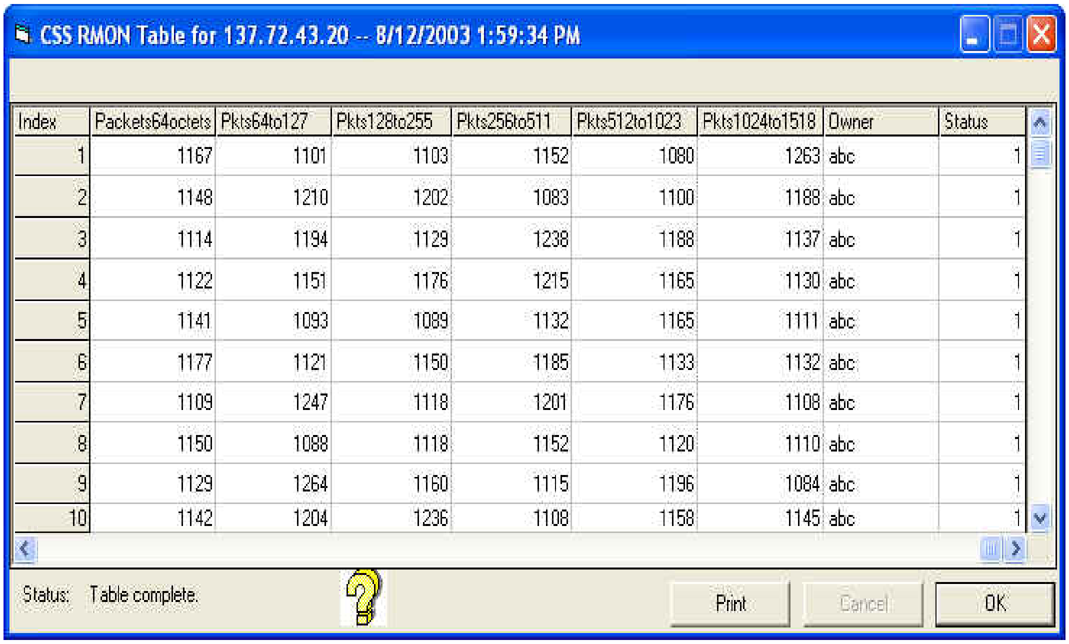

CSS RMON Table for 137.72.43.20 -- 8/12/2003 1:59:34 PM

Index	Packets64octets	Pkts64to127	Pkts128to255	Pkts256to511	Pkts512to1023	Pkts1024to1518	Owner	Status
1	1167	1101	1103	1152	1080	1263	abc	1
2	1148	1210	1202	1083	1100	1188	abc	1
3	1114	1194	1129	1238	1188	1137	abc	1
4	1122	1151	1176	1215	1165	1130	abc	1
5	1141	1093	1089	1132	1165	1111	abc	1
6	1177	1121	1150	1185	1133	1132	abc	1
7	1109	1247	1118	1201	1176	1108	abc	1
8	1150	1088	1118	1152	1120	1110	abc	1
9	1129	1264	1160	1115	1196	1084	abc	1
10	1142	1204	1236	1108	1158	1145	abc	1

Status: Table complete.

CLEVER eRoute CSS RMON Reporting

The Number of Octets field shows the total number of octets received on the network, and can be used to estimate Ethernet utilization. The Packet information (broadcast, multicast, undersized, oversized), and the brackets of packet sizes (up to 64 bytes, 65-127 bytes, 128-255 bytes, 256-511 bytes, 512-1023 bytes, 1024-1518 bytes) can provide packet volume and distribution information.

ICMP MIB Report

The Internet Control Message Protocol (ICMP) objects are the input and output error and control message statistics for the IP layer. The report groups packets by IN and OUT categories, and provides a graphical report format.

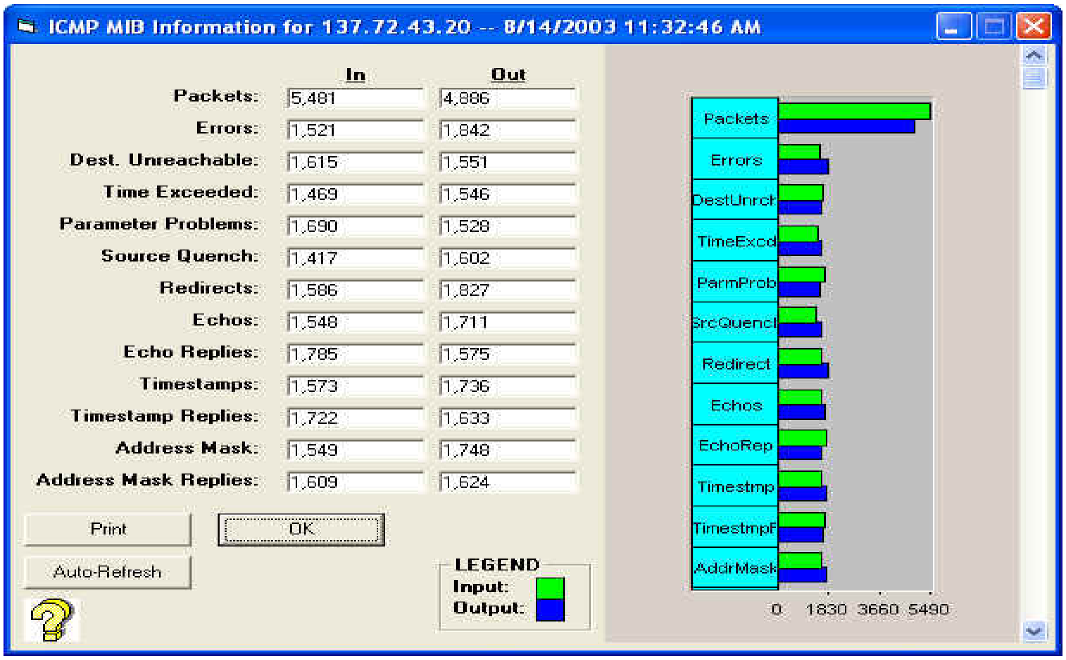

ICMP MIB Information for 137.72.43.20 -- 8/14/2003 11:32:46 AM

	In	Out
Packets:	5,481	4,886
Errors:	1,521	1,842
Dest. Unreachable:	1,615	1,551
Time Exceeded:	1,469	1,546
Parameter Problems:	1,690	1,528
Source Quench:	1,417	1,602
Redirects:	1,586	1,827
Echos:	1,548	1,711
Echo Replies:	1,785	1,575
Timestamps:	1,573	1,736
Timestamp Replies:	1,722	1,633
Address Mask:	1,549	1,748
Address Mask Replies:	1,609	1,624

LEGEND
Input: 🟩
Output: 🟦

CLEVER eRoute ICMP MIB Reporting

The Packets data field provides the total number of ICMP packets received and sent. The Errors field shows the number of receiving/sending ICMP errors (e.g., bad ICMP checksum, bad length, inability of IP to route, etc) used for determining causes of errors. The number of time-outs and time-stamps gives an indication as to the number of ICMP request and reply messages received and sent for coordinating the transmission sequence.

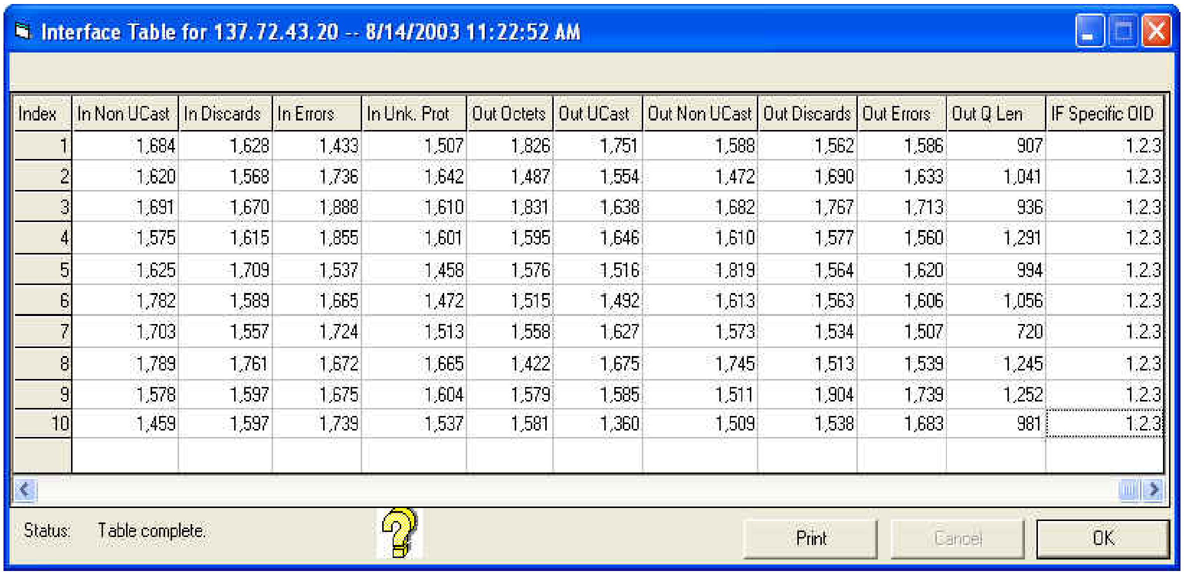

Index	In Non UCast	In Discards	In Errors	In Unk. Prot	Out Octets	Out UCast	Out Non UCast	Out Discards	Out Errors	Out Q Len	IF Specific OID
1	1,684	1,628	1,433	1,507	1,826	1,751	1,588	1,562	1,586	907	1.2.3
2	1,620	1,568	1,736	1,642	1,487	1,554	1,472	1,690	1,633	1,041	1.2.3
3	1,691	1,670	1,888	1,610	1,831	1,638	1,682	1,767	1,713	936	1.2.3
4	1,575	1,615	1,855	1,601	1,595	1,646	1,610	1,577	1,560	1,291	1.2.3
5	1,625	1,709	1,537	1,458	1,576	1,516	1,819	1,564	1,620	994	1.2.3
6	1,782	1,589	1,665	1,472	1,515	1,492	1,613	1,563	1,606	1,056	1.2.3
7	1,703	1,557	1,724	1,513	1,558	1,627	1,573	1,534	1,507	720	1.2.3
8	1,789	1,761	1,672	1,665	1,422	1,675	1,745	1,513	1,539	1,245	1.2.3
9	1,578	1,597	1,675	1,604	1,579	1,585	1,511	1,904	1,739	1,252	1.2.3
10	1,459	1,597	1,739	1,537	1,581	1,360	1,509	1,538	1,683	981	1.2.3

Status: Table complete.

CLEVER eRoute Interface Reporting

Interface Report

The Interface report contains information regarding the entity's interfaces. Each interface is treated based on the pertinent sub-network.

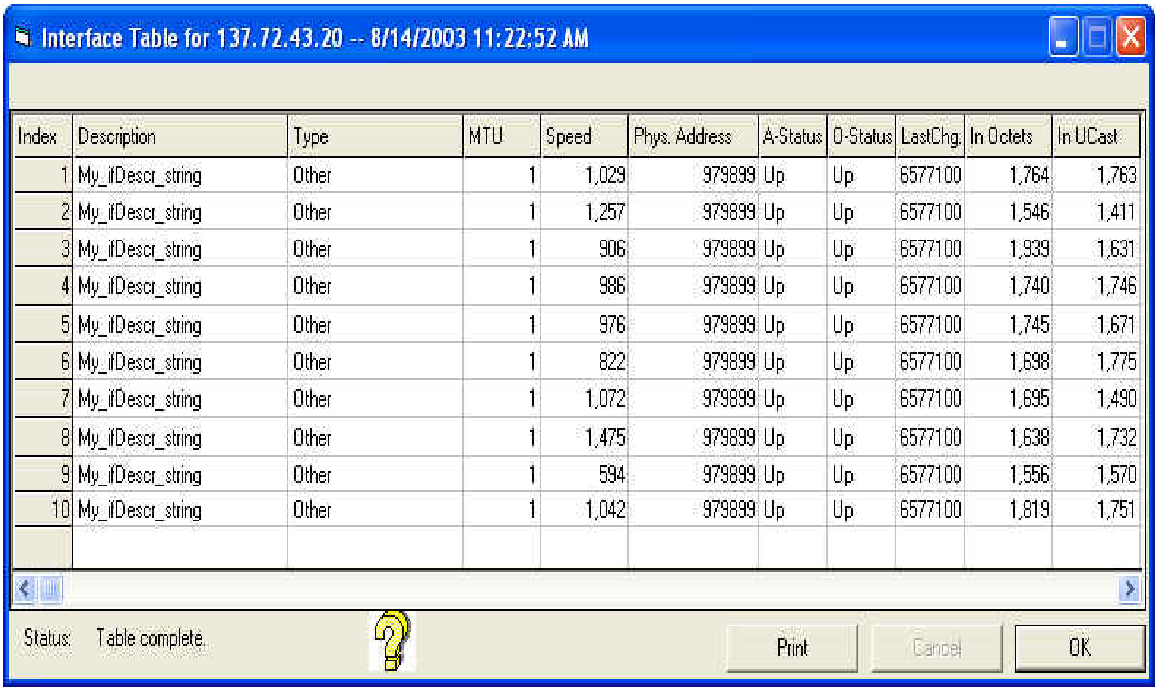

Index	Description	Type	MTU	Speed	Phys. Address	A-Status	O-Status	LastChg.	In Octets	In UCast
1	My_ifDescr_string	Other	1	1,029	979899	Up	Up	6577100	1,764	1,763
2	My_ifDescr_string	Other	1	1,257	979899	Up	Up	6577100	1,546	1,411
3	My_ifDescr_string	Other	1	906	979899	Up	Up	6577100	1,939	1,631
4	My_ifDescr_string	Other	1	986	979899	Up	Up	6577100	1,740	1,746
5	My_ifDescr_string	Other	1	976	979899	Up	Up	6577100	1,745	1,671
6	My_ifDescr_string	Other	1	822	979899	Up	Up	6577100	1,698	1,775
7	My_ifDescr_string	Other	1	1,072	979899	Up	Up	6577100	1,695	1,490
8	My_ifDescr_string	Other	1	1,475	979899	Up	Up	6577100	1,638	1,732
9	My_ifDescr_string	Other	1	594	979899	Up	Up	6577100	1,556	1,570
10	My_ifDescr_string	Other	1	1,042	979899	Up	Up	6577100	1,819	1,751

Status: Table complete.

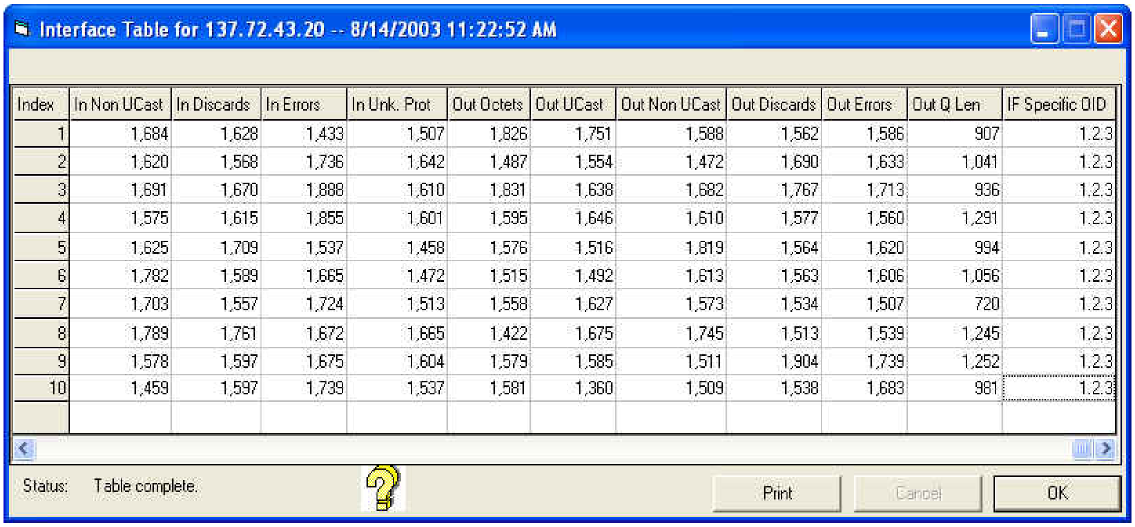

Interface Table for 137.72.43.20 -- 8/14/2003 11:22:52 AM

Index	In Non UCast	In Discards	In Errors	In Unk. Prot	Out Octets	Out UCast	Out Non UCast	Out Discards	Out Errors	Out Q Len	IF Specific OID
1	1,684	1,628	1,433	1,507	1,826	1,751	1,588	1,562	1,586	907	1.2.3
2	1,620	1,568	1,736	1,642	1,487	1,554	1,472	1,690	1,633	1,041	1.2.3
3	1,691	1,670	1,888	1,610	1,831	1,638	1,682	1,767	1,713	936	1.2.3
4	1,575	1,615	1,855	1,601	1,595	1,646	1,610	1,577	1,560	1,291	1.2.3
5	1,625	1,709	1,537	1,458	1,576	1,516	1,819	1,564	1,620	994	1.2.3
6	1,782	1,589	1,665	1,472	1,515	1,492	1,613	1,563	1,606	1,056	1.2.3
7	1,703	1,557	1,724	1,513	1,558	1,627	1,573	1,534	1,507	720	1.2.3
8	1,789	1,761	1,672	1,665	1,422	1,675	1,745	1,513	1,539	1,245	1.2.3
9	1,578	1,597	1,675	1,604	1,579	1,585	1,511	1,904	1,739	1,252	1.2.3
10	1,459	1,597	1,739	1,537	1,581	1,360	1,509	1,538	1,683	981	1.2.3

Status: Table complete.

CLEVER eRoute Interface Reporting

The Type of interface is based on the protocol immediately below the network layer in the protocol stack. The Maximum Transmission Unit (MTU) limits the size of the largest datagram that can be sent/received on the interface. The Speed is an estimate of the interface's current bandwidth in bits per second. This table also provides a set of causes for IN and OUT packets of the interface (e.g., Discards, Errors, Unsupported, Non-Unicast, etc.). The OUT Q Len field indicates the length of the output packet queue, which is very helpful in determining if the output queue size is properly chosen. The total numbers of IN and OUT packets in octets are a good indication of the performance of the interface.

IP MIB Report

The data fields of the IP MIB report, shown below, provide several useful statistics.

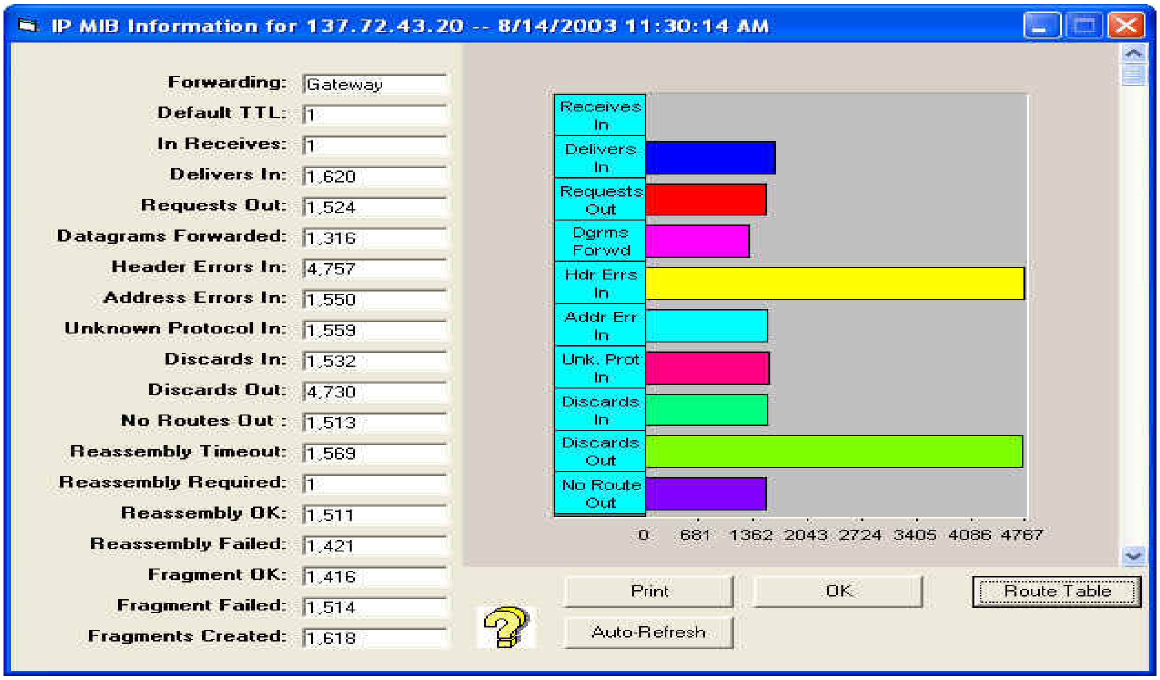

IP MIB Information for 137.72.43.20 -- 8/14/2003 11:30:14 AM

Forwarding:	Gateway
Default TTL:	1
In Receives:	1
Delivers In:	1,620
Requests Out:	1,524
Datagrams Forwarded:	1,316
Header Errors In:	4,757
Address Errors In:	1,550
Unknown Protocol In:	1,559
Discards In:	1,532
Discards Out:	4,730
No Routes Out :	1,513
Reassembly Timeout:	1,569
Reassembly Required:	1
Reassembly OK:	1,511
Reassembly Failed:	1,421
Fragment OK:	1,416
Fragment Failed:	1,514
Fragments Created:	1,618

CLEVER eRoute IP MIB Reporting

The Time-To-Live (TTL) field in the IP header of datagram shows the TTL value, ranging from 1 to 255. The Forwarding field indicates whether this entity is acting as an IP gateway. The value 1 indicates to forward the received packet and a 2 means do not forward. The Reassembly information, such as the number of requests, time-outs, failed and successful counts, would provide the effectiveness of the datagram reassembly process. The Fragment field gives an indication of IP datagram fragmentation information.

IP Route Report

The IP Route Report contains an entry for each route presently known to this segment or node.

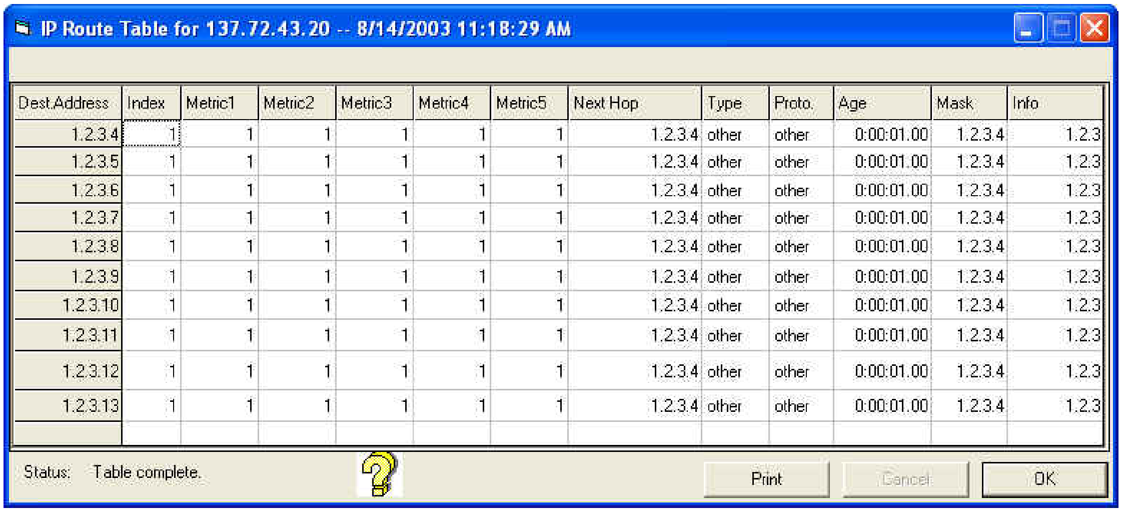

IP Route Table for 137.72.43.20 -- 8/14/2003 11:18:29 AM

Dest.Address	Index	Metric1	Metric2	Metric3	Metric4	Metric5	Next Hop	Type	Proto.	Age	Mask	Info
1.2.3.4	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.5	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.6	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.7	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.8	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.9	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.10	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.11	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.12	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3
1.2.3.13	1	1	1	1	1	1	1.2.3.4	other	other	0:00:01.00	1.2.3.4	1.2.3

Status: Table complete.

CLEVER eRoute IP Route Reporting

There are five metrics provided by the IP Route report. The first metric is the primary routing parameter that identifies the routing-protocol specified in the router IP-Route-Proto value. If it is not used, its value should be set to –1. The second metric contains an alternate routing value that is determined by the IP-Route-Proto value. If it is not used, it should be set to –1. The third, fourth, and fifth metrics provide alternate routing information based on the value in IP-Route-Proto field.

The remaining fields are Next Hop, Type, Proto, Age, and Mask. The Next Hop field contains the IP address of the next hop of this route that the received packet can be sent to. The Type field provides the routing action. This field can have one of four possible values:

1 - Indicates not applicable
2 - Shows an invalidated route
3 - Means to route only to the directly attached sub-network
4 - Means to route to a non-local node/sub-network.

The Proto field is a routing mechanism that the route would be able to learn. There are four possible values:

1 - Means no action
2 - Means a non-protocol or manually configured entry
3 - Indicates a network management protocol
4 - Means to learn from ICMP data field.

The Age field contains the number of seconds since this route was last updated or otherwise determined to be correct. The Mask field provides the validation of the destination address via the logically ANDed comparison.

CONCLUDING REMARKS

TCP/IP route performance analysis can yield extremely beneficial data to better tune an open enterprise network environment, resulting in improved performance and reliability, as well as significant cost savings. It can be used with other TCP/IP performance analysis tools to effectively manage the performance of the TCP/IP network, particularly when it contains some TN3270-based networks. Given the nature of dynamic routing and its impact on problem diagnosis, performance tuning, and capacity planning, TCP/IP networks are understandably challenging to manage. Without the means to collect and analyze data in a systematic framework that shows the behavior of routes and segments over a period of time, symptoms may be easily misinterpreted. Thorough route performance analysis allows symptoms to be interpreted more accurately the first time, preventing the misuse of human resources by undertaking tasks that do not address the real problems. In addition, it can also prevent the unnecessary acquisition of additional hardware resources.

Through the description of data collection architectures, data elements, and examples, we hope we have demonstrated that superior data collection and analysis of TCP/IP route performance can indeed provide valuable input to address not only the short-term problems, but also to allow preparation to meet enterprise network long-term growth objectives. AES CLEVER *e*Route was created to uniquely answer to these needs.