




(72) DRYSDALE, R. SCOTT, US
(72) ENNIS, JAMES D., JR., US
(72) HASSELKUS, JOHN E., US
(72) NISBET, THOMAS R., US
(72) TROUTMAN, ROBERT C., US
(72) REHBEHN, KEN, US
(72) WHEELER, DAVID, US
(71) VISUAL NETWORKS TECHNOLOGIES, INC., US

(51) Int.Cl.⁶ H04J 3/14

(30) 1997/11/07 (60/064,620) US


(54) **PROCEDE ET APPAREIL PERMETTANT D'EFFECTUER DES
ANALYSES DE NIVEAU DE SERVICE DE PARAMETRES DE
LA TENUE DU RESEAU DE COMMUNICATION AUX
DONNEES**

(54) **METHOD AND APPARATUS FOR PERFORMING SERVICE
LEVEL ANALYSIS OF COMMUNICATIONS NETWORK
PERFORMANCE METRICS**



(57) L'invention concerne un système de transmission de données qui comprend des sondes (sondes A, B) reliées entre un site (sites A, B) d'utilisateur final et un réseau (12) de commutation de données. Chaque sonde est reliée au réseau (12) de commutation via un canal (20, 21, 26 ou 27) d'accès dans laquelle les circuits de transmission établissent des trajets entre les sites à travers le canal (20, 21, 26 ou 27) d'accès et le réseau (12) de commutation. Les sondes capturent et retransmettent le déplacement de données entre les sites via des circuits de transmission et peuvent par conséquent insérer des messages d'analyses de niveau de

(57) A data transmission system includes probes (probe A, probe B) connected between end user sites (site A, site B) and a data switching network (12). Each probe is connected to the switching network (12) via an access channel (20, 21, 26 or 27) wherein transmission circuits establish paths between the sites through the access channel (20, 21, 26 or 27) and switching network (12). The probes capture and retransmit data traveling between the sites over respective transmission circuits, and can thereby insert service level analysis (SLA) messages into data traffic in order to actively communicate network performance information to other



(21) (A1) **2,307,766**
(86) 1998/11/06
(87) 1999/05/20

service (SLA) dans le trafic de données de manière à communiquer activement à d'autres sondes des informations sur la tenue du réseau aux données. Les sondes recueillent périodiquement, pour chaque circuit de transmission, des mesures concernant un ou plusieurs paramètres de tenue de réseau aux données, comprenant le temps de transmission aller-retour (RTE), le rapport de distribution de données (DDR) et la disponibilité du réseau. Pendant chaque cycle de mesure SLA, une séquence de messages SLA est échangée sur chaque circuit de transmission, les messages contenant des données utilisées pour déterminer le RTD et DDR.


probes. For each transmission circuit, the probes periodically collect measurements related to one or more network performance metrics, including round-trip delay (RTD), data delivery ratio (DDR) and network availability. During each SLA measurement cycle, a sequence of SLA messages is exchanged over each transmission circuit, which messages contain data used to determine RTD and DDR.

**PCT**WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04J 3/14	A1	(11) International Publication Number: WO 99/25085 (43) International Publication Date: 20 May 1999 (20.05.99)
(21) International Application Number: PCT/US98/23891 (22) International Filing Date: 6 November 1998 (06.11.98) (30) Priority Data: 60/064,620 7 November 1997 (07.11.97) US (71) Applicant (for all designated States except US): VISUAL NETWORKS, INC. [US/US]; Suite 220, 2092 Gaither Road, Rockville, MD 20850 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): DRYSDALE, R., Scott [US/US]; 3 McMurray, Frederick, MD 21701 (US). EN-NIS, James, D., Jr. [US/US]; 11204 Jeffro Court, Ijamsville, MD 21754 (US). HASSELKUS, John, E. [US/US]; 13215 Meander Cove Drive, Germantown, MD 20874 (US). NIS-BET, Thomas, R. [US/US]; 10215 Tarpley Court, Ellicott City, MD 21042 (US). TROUTMAN, Robert, C. [US/US]; 11 Melmark Court, Gaithersburg, MD 20878 (US). RE-HBEHN, Ken [US/US]; 4802 Strathmore Avenue, Garret Park, MD 20896 (US). WHEELER, David [US/US]; 19009 Laurel Grove Terrace, Germantown, MD 20874 (US).		(74) Agents: FINNAN, Patrick, J. et al.; Epstein, Edell & Retzer, Suite 400, 1901 Research Boulevard, Rockville, MD 20850-3164 (US). (81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: METHOD AND APPARATUS FOR PERFORMING SERVICE LEVEL ANALYSIS OF COMMUNICATIONS NETWORK PERFORMANCE METRICS



(57) Abstract

A data transmission system includes probes (probe A, probe B) connected between end user sites (site A, site B) and a data switching network (12). Each probe is connected to the switching network (12) via an access channel (20, 21, 26 or 27) wherein transmission circuits establish paths between the sites through the access channel (20, 21, 26 or 27) and switching network (12). The probes capture and retransmit data traveling between the sites over respective transmission circuits, and can thereby insert service level analysis (SLA) messages into data traffic in order to actively communicate network performance information to other probes. For each transmission circuit, the probes periodically collect measurements related to one or more network performance metrics, including round-trip delay (RTD), data delivery ratio (DDR) and network availability. During each SLA measurement cycle, a sequence of SLA messages is exchanged over each transmission circuit, which messages contain data used to determine RTD and DDR.

METHOD AND APPARATUS FOR PERFORMING SERVICE
LEVEL ANALYSIS OF COMMUNICATIONS
NETWORK PERFORMANCE METRICS

CROSS REFERENCE TO RELATED APPLICATION

This application claims priority from U.S. Provisional Patent Application Serial No. 60/064,620, entitled "Method And Apparatus For Measurement of Network Availability, Data Delivery Ratio and Round Trip Delay In Communications Networks," filed November 7, 1997. The disclosure of that provisional patent application is incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to a method and apparatus for monitoring data transmission through a communications network while the communications network is in service. More particularly, the present invention relates to a network monitoring system having endpoint probes that transmit and receive message data over a packetized switching network to exchange information allowing the probes to measure network performance metrics, such as network availability, data delivery ratio and round trip delay in the communications network.

Description of the Related Art

Both from an end user and service provider standpoint, there is an increasing need to accurately measure operational performance of data communications networks. Communications networks, especially packetized data networks, are currently utilized in various applications for transmission and reception of data between parties at different locations. A typical data transmission system includes a plurality of end user sites and a data packet switching network, which resides between the sites to facilitate communications. Each site is connected to the

1 switching network via an access channel (i.e., a channel connecting a site to a
2 communications system), wherein transmission circuits, preferably virtual circuits,
3 establish paths between the sites through the access channel and the switching
4 network.

5 Packetized data networks typically format data into packets for transmission
6 from one site to another. In particular, the data is partitioned into separate packets
7 at a transmission site, wherein the packets usually include headers containing
8 information relating to packet data and routing. The packets are transmitted to a
9 destination site in accordance with any of several conventional data transmission
10 protocols known in the art (e.g., Asynchronous Transfer Mode (ATM), Frame Relay,
11 High Level Data Link Control (HDLC), X.25, IP tunneling, etc.), by which the
12 transmitted data is restored from the packets received at the destination site.

13 Packetized data communications are especially appealing for common carrier
14 or time-shared switching systems, since a packet transmission path or circuit is
15 unavailable only during the time when a packet utilizes the circuit for transmission to
16 the destination site, thereby permitting other users to utilize that same circuit when
17 the circuit becomes available (i.e., during intervening periods between packet
18 transmissions). The access channel and each individual transmission circuit
19 typically have a maximum data carrying capacity or bandwidth that is shared among
20 the various users of the network. The access channel utilization is typically
21 measured as an aggregate of the individual circuit utilizations and has a fixed
22 bandwidth, while the individual circuits may be utilized by several users wherein
23 each user may utilize an allocated portion of the circuit.

24 Typically, when a party needs to send and receive data over distances, the
25 party (end user) enters into a service contract with a service provider to provide
26 access to a data communications network. Depending on an individual end user's
27 needs, the service contract may include provisions that guarantee certain minimum
28 performance requirements that the service provider must meet. For example, if the
29 end user expects to send and receive a certain amount of data on a regular basis,
30 the end user may want the service provider to guarantee that a certain minimum
31 bandwidth will be available to the end user at all times. Certain end user

1 applications are sensitive to transmission delays and/or the loss of data within the
2 network (i.e., failure to successfully deliver data packet(s) to their destination).
3 Specifically, while loss of data packets can generally be detected by end users (via
4 information provided in the data transmission protocol), and lost packets can be
5 retransmitted, certain applications cannot function when the percentage of lost data
6 exceeds a given level. Thus, the end user may want the service provider to
7 guarantee that the average or minimum ratio of data units delivered by the network
8 to data units offered to the network at the far-end is above a certain percentage
9 and/or that the average or maximum transmission delays will not exceed a certain
10 duration.

11 From a service provider's perspective, it would be competitively
12 advantageous to be able to demonstrate to potential and existing end users that the
13 service provider is capable of meeting and does meet such network performance
14 metrics. Thus, the capability to provide analysis of network system performance at
15 the service level, i.e., service level analysis (SLA), particularly in the context of
16 network systems that share bandwidth between sites, would be advantageous from
17 both an end user and service provider standpoint.

18 Various systems have been proposed which provide some measure of
19 network system performance. Specifically, a number of techniques for measuring
20 round trip delay (RTD) of data transmitted between two sites is known. For
21 example, U.S. Patent No. 5,521,907 to Ennis, Jr. et al., the disclosure of which is
22 incorporated herein by reference in its entirety, discloses a system for passively
23 measuring the round trip delay of data messages sent between two sites. More
24 specifically, a console triggers probes at two sites to store data packets being sent
25 between the two sites. The probes generate unique packet signatures based on the
26 data in the packets, and time stamp the signatures. By matching signatures from
27 the two probes and comparing the corresponding timestamp values, the console can
28 determine the round trip delay between the sites. This technique requires the
29 storage, transmission and processing of a significant amount of data, particularly if
30 implemented to periodically monitor all virtual circuits existing between a set of sites.
31 That is, the passive probes cannot individually determine round trip delay, and each

1 probe must store and transmit a substantial amount of data to the console which is
2 required to correlate signature and timestamp data from different sites.

3 U.S. Patent No. 5,450,394 to Gruber et al., the disclosure of which is
4 incorporated herein by reference in its entirety, discloses a technique for determining
5 round trip delay in which measurement cells containing timestamp information are
6 sent between two nodes. A first node transmits a measurement cell with a first time
7 stamp to a second node, and the second node replies with a measurement cell
8 containing additional time stamp information which can be used by the first node to
9 determine the round trip delay. Because the technique relies, in part, on timestamps
10 already present in PM OAM (performance management operations, administration
11 and maintenance) ATM cells, the technique is specific to the ATM protocol and
12 cannot readily be adapted to other data protocols or be expanded to monitor other
13 service level performance metrics. Further, the technique does not allow both nodes
14 to measure the round trip delay of the same sequence of cells (i.e., either only one
15 of the two nodes measures round trip delay or the two nodes measure delays of
16 different transmitted cell sequences).

17 Further, while it is possible for individual switches in existing network systems
18 to indicate how many packets of data have been dropped by the switch, there are no
19 known systems capable of measuring a rate of successful (or unsuccessful) data
20 delivery on a service level, e.g., over a particular virtual circuit or to a particular end
21 user.

22 The problem of providing service level analysis of network performance is
23 complicated by the fact that many switching networks comprise interworked systems
24 using plural, different data transmission protocols (e.g., an ATM switching network
25 interworked with a Frame Relay switching network), thereby forming a so-called
26 "interworked" network. Such interworked networks are becoming more common,
27 and present an additional challenge to designing a service level analysis tool that
28 employs a standard message structure and messaging protocol useful for
29 communicating between any two sites. Existing systems relying on inter-site or
30 inter-probe messages to assess system performance are generally incapable of
31 operating across interworked networks.

6 **SUMMARY OF THE INVENTION**

28 A further object of the present invention is to provide a tool capable of
29 performing service level analysis with any conventional data transmission protocol.

1 Yet a further object of the present invention is to transmit service level
2 analysis (SLA) messages over transmission circuits connecting user endpoints that
3 can be encapsulated in a single, standard data unit of any conventional data
4 transmission protocol.

5 Still a further object of the present invention is to provide service level
6 analysis of network performance metrics by transmitting SLA messages over circuits
7 connecting endpoint sites without requiring synchronization between local site/probe
8 clocks.

9 Another object of the present invention is to monitor network availability of
10 transmission circuits connecting endpoint sites.

11 The aforesaid objects are achieved individually and in combination, and it is
12 not intended that the present invention be construed as requiring two or more of the
13 objects to be combined unless expressly required by the claims attached hereto.

14 The system of the present invention provides a service level analysis (SLA)
15 capability for communications networks, especially any packetized, circuit-oriented
16 data networks, wherein packets (or units) of data are transmitted to a destination site
17 in accordance with one or a combination of data transmission protocols, including,
18 but not limited to: Frame Relay; ATM and X.25. According to the present invention,
19 a data transmission system includes endpoint probes connected between
20 corresponding endpoint user sites and a data switching network. Each probe is
21 connected to the switching network via an access channel wherein transmission
22 circuits, preferably virtual circuits (e.g., a permanent virtual circuit (PVC), a switched
23 virtual circuit (SVC) or a tunnel through an IP network), establish paths between the
24 sites/probes through the access channel and switching network. A virtual circuit is
25 basically a path established in a packet switching network to transfer data to a
26 specific destination or site.

27 Each probe captures and retransmits data traveling between its site and other
28 sites over respective transmission circuits (such as virtual circuits (VCs)) of the
29 switching network, and can thereby insert service level analysis (SLA) messages
30 into the data traffic in order to actively generate and communicate network
31 performance information to other probes.

1 Each probe periodically stores data collected for each VC associated with the
2 site, which data is used to monitor certain network performance metrics on a per-VC
3 basis. Specifically, each probe performs service level analysis, including monitoring
4 of one or more of the following system performance metrics: round-trip delay, data
5 delivery ratio and availability. These performance metrics are measured for each
6 virtual circuit connecting endpoint sites through the network, although these metrics
7 are affected by the access line and access channel.

8 During each SLA measurement cycle, whose duration can be set at a desired
9 predetermined value, such as every fifteen minutes, a probe exchanges a sequence
10 of SLA messages over each VC connecting the probe to other probes. The
11 sequence of messages contains timestamp and counter data that allows the probe
12 to determine the round trip delay of the SLA messages as well as the number of
13 data units offered to the network (bound for the probe's site) and the number of data
14 units actually delivered to the site for each VC connected to the site. The probes are
15 not required to be time synchronized in order to exchange these SLA messages.

16 In accordance with another aspect of the present invention, the SLA message
17 contents and protocol are designed to allow the SLA messages to be encapsulated
18 in a single, standard data unit of any conventional data transmission protocol, e.g.,
19 an ATM cell or a Frame Relay frame. Consequently, the messaging system can be
20 used with any data transmission protocol and in interworked networks without
21 modification of the message data payload.

22 The SLA measurements collected by the probes can be transmitted to a
23 console in communication with the probes for processing, display, and archiving.
24 The SLA measurements provided by the probes of the present invention allows end
25 users and service providers to know whether the network performance is meeting
26 requirements set forth in a customer service agreement.

27 The above and still further objects, features and advantages of the present
28 invention will become apparent upon consideration of the following detailed
29 description of specific embodiments thereof, particularly when taken in conjunction
30 with the accompanying drawings wherein like reference numerals in the various
31 figures are utilized to designate like components.

1

2

BRIEF DESCRIPTION OF THE DRAWINGS

3

4

5

Fig. 1 is a functional block diagram of a data transmission system having service level analysis probes disposed between endpoint sites and a switching network in accordance with an exemplary embodiment of the present invention.

6

7

Fig. 2 is a diagram illustrating the time sequence and contents of a sequence of inter-probe messages sent between two probes of the present invention.

8

9

10

11

Figs. 3A and 3B illustrate a functional flow chart corresponding to the diagram of Fig. 2, indicating operations performed by the probes of the present invention in order to capture and store measurements supporting service level analysis of network performance metrics.

12

13

14

Fig. 4 is a diagram illustrating an inter-probe message structure used in accordance with an exemplary embodiment of the present invention to transmit measurements supporting service level analysis of network performance metrics.

15

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

A system for monitoring performance and providing service level analysis (SLA) capability for data communications networks is illustrated in Fig. 1. Specifically, an exemplary data transmission system 10 includes two sites (A and B) and a packet switching network 12 to facilitate communications between the sites. Site A is connected to network 12 via a probe A, while site B is connected to network 12 via another probe B. Site A is connected to probe A by communication lines 20 and 21, probe A is connected to network 12 by communication lines 22 and 23, network 12 is connected to probe B by communication lines 24 and 25, and probe B is connected to site B by communication lines 26 and 27. The data transmission system 10 typically includes conventional telecommunications line types, such as T3, OC-3, North American T1 (1.544 Mbits/second), CCITT (variable rate), 56K or 64K North American Digital Dataphone Service (DDS), and a variety of data communications connections, such as V.35, RS-449, EIA 530, X.21 and RS-232. Sites A and B are each capable of transmitting and receiving data packets in various protocols utilized by communication lines 20, 21, 26 and 27, such as Asynchronous

1 Transfer Mode (ATM), Frame Relay, High Level Data Link Control (HDLC) and X.25.
2 Each line 20-27 represents a respective transmission direction as indicated by the
3 arrows. For example, the arrows on communication lines 20 and 27 represent
4 transmissions from sites A and B to probes A and B, respectively, while the arrows
5 on communication lines 21 and 26 represent transmissions to sites A and B from
6 probes A and B, respectively. Similarly, the arrows on communication lines 22 and
7 25 represent transmissions from probes A and B to switching network 12,
8 respectively, while the arrows on communication lines 23 and 24 represent
9 transmissions to probes A and B from switching network 12, respectively.

10 Generally, site A and site B utilize switching network 12 to communicate with
11 each other, wherein each site is connected to switching network 12 via an access
12 channel having transmission circuits, preferably virtual circuits, that establish paths
13 between the sites through the access channel and switching network. The access
14 channel refers to the lines utilized by each site to communicate with the switching
15 network (i.e., communication lines 20-27), while a virtual circuit is basically a path
16 established through a packetized data switching network that transfers data to a
17 specific endpoint or site.

18 As used herein, the term "packet" (e.g., as used in "packetized switching
19 network") does not imply any particular transmission protocol and can refer to units
20 or segments of data in a system using, for example, any one or combination of the
21 above-listed data transmission protocols (or other protocols). However, since the
22 term "packet" is often associated with only certain data transmission protocols, to
23 avoid any suggestion that the system of the present invention is limited to any
24 particular data transmission protocols, the term "protocol data unit" (PDU) will be
25 used herein to refer the unit of data being transmitted from sites A and B through
26 switching network 12. Thus, for example, a PDU can be carried on a frame in the
27 Frame Relay protocol, a related set of cells in the ATM protocol, or a packet in the IP
28 protocol.

29 As shown in Fig. 1, probes A and B are respectively disposed between
30 switching network 12 and sites A and B. Probes A and B can be located at sites A
31 and B or at any point between switching network 12 and sites A and B. Probes A

1 and B allow PDUs being sent between sites A and B via switching network 12 to
2 pass through the probe and also insert inter-probe message PDUs into the data
3 traffic. As used herein, the term "to the switching network" refers to the direction of
4 data traveling to the switching network (e.g., data traveling on communication lines
5 20, 22, 25 and 27), while the term "from the switching network" refers to the direction
6 of data traveling from the switching network (e.g., data traveling on communication
7 lines 21, 23, 24 and 26). The terms "arriving", "to switching network", "departing",
8 and "from switching network" are all relative and are employed to imply transmission
9 direction.

10 For illustrative purposes, only two sites (A and B) are shown in Fig. 1.
11 However, it will be understood that the data communication system can include
12 numerous sites, wherein each site is generally connected to multiple other sites over
13 corresponding transmission circuits, such as virtual circuits (VCs).

14 In accordance with an exemplary embodiment of the present invention,
15 probes A and B actively perform service level analysis (SLA) by exchanging a
16 sequence of inter-probe SLA messages that are inserted into the data traffic
17 traveling over switching network 12 between sites, and that contain information
18 allowing the probes to monitor one or more of the following system performance
19 metrics on a per-virtual-circuit basis for each virtual circuit connecting each site to
20 other sites: round trip delay, data delivery ratio and network availability. Round-trip
21 delay (RTD) can be defined as the duration of time required for a PDU to go from a
22 first end of a virtual circuit (VC) to the second end plus the time required for return
23 PDU to go from the second end of the VC to the first end. This round trip delay does
24 not include the far-end turn around time (i.e., the duration between the receive time
25 of the first PDU at the second end and the transmit time of the return PDU from the
26 second end). Optionally, components of the delay that are not attributable to the
27 switching network, such as those related to the access line rate (i.e. serialization
28 delays), can be excluded from the computed RTD. Data Delivery Ratio (DDR) is the
29 ratio of the number of PDUs delivered on a transmission circuit by the network, to
30 the number of PDUs offered to the network on that transmission circuit. The DDR
31 can be computed from the aggregate of the PDUs offered and delivered in both

1 directions of the circuit or as separate, one-way ratios, as described in greater detail
2 below. Availability is defined as the percent of the time that the network (or virtual
3 circuit) is capable of accepting and delivering user PDUs within the total time of a
4 measurement period. Unlike RTD and DDR, Availability can be determined
5 passively, without inserting messages into the data traffic, as explained below.

6 The SLA measurements performed by the probes to monitor these network
7 performance metrics apply specifically to the virtual circuits through the switching
8 network, although they are affected by the access line and access channel. The
9 measurements that are accumulated on the probes of the present invention support
10 service level analysis. Each measurement is preferably accumulated on a
11 continuous basis with a minimum of one historical sample per fixed historical time
12 interval, e.g., once every fifteen minutes. As will become evident, the
13 measurements are collected in a manner that scales to very large networks without
14 performance degradation.

15 In order to acquire the service level analysis (SLA) measurements required to
16 determine round trip delay (RTD) and data delivery ratio (DDR), endpoint probes A
17 and B insert data into the user data stream being transmitted through the network
18 over the managed transmission circuit connecting sites A and B, which circuit can be
19 a virtual circuit (VC). Probes A and B also insert similar data over the other VCs
20 connecting sites A and B to other sites.

21 More specifically, each site's probe separately maintains a periodic historical
22 measurement cycle (which can be, for example, 15 minutes long) which is not
23 required to be synchronized with the periodic measurement cycles of other probes.
24 During each periodic cycle, each probe acquires one set of SLA measurements for
25 each VC. At an activation time within the cycle (which is an arbitrary, predetermined
26 time after the beginning of the measurement cycle, e.g., the mid-point of each
27 measurement cycle), each probe initiates an inter-probe message sequence on
28 each VC that has not already had an inter-probe message sequence initiated by a
29 far-end probe. That is, because the periodic cycles of the endpoint probes in the
30 communications system are not synchronized, when the activation time within a
31 cycle occurs for a particular probe (probe A), certain far-end probes connected to

1 probe A over certain VCs may have already reached the activation times within their
2 measurement cycles, and therefore would have initiated a message sequence for
3 these VCs. Other far-end probes connected to probe A over other VCs may not
4 have reached the activation times within their cycles at the activation time of probe
5 A's cycle, and therefore would not have initiated a message sequence for these
6 other VCs. For illustrative purposes, it is assumed in the exemplary embodiment
7 that probe A reaches the activation time within its periodic measurement cycle prior
8 to probe B reaching the activation time within its periodic measurement cycle, and
9 therefore probe A initiates the inter-probe messaging sequence with probe B over
10 the VC connecting probes A and B. However, it will be understood that the
11 hereafter-described message sequence is initiated for every VC by the one of the
12 VC's two endpoint probes that first reaches the activation time within its periodic
13 historical measurement cycle.

14 Each probe acquires certain information for each VC during each
15 measurement cycle in order to support SLA. Specifically, in order for each probe to
16 be able to determine the round trip delay (RTD) for a VC, the probe must know the
17 transmit and receive times of an SLA message PDU transmitted in one direction
18 over the VC and the transmit and receive times of a reply SLA message PDU
19 transmitted in the opposite direction over the VC.

20 In order to determine the data delivery ratio (DDR), for each monitored VC,
21 each probe maintains a counter that accumulates the count of the total number of
22 PDUs delivered ("delivered PDUs") by the network over the VC (e.g., probe A
23 maintains a count of the number of PDUs delivered by the network from probe B and
24 vice versa). Additionally, each probe maintains a counter, per VC, that accumulates
25 the total number of PDUs offered by the probe.

26 To support SLA, probe A acquires and stores in each measurement cycle a
27 count of the PDUs bound for site A that were offered to the network by probe B (not
28 initially known by probe A, but known by probe B), and a count of the PDUs
29 originating from site B that were delivered by the network to probe A (known by
30 probe A from its counters). Similarly, probe B acquires and stores in each
31 measurement cycle a count of the PDUs bound for site B that were offered to the

1 network by probe A (not initially known by probe B, but known by probe A), and a
 2 count of the PDUs originating from site A that were delivered by the network to
 3 probe B (known by probe B from its counters). Probes A and B acquire these
 4 measurements for all other VCs as well.

5 Referring to Figs. 2, 3A and 3B, acquisition of the aforementioned SLA
 6 measurements by endpoint probes A and B is described. When probe A reaches
 7 the activation time of a periodic measurement cycle, probe A sends an "initial
 8 request" message M_1 to probe B (see Fig. 2, Fig. 3A, blocks 40-44). The actual
 9 transmit time T_1 of message M_1 is not known by probe A until the message M_1 is
 10 sent; accordingly, message M_1 does not contain timestamp information indicating
 11 the time of its transmission.

12 Once the transmit time T_1 of message M_1 is known, probe A records a count
 13 ($OFF_A(i)$) of the PDUs (bound for probe B) that probe A has offered to the network
 14 up to time T_1 . The index i denotes the present periodic measurement cycle. The
 15 PDU count $OFF_A(i)$ can be a "delta" count, meaning that the count is the number of
 16 PDUs offered to the network from the time the last initial request message was sent
 17 (at the activation time within the previous $(i-1)$ periodic historical measurement cycle)
 18 up to time T_1 . However, for reasons that will become evident, the PDU count
 19 $OFF_A(i)$ and the other PDU counts recorded by the probes are preferably "raw"
 20 counts, meaning that the counts are running totals counted from the time of system
 21 boot up or the last counter roll-over. The use of raw counts requires the probes to
 22 maintain a record of the raw counts recorded during the message sequence in the
 23 previous measurement cycle $(i-1)$ in order to compute the delta PDU counts that
 24 pertain to the present data collection interval (which runs from the present activation
 25 time to the previous activation time).

26 As explained above, the count ($OFF_A(i)$) is needed by probe B to support SLA;
 27 however, since $OFF_A(i)$ is not known until after the transmission time T_1 of message
 28 M_1 is known, $OFF_A(i)$ (like T_1) cannot be sent to probe B in message M_1 . Time T_1 is
 29 measured relative to a local clock of probe A, which clock need not be synchronized
 30 with a corresponding local clock of probe B.

1 As shown in Fig. 2, message M_1 travels over the network VC and is received by
2 probe B at time T_2 (as measured relative to a local clock of probe B). In response to
3 reception of message M_1 , probe B records a count ($DEL_B(i)$) of the PDUs sent by
4 probe A that have been delivered by the network to probe B up to time T_2 (see Fig.
5 3A, block 46). Notably, because this PDU count $DEL_B(i)$ and the PDU count $OFF_A(i)$
6 measured by probe A are defined by the transmit and receive times of the same
7 message (M_1), these counts relate to correlated time periods at probes A and B;
8 thus, these counts can be directly compared (once converted to delta counts) to
9 obtain a meaningful measure of one-way (A to B) data delivery ratio over the VC
10 during a given data collection period.

11 In response to reception of initial request message M_1 , probe B sends a
12 "reply_B" message M_2 to probe A (see Fig. 2, Fig. 3A, blocks 48) at time T_3 . The
13 actual transmit time T_3 of message M_2 is not known by probe B until the message M_2
14 is sent; accordingly, message M_2 does not contain timestamp information indicating
15 the time of its transmission. However, because the receive time T_2 of initial request
16 message M_1 is known by probe B, a timestamp indicating the value of time T_2 is sent
17 by probe B to probe A in reply_B message M_2 .

18 Once the transmit time T_3 of message M_2 is known, probe B records a count
19 ($OFF_B(i)$) of the PDUs (bound for probe A) that probe B has offered to the network
20 up to time T_3 (see Fig. 3A, block 48). As explained above, this count ($OFF_B(i)$) is
21 needed by probe A to support SLA; however, since $OFF_B(i)$ is not known until after
22 the transmission time T_3 of message M_2 is known, $OFF_B(i)$ (like T_3) cannot be sent to
23 probe A in message M_2 . Like time T_2 , time T_3 is measured relative to the local clock
24 of probe B.

25 As shown in Fig. 2, message M_2 travels over the network VC and is received
26 by probe A at time T_4 (as measured relative to probe A's local clock). In response
27 reception of message M_2 , probe A records a PDU count ($DEL_A(i)$) of the PDUs sent
28 by probe B that have been delivered by the network to probe A up to time T_4 (see
29 Fig. 3A, block 50). Because this PDU count $DEL_A(i)$ and the PDU count $OFF_B(i)$
30 measured by probe B are defined by the transmit and receive times of the same
31 message (M_2), they relate to correlated time periods at probes A and B; thus, these

1 counts can be directly compared (once converted to delta counts) to obtain a
 2 meaningful measure of one-way (B to A) data delivery ratio over the VC over a given
 3 data collection period.

4 As shown in Fig. 2, as this point in time (after T_4), probe A knows the values
 5 of times T_1 , T_2 , T_4 , and PDU counts $OFF_A(i)$ and $DEL_A(i)$. To provide probe B with
 6 the aforementioned information necessary to support SLA, in response to reception
 7 of $reply_B$ message M_2 , probe A sends a "reply_A" message M_3 to probe B (see Fig. 2,
 8 Fig. 3A, block 52) at time T_5 containing timestamps indicating the values of times T_1
 9 and T_4 as well as the value of count $OFF_A(i)$. Reply_A message M_3 travels over the
 10 network VC and is received by probe B at time T_6 (Fig. 3A, block 54). At this point,
 11 probe B has acquired all of the measurements required to support SLA, i.e., the
 12 timestamps indicating the values of times T_1 through T_4 and PDU counts $OFF_A(i)$ and
 13 $DEL_B(i)$. Accordingly, probe B can compute the round trip delay (RTD) as the
 14 difference between the overall round trip time ($T_4 - T_1$) less the far-end turn around
 15 time ($T_3 - T_2$):

$$16 \quad RTD = (T_4 - T_1) - (T_3 - T_2) = (T_4 - T_3) + (T_2 - T_1) \quad (1)$$

17 Note that this RTD calculation is valid despite the fact that probe A's reference
 18 clock is not synchronized with probe B's reference clock due to the subtraction of T_1
 19 from T_4 and T_2 from T_3 .

20 Further probe B can compute, from the raw PDU counts, the delta PDU counts
 21 for the present measurement cycle i (i.e., the counts of the PDUs accumulated in the
 22 collection data period ending at the time of the message exchange) as:

$$23 \quad \Delta OFF_A(i) = OFF_A(i) - OFF_A(i-1) ; \Delta DEL_B(i) = DEL_B(i) - DEL_B(i-1) \quad (2)$$

24 where $OFF_A(i-1)$ is the recorded (raw) count (sent to probe B in the last
 25 measurement cycle) of the number of PDUs bound for probe B that were offered by
 26 probe A to the network up to the transmit time of the initial request message M_1 sent
 27 by probe A in the previous measurement cycle ($i-1$), and $DEL_B(i-1)$ is the recorded
 28 (raw) count of the number of PDUs sent by probe A that were delivered by the
 29 network to probe B up to the receive time of the initial request message M_1 received
 30 by probe B in the previous measurement cycle ($i-1$).

1 As previously mentioned, PDU counts $\Delta\text{OFF}_A(i)$ and $\Delta\text{DEL}_B(i)$ can be directly
 2 compared to obtain a measure of the one-way VC data delivery ratio. For example,
 3 the one-way DDR can be computed as:

$$4 \quad \text{DDR}_{A\text{-to-B}} = \Delta\text{DEL}_B(i) / \Delta\text{OFF}_A(i) \quad (3)$$

5 The data delivery ratio can also be computed as a two-way performance metric.
 6 In this case, probe B simply stores counts $\Delta\text{OFF}_A(i)$ and $\Delta\text{DEL}_B(i)$ for later processing
 7 by a console 16 (Fig. 1), which also receives counts $\Delta\text{OFF}_B(i)$ and $\Delta\text{DEL}_A(i)$ from
 8 probe A. Accordingly, at the end of each of probe B's periodic historical
 9 measurement cycle, the computed RTD and PDU delta counts $\Delta\text{OFF}_A(i)$ and
 10 $\Delta\text{DEL}_B(i)$ (and/or the one-way DDR) are stored by probe B for later transmission to
 11 console 16 (Fig. 3B, block 62).

12 To provide probe A with the aforementioned information necessary to support
 13 SLA, in response to reception of the reply_A message M_3 , probe B sends a "final
 14 reply" message M_4 to probe B (see Fig. 2, Fig. 3B, block 56) at time T_7 containing a
 15 timestamp indicating the value of time T_3 as well as the value of PDU count $\text{OFF}_B(i)$.
 16 Final reply message M_4 travels over the network VC and is received by probe A at
 17 time T_8 (Fig. 3B, block 58). At this point, probe A has acquired all of the
 18 measurements required to support SLA, i.e., the timestamps indicating the values of
 19 times T_1 through T_4 and PDU counts $\text{OFF}_B(i)$ and $\text{DEL}_A(i)$. Accordingly, probe A can
 20 compute the round trip delay (RTD) using equation (1). Note that probes A and B
 21 both measure the round trip delay of the same set of messages, i.e., messages M_1
 22 and M_2 .

23 Further probe A can compute the delta PDU counts for the present
 24 measurement cycle i as:

$$25 \quad \Delta\text{OFF}_B(i) = \text{OFF}_B(i) - \text{OFF}_B(i-1) ; \Delta\text{DEL}_A(i) = \text{DEL}_A(i) - \text{DEL}_A(i-1) \quad (4)$$

26 where $\text{OFF}_B(i-1)$ is the recorded (raw) count (sent to probe A in the last
 27 measurement cycle) of the number of PDUs bound for probe A that were offered by
 28 probe B to the network up to the transmit time of the reply_B message M_2 sent by
 29 probe B in the previous measurement cycle ($i-1$), and $\text{DEL}_A(i-1)$ is the recorded (raw)
 30 count of the number of PDUs sent by probe B that were delivered by the network to

1 probe A up to the receive time of the reply_B message M₂ received by probe A in the
2 previous measurement cycle (i-1).

3 PDU counts $\Delta\text{OFF}_B(i)$ and $\Delta\text{DEL}_A(i)$ can be directly compared to obtain a
4 measure of the one-way VC data delivery ratio. For example, the one-way DDR can
5 be computed as:

$$6 \quad \text{DDR}_{B \rightarrow A} = \Delta\text{DEL}_A(i) / \Delta\text{OFF}_B(i) \quad (5)$$

7 Again, the data delivery ratio can also be computed as a two-way performance
8 metric. In this case, probe A simply stores counts $\Delta\text{OFF}_B(i)$ and $\Delta\text{DEL}_A(i)$ for later
9 processing by a console, which also receives counts $\Delta\text{OFF}_A(i)$ and $\Delta\text{DEL}_B(i)$ from
10 probe B. At the end of each of probe A's periodic historical measurement cycles, the
11 computed RTD and PDU counts $\Delta\text{OFF}_B(i)$ and $\Delta\text{DEL}_A(i)$ (and/or the one-way DDR)
12 are stored by probe A for later transmission to console 16 (Fig. 3B, block 60).

13 One advantage of transmitting raw PDU counts rather than delta PDU counts in
14 the SLA message is that the loss of an SLA message PDU does not result in the
15 loss of PDU count data. If a probe does not receive PDU count information within a
16 measurement cycle due to a lost message PDU (message M₃ or M₄), the probe can
17 simply wait until the next measurement cycle and compute the delta PDU counts for
18 two consecutive data collection time intervals rather than the normal single data
19 collection time interval.

20 Another advantage of maintaining raw PDU counts rather than delta PDU
21 counts is that message contention situations are more easily managed. Specifically,
22 it is possible that two probes initiate message sequences with each other
23 substantially simultaneously (e.g., probe A transmits an initial request message M₁
24 and probe B subsequently transmits an initial request message M₁ prior to reception
25 of probe A's message). Under these circumstances, when raw counts are used, the
26 delta counts can be computed simply by subtracting the last-received PDU count
27 values from the raw counts from the previous measurement cycle.

28 As an enhancement to the data delivery ratio measurement described above,
29 the data delivery ratio can be determined for plural throughput ranges for certain
30 data transmission protocols. For example, additional counters can be used to
31 separately accumulate counts of PDUs where the throughput is below a committed

1 data rate, and PDUs where the throughput is between the committed and burst data
2 rate. This allows reports of DDR for each throughput range. The committed or peak
3 rate of data flow can be specified differently for different transmission protocols. For
4 example, for ATM, the peak rate of data flow is the Peak Cell Rate (PCR), and for
5 Frame Relay the peak rate of data flow is, $B_c + B_e$ (where B_c is the number of bits
6 the service provider is committed to deliver during a time interval, and B_e is the
7 number of bits in excess of B_c that the service provider might deliver during the time
8 interval). Thus, for example, for the Frame Relay protocol, the count of non-burst
9 offered frames (PDUs) can be expanded to add other two other counters: frames in
10 seconds where the throughput is below B_c ; and frames in seconds where the
11 throughput is between B_c and $B_c + B_e$.

12 Optionally, to minimize network impact, inter-probe messaging can be
13 enabled/disabled on a per-VC basis. This allows the user to disable the message
14 generation on VCs that do not have probes on each end.

15 In addition to measuring round trip delay (RTD) and data delivery ratio (DDR),
16 the probes of the present invention can also calculate network availability. Unlike
17 RTD and DDR, Availability is measured passively, i.e., without the assistance of SLA
18 messages injected into the data traffic. Instead determination of Availability relies on
19 analyzing information provided by the network operating under a particular protocol.

20 When the probe is operating under the Frame Relay data transmission
21 protocol, the availability of a circuit can be determined from the local management
22 interface LMI activity between the endpoint user equipment and the network
23 switches.

24 A circuit is deemed to be capable of accepting and delivering user frames
25 whenever one of the following conditions is present on both of its ends: a local
26 management interface (LMI) STATUS message is generated by the network in
27 response to each STATUS ENQUIRY message, and the LMI STATUS message
28 indicates that the circuit is valid and active; or No LMI STATUS ENQUIRY message
29 has ever been seen by the probe (in other words, no LMI is in use on the link).

30 More specifically, at any one time, a circuit end-point can be in any one of the
31 following states:

- 1 1. Link Down - user side (LMI requests are not being generated)
- 2 2. Link Down - network side (no LMI response to user generated requests)
- 3 3. Link Up, Data Link Connection Identifier (DLCI) invalid
- 4 4. Link Up, DLCI inactive
- 5 5. Link Up, DLCI active

6 From the service provider's standpoint, the circuit is considered available if both
7 ends of the VC are in state 1 or 5. To facilitate this determination, the probe
8 maintains a counter for each circuit, that accumulates a count of the number of
9 seconds that that end of the circuit was ever not in either of these states. If a probe
10 is assuming the role of the user's equipment when the user's equipment fails to
11 participate, state 1 cannot occur because the probe itself guarantees that LMI
12 requests are generated. Note that, because there is no required time
13 synchronization between probes, the console cannot accurately "line up" periodic
14 historical samples from both ends of the VC, so the console generally must pick one
15 end of each circuit on which to base the availability (e.g., probe A in Fig. 1).
16 Alternatively, the console can query the probes at both ends of a circuit and use the
17 weighted average of the two available times.

18 In the case of many Frame Relay switching networks, circuit status is transferred
19 through the network so that, if the far-end site is in state 1, the near-end DLCI is
20 marked inactive. If looking only at the near-end DLCI's statistics, the SLA report will
21 show that the VC was down even though it was caused by far-end user equipment
22 failure. If the probes act in the role of the user equipment, when it fails to LMI poll,
23 this can be avoided.

24 In the case of ATM networks, availability cannot be derived from LMI messages.
25 In some cases, LMI does not exist on an ATM network. OAM cell flows and cell
26 delineation are the two factors that can be used to determine if a VC is active or not,
27 and therefore available.

28 The first factor in availability is cell delineation. If the network is out of cell
29 delineation, all VCs on that network are unavailable. The second factor in deciding
30 availability is AIS or RDI cell flow. These OAM cells are sent into the access line on

1 a VC to inform the network that service on that VC is not available. Any second
2 where an AIS or RDI cell is received is considered an unavailable second.

3 Availability is measured on a per-circuit basis and stored in the probe as a count
4 of the number of seconds in which the circuit was not available for any part of that
5 second.

6 From time to time (e.g., a time period equal to or greater than periodic
7 measurement cycle, up to several days), probes A and B download the
8 aforementioned historical data samples taken at periodic intervals to a console (Fig.
9 1). For example, the probes collect the SLA parameters in the standard periodic
10 measurement intervals with a two day history. Data can be periodically collected
11 and stored in a database in accordance with a data collection process running on a
12 database server.

13 The format and content of the measurements provided to the console for
14 each periodic measurement interval is shown in Table 1.

Table 1

Object Name	Type	Description
Round-trip delay	16-bit Integer	A sample of the number of milliseconds of round-trip delay on this circuit. This is the RTD the inter-probe messages experienced during the history interval. If this value is -1, the measurement is not available.
Unavailable seconds	16-bit Counter	Number of seconds this circuit was not available during the history interval.
Far-end offered PDUs	32-bit Counter	Number of non-burst PDUs offered on the far-end of this circuit since the last successful SLA conversation. If this statistic is not available for the interval, its value is 0.
Delivered PDUs	32-bit Counter	Number of PDUs received during the SLA measurement interval. If the Far-end offered PDUs statistic is not available for this interval, this value is also 0.

1 Referring to Fig. 1, console 16 is in communication with probe A via a
2 communication line 28 (e.g., a local area network (LAN)) and with probe B via a
3 communication line 29, and retrieves the information collected by probes A and B to
4 process the information for display and/or recordation. While shown for illustrative
5 purposes as separate communication lines, it will be understood that the
6 communication lines connecting probes A and B can be virtual circuits over the
7 switching network or any other suitable interconnection, and the present invention is
8 not limited to any particular mechanization for interconnecting the probes and
9 console. Console 16 is typically implemented by a conventional personal computer;
10 however, other forms of computers, such as a Sun, Hewlett Packard, or IBM Unix
11 workstation, may also be equipped and utilized as the console in substantially the
12 same manner described below. Specifically, console 16 may be implemented by an
13 IBM-compatible personal computer preferably equipped with a mouse, monitor,
14 keyboard and base. The base commonly contains the processors, memory and
15 communications resources, such as internal/external modems or other
16 communications cards for the console. The console includes software for analyzing
17 the data collected by the corresponding probe or probes and displaying the
18 information to an operator, for example, in a manner similar to that described in U.S.
19 Patent Application Serial No. 08/746,416, the disclosure of which is incorporated
20 herein by reference in its entirety.

21 Further, console 16 can utilize short term and long term databases to
22 maintain data for extended periods of time. The databases may be implemented by
23 any conventional or commercially available database. Console 16 may operate as a
24 stand-alone console coupled to the probes, or in a client/server configuration
25 wherein a server (i.e., a computer system as described above preferably utilizing a
26 windows NT environment) performs substantial interactions with the probe and
27 conveys probe information to its clients (i.e., computer systems as described above
28 preferably utilizing a Windows 95 NT or Unix environment). Users may also
29 communicate with the probe directly for the data collection requests and for
30 providing configuration parameters to the probe.

1 While console 16 shown in Fig. 1 as a single console, it will be understood
2 that the present invention is not limited to a single-console embodiment, and more
3 than one console can be used to process and display monitoring information. For
4 example, two or more consoles at different sites or locations can be used to receive
5 data from one or more probes. In such a plural-console arrangement, the consoles
6 would preferably be interconnected (e.g., via the network or by other communication
7 lines) so that each console can display or record all of the monitoring information of
8 relevance to a related site or set of sites. Accordingly, the term "console," as used
9 herein, can be a single unit or plural, interconnected units located at one or more
10 sites or locations, wherein the unit(s) process, display, record and/or archive system
11 performance data in any format, such as those disclosed in U.S. Patent Application
12 Serial No. 08/746,4160.

13 Console 16 receives the delta PDU counts from both probes A and B, thereby
14 allowing console 16 to compute the two-way data delivery ratio for the VC
15 connecting sites A and B. Specifically, DDR is calculated from the aggregate of the
16 counts from both probes by:

$$17 \quad \text{DDR} = (\Delta\text{DEL}_A + \Delta\text{DEL}_B) / (\Delta\text{OFF}_A + \Delta\text{OFF}_B) \quad (6)$$

18 Again, note that calculation of DDR by this method works even if inter-probe
19 message PDUs are lost, because the inter-probe messages carry raw counts, and
20 the delta PDU counts are computed by the probes themselves. The round trip delay
21 (RTD) and DDR performance metrics can then be manipulated, recorded and
22 displayed by the console 16 in any suitable manner to depict the SLA performance
23 at particular times or over a period of time.

24 Note that, for a single transmission circuit, both of the probes at the respective
25 ends of that circuit will generally provide the RTD measurement (i.e., two identical
26 measurements). This redundancy allows the console to receive a valid RTD even
27 when one of the two probes fails to receive the necessary timestamp information to
28 calculate the RTD in a measurement cycle (e.g., when the final message PDU is
29 lost).

1 By way of non-limiting example, the SLA data can be presented in single
 2 tabular report as represented in the example below. Numerous other reports can be
 3 prepared from the accumulated data.

4

5 Service Level Verification Report

6 Network: Acme Janitorial Supply
 7 Period: 7/1/97-7/31/97

8 Circuit Designation	Availability	Delay	DDR
9 NY-to-Boston 718	100.000%	78ms	99.980%
10 NY-to-Detroit 719	99.943%	105ms	98.990%
11 NY-to-Charlotte 722	96.370%	113ms	99.567%
12 NY-to-Atlanta 725	97.065%	145ms	99.101%
13 NY-to-Los Angeles 731	98.950%	138ms	98.781%
14 Overall Average	98.850%	114ms	99.210%

16 Where the service offering describes several classes of service, each having
 17 their own service level, reports can be by service class. This can be implemented by
 18 creating each service class as a separate network. For example, all of the circuits
 19 with the SNA class of service can be configured to be in the "SNA" network.
 20 Preferably, the database stores a service class parameter for each circuit. This can
 21 be used to create reports by service class.

22 Optionally, the user can configure a single, weekly maintenance window of
 23 time that is automatically excluded from the SLA measurements. The window is
 24 specified by the day of the week, its start time and its duration. To automatically
 25 exclude periods of unavailability due to maintenance activity that occurs outside of a
 26 regularly scheduled maintenance window, custom integration to a trouble-ticketing
 27 system is required. Alternatively, a special type of data editor can be used to
 28 manually control exclusion of measurements during periods of unavailability due to
 29 maintenance.

1 While the above-described messaging sequence and collected measurement
2 data support both round-trip delay (RTD) and data delivery ratio (DDR) network
3 performance metrics, it will be understood that these metrics are separable and
4 need not be determined in conjunction with each other, i.e., RTD can be determined
5 without determining DDR, and DDR can be determined without determining RTD.
6 Moreover, while equations (1) through (6) set forth specific expressions for
7 determining RTD and DDR, the present invention is not limited to these specific
8 expressions, and the timestamp and PDU count measurements can be used to
9 calculate different or modified measures of network delay and percentages of
10 successful (or unsuccessful) data transmission.

11 In accordance with another aspect of the present invention, the inter-probe
12 message structure used to transmit inter-probe messages M_1 through M_4 is designed
13 to allow each inter-probe message to be sent within a single PDU through the
14 switching network, irrespective of the data transmission protocol(s) employed in the
15 network, including ATM, Frame Relay, etc. This feature of the present invention
16 advantageously permits the messaging scheme to be used with any conventional
17 packet switching network, including interworked switching networks. This feature is
18 achieved by minimizing the number of bytes required to send a single inter-probe
19 message, so that the message can be encapsulated within the smallest PDUs used
20 in the network, e.g., a single ATM cell. Fig. 4 illustrates the data payload of an inter-
21 probe message according to the exemplary embodiment of the present invention.
22 As seen in Fig. 4, the inter-probe message payload to be encapsulated within a PDU
23 includes 20 bytes.

24 A one-byte "Version #" field identifies the product version number to which the
25 message structure corresponds (e.g., a "2" indicates that the message structure
26 corresponds to the second release of the product).

27 A one-byte "Message Type" field identifies the current state of the overall SLA
28 message exchange, i.e., it identifies whether the message is M_1 , M_2 , M_3 or M_4 . A
29 value of 1 indicates that the message is the "initial request" message M_1 ; a value of
30 2 indicates that the message is the "reply_B" message M_2 ; a value of 3 indicates that
31 the message is the "reply_A" message M_3 ; and a value of 4 indicates that the

1 message is the "final reply" message M_4 . Even if the probes on both ends of the VC
2 begin an SLA message exchange with the same sequence number, the "Message
3 Type" field is used to resolve the potential ambiguity.

4 A two-byte "Sequence #" field identifies the SLA message exchange of which
5 the message is a part (somewhat analogous to the index "i" described above to
6 denote the measurement cycle).

7 A four-byte "SlaUpTime" field is a count of seconds since the probe booted,
8 SLA was enabled, or this particular circuit became available. The SlaUpTime of the
9 reporting probe is included in the inter-probe message so that reboots, SLA
10 enabling, and circuit reconfigurations can be detected and accounted for in the
11 probes' calculations of delta PDU counts.

12 A four-byte "Last RX timestamp" field is the timestamp (in milliseconds and
13 1024ths of a millisecond) of when the probe received its last SLA message PDU.
14 Specifically, for the reply_B message M_2 , this field contains the timestamp indicating
15 the value of time T_2 , and for the reply_A message M_3 , this field contains the timestamp
16 indicating the value of time T_4 . For messages M_1 and M_4 , this field contains no valid
17 data.

18 A four-byte "Last TX timestamp" field is the timestamp (in milliseconds and
19 1024ths of a millisecond) of when the probe offered its last SLA message PDU to
20 the network. Specifically, for the reply_A message M_3 , this field contains the
21 timestamp indicating the value of time T_1 , and for final reply message M_4 , this field
22 contains the timestamp indicating the value of time T_3 . For messages M_1 and M_2 ,
23 this field contains no valid data.

24 Both timestamp fields consist of 22 bits of integer milliseconds (upper 22 bits)
25 and 10 bits of 1024ths of a millisecond (lower 10 bits). This allows for about 69
26 minutes between timestamps before they wrap. Internal timestamp values are in
27 milliseconds, but the fractional part allows for the potential of higher precision
28 timestamps, and allows the serialization delay and internal delay factors to be
29 combined with less accumulated round off error. The timestamps may be relative to
30 the probe SlaUpTime.

1 A four-byte "non-burst offered PDUs" field is a free-running (raw) count of the
2 number of PDUs offered to the network on the transmission circuit, recorded at the
3 time the sending probe sent its first message in the message sequence.
4 Specifically, for the reply_A message M₃, this field contains the raw count, since
5 counter initialization/rollover, of the PDUs offered by probe A up to time T₁ (the time
6 at which probe A transmitted message M₁). For the final reply message M₄, this field
7 contains the raw count, since counter initialization/rollover, of the PDUs offered by
8 probe B up to time T₃ (the time at which probe B transmitted message M₂). For
9 messages M₁ and M₂ in the message sequence, this field contains no valid data.
10 The SlaUpTime is used to account for counter rollover and probe reboots.

11 32-bit unsigned counters can be used by probes A and B to store the PDU
12 counts in order to ensure infrequent counter rollovers with these counts. Assuming
13 a worst case of 128 octets frames at full DS3 (45MBps), the 32-bit PDU counter will
14 roll over only once every 27 hours. $45 \text{ MBps} / (128 * 8) = 44,000 \text{ fps}$. $2^{32} / 44,000 =$
15 $97,600 \text{ secs} = 27 \text{ hours}$. If the rollover rate becomes unmanageable for higher
16 transmission rates, this counter can be extended to 64 bits in a straight-forward
17 manner.

18 The above-described message data payload is encapsulated in a PDU for
19 transmission over the switching network. For probes operating under the DS3 ATM
20 protocol, the size and format of the SLA message PDU is fixed. For probes
21 operating under the Frame Relay protocol, the SLA message PDU size can be user
22 selectable.

23 Specifically, for a 53 byte ATM message PDU, the PDU contains the following: a
24 5 byte ATM header; 8 bytes of encapsulation; 20 bytes of message payload; 12
25 bytes pad; and an 8 byte ATM AAL5 trailer. For a Frame Relay message PDU, the
26 PDU contains the following: 2 bytes of Frame Relay header; 8 bytes of
27 encapsulation; 20 bytes of payload; 0 to 224 bytes of pad (to make total frame size
28 32 to 256 bytes); and 2 bytes of Frame Relay FCS.

29 As will be understood from the foregoing, a complete SLA message exchange
30 involves 4 messages being sent: 1) probe A sends an "initial request" message M₁
31 to probe B; 2) probe B replies with a "reply_B" message M₂; probe A replies with a

- 1 "reply_A" message M₃; and probe B replies with a "final reply" message M₄. The valid
 2 fields in the four messages are summarized in Table 2 (also see Fig. 2).

Table 2

Message Type Value

Fields	1	2	3	4
T2 (last received)		Valid		
T4 (last received)			Valid	
T1 (last sent)			Valid	
T3 (last sent)				Valid
Offered PDUs			Valid	Valid

- 3 If at any point in an SLA message exchange, an expected message doesn't
 4 arrive within a predetermined period of time, e.g., 10 seconds, the entire SLA
 5 conversation is terminated. On the initiating side, the conversation is reinitiated.
 6 This initial request is sent three times before giving up. The sequence number for
 7 each subsequent initial request is one greater than the previous initial request
 8 sequence number.

- 9 The inter-probe messages must be uniquely identifiable from the user traffic so
 10 that they can be transparently inserted into and extracted from any legal Frame
 11 Relay, ATM or other protocol data stream.

- 12 A probe operating in the Frame Relay protocol can use RFC 1490 and SNAP
 13 ethertype encapsulation with a proprietary (e.g., Visual Networks, Inc.) SLA
 14 ethertype value for the inter-probe messages. This encapsulation method is used
 15 for SLA messages regardless of the user's network encapsulation configuration.

- 16 An exemplary eight byte RFC 1490 encapsulation header for Frame Relay (all
 17 values in Hex) is shown in Table 3.

Table 3

Control	PAD	NLPID	OUI	Ethertype
03	00	80	00 00 00	88 65

1 The "Ethertype" field in the last two bytes of the SNAP header identifies a
2 proprietary protocol type of the packet. If the user data is also encapsulated with
3 RFC 1490, there is no problem in differentiating the user data from the probe
4 messages. If the user data is encapsulated with Cisco/EtherType, the probe
5 messages contain an invalid EtherType (03 00); thus, differentiation is easily
6 achieved. Even if some other (proprietary) user data encapsulation is in use, it is
7 highly unlikely that user data would be erroneously identified as an SLA message,
8 because all 8 encapsulation bytes would be required to match. As a safeguard,
9 however, a local configuration setting allows the user to disable SLA messaging.

10 A probe operating in the ATM protocol can use a similar technique based on
11 RFC 1483 encapsulation. The RFC1483 header is followed by SNAP-Ethertype and
12 a proprietary SLA EtherType value. The resulting eight byte RFC 1483
13 encapsulation header for ATM (all values in Hex) is shown in Table 4.

Table 4

DSAP	SSAP	Control	OUI	EtherType
AA	AA	03	00 00 00	88 65

14 While the above messaging sequence protocol and message structure are
15 particularly advantageous because of the ability to encapsulate each message within
16 an single PDU of any conventional data transmission protocol, it will be understood
17 that the present invention can be carried out with other message protocols and
18 message structures. By way of example, in accordance with another embodiment of
19 the present invention, the aforementioned time stamps and PDU counts are
20 transmitted using an alternative messaging scheme which is particularly suited to the
21 Frame Relay protocol.

22 Specifically, probe A sends a first message M_1 containing a timestamp of the
23 transmit time, T_1 , and the count of the number of PDUs offered ($OFF_A(i)$) by site A
24 (bound for site B) to the network up to time T_1 . Since the time T_1 is not precisely
25 known until message M_1 is actually transmitted, the transmitted values of T_1 and
26 $OFF_A(i)$ are estimated.

1 When the message M_1 is received by probe B on the other end of the circuit,
2 probe B records a timestamp of the receive time, T_2 , of the message M_1 . Probe B
3 then sends a reply message M_2 back to probe_A at time T_3 . Message M_2 includes
4 timestamps indicating the values of times T_2 and T_3 as well as the count of number
5 of PDUs offered ($OFF_B(i)$) by site B (bound for site A) to the network up to time T_3 .
6 Once again, time T_3 and PDU count $OFF_B(i)$ must be estimated, since they are not
7 known precisely until the actual transmit time T_3 of the message M_2 . Alternatively,
8 timestamps T_1 and T_3 must be taken as close as possible to the actual time the
9 messages are transmitted so that time spent waiting for idle time (a break in the user
10 traffic) is not included in the measurement. Upon receiving the reply message M_2 at
11 time T_4 , probe A can compute the round-trip delay in accordance with equation (1).
12 This final result is then sent back to probe B in one final message M_3 , so that probe
13 B can record the result as well (probe B is not sent a timestamp for time T_4 and
14 cannot, therefore, calculate RTD).

15 The above-described inter-probe messages conform to message format shown
16 in Table 5.

Table 5

Byte 0	Byte 1	Byte 3	Byte 4
Version #		Message Type	
Sequence #			
T ₁			
T ₂			
T ₃			
RTD (ms)			
SlaUpTime			
Non Burst Offered PDUs			

17 ... pad with 0x00 to 200 bytes (including encapsulation layers)

18 The fields T_1 , T_2 and T_3 contain the timestamps for times T_1 , T_2 and T_3 ,
19 respectively, and the RTD field contains the round trip delay time. The
20 Measurement Type field indicates whether the message is an M_1 , M_2 or M_3 type
21 message. The remaining fields contain information similar to the information
22 contained in the corresponding fields described above in conjunction with Fig. 4.

1 As will be understood from above, a complete message exchange sequence
 2 includes three messages: 1) probe A sends an "initial request" message to probe B;
 3 2) probe B replies with a "reply" message; and 3) probe A replies with a "result reply"
 4 message. The valid fields in the four messages are summarized in Table 6.

Table 6

Message Type value

Fields	0	1	2
Sequence #	Valid	Valid	Valid
T ₁	Valid	Valid	Valid
T ₂		Valid	Valid
T ₃		Valid	Valid
RTD			Valid
Offered PDUs	Valid	Valid	

5 To accommodate service level agreements that specify delay in terms of
 6 frame sizes other than 200 bytes, the inter-probe message size can be user
 7 programmable. Having described preferred embodiments of a new and improved
 8 method and apparatus for performing service level analysis of communications
 9 network performance metrics, it is believed that other modifications, variations and
 10 changes will be suggested to those skilled in the art in view of the teachings set forth
 11 herein. It is therefore to be understood that all such variations, modifications and
 12 changes are believed to fall within the scope of the present invention as defined by
 13 the appended claims.

PCT/US 98/23891
IPEA/US 28 MAY 1999What is claimed is:

1. An apparatus for acquiring data used to perform service level analysis of network performance metrics in a data transmission system comprising at least a first site and a second site and at least a first data transmission circuit establishing a communication path through a switching network between the first and second sites,
5 the apparatus comprising:

at least a first data acquisition device and a second data acquisition device respectively associated with said first and second sites;

said first data acquisition device inserting a first inter-device message into
10 data traffic bound for the second site that is offered by the first site to the switching network over the first data transmission circuit, the first inter-device message delimiting a set of the data traffic bound for the second site from subsequent data traffic;

said second data acquisition device receiving the first inter-device message
15 with data traffic originating from the first site and delivered by the switching network to the second site over the first data transmission circuit;

said first data acquisition device determining a count OFF_A indicating an amount of data traffic offered by the first site in the set delimited by the first inter-device message;

20 said second data acquisition device determining a count DEL_B indicating an amount of data traffic, in the set delimited by the first inter-device message, that is delivered by the switching network to the second site over the first data transmission circuit.

2. The apparatus of claim 1, wherein said first data acquisition device
25 transmits the count OFF_A to said second data acquisition device, and said second data acquisition device stores data delivery performance data based on the count OFF_A and the count DEL_B .

3. The apparatus of claim 1, wherein:

said second data acquisition device inserts a second inter-device message
30 into data traffic bound for the first site that is offered by the second site to the switching network over the first data transmission circuit, the second inter-device

message delimiting a set of the data traffic bound for the first site from subsequent data traffic;

said first data acquisition device receiving the second inter-device message with data traffic originating from the second site and delivered by the switching
5 network to the first site over the first data transmission circuit;

said second data acquisition device determining a count OFF_B indicating an amount of data traffic offered by the second site in the set delimited by the second inter-device message; and

said first data acquisition device determining a count DEL_A indicating an
10 amount of data traffic, in the set delimited by the second inter-device message, that is delivered by the switching network to the first site over the first data transmission circuit.

4. The apparatus of claim 3, wherein said second data acquisition device transmits the count OFF_B to said first data acquisition device, and said first data
15 acquisition device stores data delivery performance data based on the count OFF_B and the count DEL_A .

5. The apparatus of claim 3, wherein:

said first data acquisition device transmits a plurality of first inter-device messages, such that successive first inter-device messages define first data
20 collection intervals corresponding to successive sets of data traffic bound for the second site, said first data acquisition device determining the count OFF_A for each of the first data collection intervals, and said second data acquisition device determining the count DEL_B for each of the first data collection intervals;

said second data acquisition device transmits a plurality of second inter-
25 device messages, such that successive second inter-device messages define second data collection intervals corresponding to successive sets of data traffic bound for the first site, said second data acquisition device determining the count OFF_B for each of the second data collection intervals, and the first data acquisition device determining the count DEL_A for each of the second data collection intervals.

30 6. The apparatus of claim 5, wherein:

for each of the first data collection intervals, at least one of said first and second data acquisition devices determines a count ΔDEL_A indicating a number of protocol data units (PDUs) originating from the second site that are delivered by the switching network to the first site over the first data transmission circuit during the first data collection interval, and a count ΔOFF_B indicating a number of PDUs bound for the first site that are offered by the second site to the switching network over the first data transmission circuit during the first data collection interval; and

for each of the second data collection intervals, at least one of said first and second data acquisition devices determines a count ΔDEL_B indicating a number of PDUs originating from the first site that are delivered by the switching network to the second site over the first data transmission circuit during the second data collection interval, and a count ΔOFF_A indicating a number of PDUs bound for the second site that are offered by the first site to the switching network over the first data transmission circuit during the second data collection interval.

7. The apparatus of claim 6, further comprising a console in communication with said first and second data acquisition devices, wherein at least one of said console, said first data acquisition device and said second data acquisition device determines data delivery performance of the first data transmission circuit as a function of the counts ΔDEL_A , ΔOFF_B , ΔDEL_B and ΔOFF_A .

8. The apparatus of claim 7, wherein at least one of said console, said first data acquisition device and second data acquisition device calculates a data delivery ratio of a number of PDUs delivered on the first data transmission circuit by the switching network to a number of PDUs offered to the switching network on the first transmission circuit as $(\Delta DEL_A + \Delta DEL_B) / (\Delta OFF_A + \Delta OFF_B)$.

9. The apparatus of claim 7, wherein at least one of said console, said first data acquisition device and said second data acquisition device determines a round trip delay of data transmitted over the first data transmission circuit based on the transmit and receive times of the first and second inter-device messages.

10. The apparatus of claim 7, wherein at least one of said console, said first data acquisition device and said second data acquisition device determines an amount of time that the first data transmission circuit was available during a predetermined time period.

5 11. The apparatus of claim 6, wherein:

said first data acquisition device computes a one-way data delivery performance metric for the first data transmission circuit as a function of the counts ΔDEL_A and ΔOFF_B ; and

10 said second data acquisition device computes a one-way data delivery performance metric for the first data transmission circuit as a function of the counts ΔDEL_B and ΔOFF_A .

12. The apparatus of claim 11, wherein:

15 said first data acquisition device calculates a one-way data delivery ratio of a number of PDUs originating from the second site that are delivered by the switching network to the first site over the first data transmission circuit to a number of PDUs bound for the first site that are offered by the second site to the switching network over the first data transmission circuit as $\Delta\text{DEL}_A / \Delta\text{OFF}_B$; and

20 said second data acquisition device calculates a one-way data delivery ratio of a number of PDUs originating from the first site that are delivered by the switching network to the second site over the first data transmission circuit to a number of PDUs bound for the second site that are offered by the first site to the switching network over the first data transmission circuit as $\Delta\text{DEL}_B / \Delta\text{OFF}_A$.

25 13. The apparatus of claim 3, wherein said first and second data acquisition devices exchange a sequence of inter-device messages over the first transmission circuit, the sequence of inter-device messages including said first and second inter-device messages, wherein the count OFF_A is transmitted from said first data acquisition device to said second data acquisition device in a message in the sequence of inter-device messages, and the count OFF_B is transmitted from said second data acquisition device to said first data acquisition device in a message in
30 the sequence of inter-device messages.

14. The apparatus of claim 13, wherein said first and second data acquisition devices form each message of the sequence of inter-device messages in accordance with an inter-device message structure, wherein each inter-device message includes: a message type field identifying the inter-device message as a particular message within the sequence; and a count field for transmitting a count indicating an amount of data traffic offered to the switching network for transmission over the first data transmission circuit.

15. The apparatus of claim 14, wherein said first and second data acquisition devices form each message of the sequence of inter-device messages in accordance with the inter-device message structure, wherein each inter-device message further includes:

at least one timestamp field for transmitting timestamp data relating to timing of a message in the sequence.

16. The apparatus of claim 15, wherein said at least one timestamp field includes:

a first timestamp field for transmitting timestamp data indicating a time at which a previous message in the sequence was received; and

a second timestamp field for transmitting timestamp data indicating a time at which a previous message in the sequence was transmitted.

17. The apparatus of claim 15, said first and second data acquisition devices form each message of the sequence of inter-device messages such that each message is capable of being encapsulated in a single ATM cell and in a single frame relay frame.

18. The apparatus of claim 3, wherein:

said first data acquisition device maintains the count OFF_A indicating a number of protocol data units (PDUs) bound for the second site that are offered by the first site to the switching network over the first data transmission circuit;

said first data acquisition device maintains the count DEL_A indicating a number of PDUs originating from the second site that are delivered by the switching network to the first site over the first data transmission circuit;

5 said second data acquisition device maintains the count OFF_B indicating a number of PDUs bound for the first site that are offered by the second site to the switching network over the first data transmission circuit; and

said second data acquisition device maintains the count DEL_B indicating a number of PDUs originating from the first site that are delivered by the switching network to the second site over the first data transmission circuit.

10 19. The apparatus of claim 18, wherein:

the counts OFF_A , OFF_B , DEL_A and DEL_B maintained by said first and second data acquisition devices are running totals counted from a time of counter initialization or a time of counter rollover; and

15 said first and second data acquisition devices subtract counts from previous data collection intervals from the counts OFF_A , OFF_B , DEL_A and DEL_B to determine numbers of PDUs offered and delivered by said first and second data acquisition devices over the first data transmission circuit during a latest data collection interval.

20 20. The apparatus of claim 18, wherein the counts OFF_A , OFF_B , DEL_A and DEL_B maintained by said first and second data acquisition devices indicate numbers of PDUs offered and delivered by said first and second data acquisition devices during individual data collection intervals.

21. The apparatus of claim 3, wherein, said first and second data acquisition devices exchange a sequence of inter-device messages, such that:

25 said first data acquisition device transmits the first inter-device message to said second data acquisition device at a transmit time T_1 , the first inter-device message being received at said second data acquisition device at time T_2 ;

said second data acquisition device determines the count DEL_B as of time T_2 ;

30 said second data acquisition device transmits the second inter-device message to said first data acquisition device at a transmit time T_3 , the second inter-device message being received at said first data acquisition device at time T_4 ;

said first data acquisition device determines the count DEL_A as of time T_4 ;
said first data acquisition device transmits a third inter-device message to said
second data acquisition device containing the count OFF_A as of time T_1 ; and
said second data acquisition device transmits a fourth inter-device message
5 to said first data acquisition device containing the count OFF_B as of time T_3 .

22. The apparatus of claim 21, wherein said first and second data acquisition
devices exchange the sequence of inter-device messages during a periodic
measurement cycle.

23. The apparatus of claim 22, wherein said first and second data acquisition
10 devices exchange the sequence of inter-device messages containing data collected
during a data collection interval terminating during the periodic measurement cycle.

24. The apparatus of claim 21, wherein:
the third inter-device message contains timestamp information indicating
values of the transmit time T_1 of the first message and the receive time T_4 of the
15 second message; and
said second data acquisition device calculates the round trip delay over the
first transmission circuit using the values of times T_1 , T_2 , T_3 and T_4 .

25. The apparatus of claim 24, wherein:
the second inter-device message contains timestamp information indicating a
20 value of the receive time T_2 of the first message;
the fourth inter-device message contains timestamp information indicating a
value of the transmit time T_3 of the second message; and
said first data acquisition device calculates the round trip delay over the first
transmission circuit using the values of times T_1 , T_2 , T_3 and T_4 .

25 26. The apparatus of claim 3, wherein:
the switching network is an interworked network employing plural, different
data transmission protocols;

PCT/US 98/2389
IP/EA/US 28 MAY 1999

said first data acquisition device constructs the first inter-device message such that the first inter-device message is capable of being encapsulated in a single protocol data unit of each of the plural data transmission protocols; and

5 said second data acquisition device constructs the second inter-device message such that the second inter-device message is capable of being encapsulated in a single protocol data unit of each of the plural data transmission protocols.

27. The apparatus of claim 26, wherein:

10 the interworked network includes an asynchronous transfer mode (ATM) network employing an ATM data transmission protocol and a frame relay network employing a frame relay data transmission protocol;

 the first site is a site on one of the ATM network and the frame relay network, and the second site is a site on the other of the ATM network and the frame relay network;

15 said first data acquisition device constructs the first inter-device message such that the first inter-device message is capable of being encapsulated in a single ATM cell and a single frame relay frame; and

20 said second data acquisition device constructs the second inter-device message such that the second inter-device message is capable of being encapsulated in a single ATM cell and a single frame relay frame.

28. An apparatus for measuring round-trip delay in a data transmission system comprising at least a first site and a second site and at least a first data transmission circuit establishing a communication path through a switching network between the first and second sites, the apparatus comprising:

25 at least a first data acquisition device and a second data acquisition device respectively associated with said first and second sites and being configured to exchange a sequence of inter-device messages, such that:

30 said first data acquisition device transmits a first message to said second data acquisition device at time T_1 , the first message being received at said second data acquisition device at time T_2 ;

said second data acquisition device transmits at time T_3 a second message to said first data acquisition, the second message being received at said first data acquisition device at time T_4 ;

5 said first data acquisition device transmits to said second data acquisition device a third message containing timestamp information indicating values of the transmit time T_1 of the first message and receive time T_4 of the second message; and

said second data acquisition device calculates the round trip delay over the first transmission circuit using the values of times T_1 , T_2 , T_3 and T_4 .

10 29. The apparatus of claim 28, wherein said second data acquisition device computes the round trip delay as:

$$(T_4 - T_3) + (T_2 - T_1).$$

30. The apparatus of claim 29, wherein:

15 said second data acquisition device transmits in the second message timestamp information indicating a value of the receive time T_2 of the first message;

said second data acquisition device transmits to said first data acquisition device a fourth message containing timestamp information indicating a value of the transmit time T_3 of the second message; and

20 said first data acquisition device calculates the round trip delay over the first transmission circuit using the values of times T_1 , T_2 , T_3 and T_4 .

31. The apparatus of claim 30, wherein said first data acquisition device computes the round trip delay as:

$$(T_4 - T_3) + (T_2 - T_1).$$

25 32. A data transmission system capable of performing service level analysis of network performance metrics, comprising:

a switching network;

at least a first site and a second site;

at least a first data transmission circuit establishing a communication path through the switching network between the first and second sites;

PCT/US 98/23891
IPEA/US 28 MAY 1999

at least a first data acquisition device and a second data acquisition device respectively associated with said first and second sites; and

a console in communication with at least one of said first and second data acquisition devices;

5 said first data acquisition device being configured to: collect first performance data related to performance of said first transmission circuit; insert a first inter-device message into data traffic bound for the second site that is offered by the first site to the switching network over the first data transmission circuit; and use the first inter-device message to delimit a set of the data traffic for which the first performance
10 data is collected;

 said second data acquisition device being configured to: receive the first inter-device message with data traffic originating from the first site and delivered by the switching network to the second site over the first data transmission circuit; and collect second performance data related to performance of said first transmission
15 circuit for the set of data traffic delimited by the first inter-device message;

 at least one of said first data acquisition device, said second data acquisition device and said console generating service level analysis data of at least one network performance metric from the first and second performance data; and

 said console providing as an output the service level analysis data.

20 33. The system of claim 32, wherein said first data acquisition device transmits the first performance data to said second data acquisition device, and said second data acquisition device stores historical performance data based on the first and second performance data.

 34. The system of claim 32, wherein:

25 said second data acquisition device is configured to: collect third performance data related to performance of said first transmission circuit; insert a second inter-device message into data traffic bound for the first site that is offered by the second site to the switching network over the first data transmission circuit; and use the second inter-device message to delimit a set of the data traffic on which the
30 third performance data is collected;

said first data acquisition device is configured to: receive the second inter-device message with data traffic originating from the second site and delivered by the switching network to the first site over the first data transmission circuit; and collect fourth performance data related to performance of said first transmission
5 circuit for the set of data traffic delimited by the second inter-device message;

at least one of said first data acquisition device, said second data acquisition device and said console generates service level analysis data of at least one network performance metric from the first, second, third and fourth performance data; and

10 said console provides as an output the service level analysis data.

35. The system of claim 34, wherein said second data acquisition device transmits the third performance data to said first data acquisition device, and said first data acquisition device stores historical performance data based on the third and fourth performance data.

15 36. The system of claim 34, wherein:

said first performance data collected by said first data acquisition device indicates an amount of data traffic offered by the first site in the set delimited by the first inter-device message;

20 said second performance data collected by said second data acquisition device indicates an amount of data traffic, in the set delimited by the first inter-device message, that is delivered by the switching network to the second site over the first data transmission circuit;

25 said third performance data collected by said second data acquisition device indicates an amount of data traffic offered by the second site in the set delimited by the second inter-device message; and

said fourth performance data collected by said first data acquisition device indicates an amount of data traffic, in the set delimited by the second inter-device message, that is delivered by the switching network to the first site over the first data transmission circuit.

37. The system of claim 34, wherein at least one of said first data acquisition device, said second data acquisition device and said console determines data delivery performance of the first data transmission circuit from the first, second, third and fourth performance data.

5 38. The system of claim 34, wherein:

said first data acquisition device determines as the first performance data a count OFF_A indicating a number of protocol data units (PDUs) bound for the second site that are offered by the first site to the switching network over the first data transmission circuit;

10 said first data acquisition device maintains as the fourth performance data a count DEL_A indicating a number of PDUs originating from the second site that are delivered by the switching network to the first site over the first data transmission circuit;

15 said second data acquisition device maintains as the third performance data a count OFF_B indicating a number of PDUs bound for the first site that are offered by the second site to the switching network over the first data transmission circuit; and

20 said second data acquisition device maintains as the second performance data a count DEL_B indicating a number of PDUs originating from the first site that are delivered by the switching network to the second site over the first data transmission circuit.

39. The system of claim 38, wherein:

25 said first data acquisition device transmits a plurality of first inter-device messages, such that successive first inter-device messages define first data collection intervals corresponding to successive sets of data traffic bound for the second site, said first data acquisition device determining the count OFF_A for each of the first data collection intervals, and said second data acquisition device determining the count DEL_B for each of the first data collection intervals;

30 said second data acquisition device transmits a plurality of second inter-device messages, such that successive second inter-device messages define second data collection intervals corresponding to successive sets of data traffic bound for the first site, said second data acquisition device determining the count

OFF_B for each of the second data collection intervals, and the first data acquisition device determining the count DEL_A for each of the second data collection intervals.

40. The system of claim 39, wherein:

5 for each of the first data collection intervals, at least one of said first data acquisition device, said second data acquisition device and said console determines a count ΔDEL_B indicating a number of PDUs originating from the first site that are delivered by the switching network to the second site over the first data transmission circuit during the first data collection interval, and a count ΔOFF_A indicating a number of PDUs bound for the second site that are offered by the first site to the
10 switching network over the first data transmission circuit during the second data collection interval; and

for each of the second data collection intervals, at least one of said first data acquisition device, said second data acquisition device and said console determines a count ΔDEL_A indicating a number of PDUs originating from the second site that are
15 delivered by the switching network to the first site over the first data transmission circuit during the second data collection interval, and a count ΔOFF_B indicating a number of PDUs bound for the first site that are offered by the second site to the switching network over the first data transmission circuit during the second data collection interval.

20 41. The system of claim 40, wherein at least one of said console, said first data acquisition device and said second data acquisition device determines data delivery performance of the first data transmission circuit as a function of the counts ΔDEL_A , ΔOFF_B , ΔDEL_B and ΔOFF_A .

25 42. The system of claim 41, wherein at least one of said console, said first data acquisition device and second data acquisition device calculates a data delivery ratio of a number of PDUs delivered on the first data transmission circuit by the switching network to a number of PDUs offered to the switching network on the first transmission circuit as $(\Delta\text{DEL}_A + \Delta\text{DEL}_B) / (\Delta\text{OFF}_A + \Delta\text{OFF}_B)$.

43. The system of claim 40, wherein:

the counts OFF_A , OFF_B , DEL_A and DEL_B maintained by said first and second data acquisition devices are running totals counted from a time of counter initialization or a time of counter rollover, and the counts ΔOFF_A , ΔOFF_B , ΔDEL_A and ΔDEL_B are determined by subtracting counts from previous data collection intervals

5 from the counts OFF_A , OFF_B , DEL_A and DEL_B .

44. The system of claim 39, wherein said first and second data acquisition devices exchange a sequence of inter-device messages over the first transmission circuit, the sequence of inter-device messages including said first and second inter-device messages, wherein at least one of the first and third performance data is

10 transmitted between said first and second data acquisition devices in a message in the sequence of inter-device messages.

45. The system of claim 44, wherein at least one of said console, said first data acquisition device and said second data acquisition device determines a round trip delay of data transmitted over the first data transmission circuit using timestamp

15 data transmitted in the sequence of inter-device messages.

46. The system of claim 44, wherein at least one of said first data acquisition device, said second data acquisition device and said console generates performance data indicating an amount of time that said first data transmission circuit was available during a predetermined time period.

20 47. The system of claim 44, wherein, said first and second data acquisition devices exchange the sequence of inter-device messages, such that:

said first data acquisition device transmits the first inter-device message to said second data acquisition device at a transmit time T_1 , the first inter-device message being received at said second data acquisition device at time T_2 ;

25 said second data acquisition device determines the count DEL_B as of time T_2 ;

said second data acquisition device transmits the second inter-device message to said first data acquisition device at a transmit time T_3 , the second inter-device message being received at said first data acquisition device at time T_4 ;

said first data acquisition device determines the count DEL_A as of time T_4 ;

said first data acquisition device transmits a third inter-device message to said second data acquisition device containing the count OFF_A as of time T_1 ; and

said second data acquisition device transmits a fourth inter-device message to said first data acquisition device containing the count OFF_B as of time T_3 .

5 48. The system of claim 47, wherein:

the second inter-device message contains timestamp information indicating a value of the receive time T_2 ;

the third inter-device message contains timestamp information indicating values of the transmit time T_1 of the first inter-device message and the receive time
10 T_4 of the second inter-device message;

the fourth inter-device message contains timestamp information indicating a value of the transmit time T_3 of the second message; and

at least one of said first data acquisition device, said second data acquisition device and said console calculates the round trip delay as:

15 $(T_4 - T_3) + (T_2 - T_1)$.

49. The system of claim 34, wherein:

said switching network is an interworked network employing plural, different data transmission protocols;

said first data acquisition device constructs the first inter-device message
20 such that the first inter-device message is capable of being encapsulated in a single protocol data unit of each of the plural data transmission protocols; and

said second data acquisition device constructs the second inter-device message such that the second inter-device message is capable of being encapsulated in a single protocol data unit of each of the plural data transmission
25 protocols.

50. The system of claim 49, wherein:

the interworked network includes an asynchronous transfer mode (ATM) network employing an ATM data transmission protocol and a frame relay network employing a frame relay data transmission protocol;

the first site is a site on one of the ATM network and the frame relay network, and the second site is a site on the other of the ATM network and the frame relay network;

5 said first data acquisition device constructs the first inter-device message such that the first inter-device message is capable of being encapsulated in a single ATM cell and a single frame relay frame; and

said second data acquisition device constructs the second inter-device message such that the second inter-device message is capable of being encapsulated in a single ATM cell and a single frame relay frame.

10 51. A method of performing service level analysis of network performance metrics in a data transmission system that includes at least first and second sites; at least first and second data acquisition devices respectively associated with the first and second sites; and at least a first data transmission circuit establishing a communication path through a switching network between the first and second sites,
15 the method comprising the steps of:

(a) collecting, at the first data acquisition device, first performance data related to performance of the first transmission circuit;

(b) collecting, at the second data acquisition device, second performance data related to performance of the first transmission circuit;

20 (c) inserting a first inter-device message into data traffic bound for the second site that is offered by the first site to the switching network over the first data transmission circuit;

(d) using the first inter-device message to delimit a set of the data traffic for which the first and second performance data are collected; and

25 (e) generating service level analysis data of at least one network performance metric from the first and second performance data.

52. The method of claim 51, further comprising the steps of:

(f) transmitting to the second data acquisition device a value of the first performance data; and

30 (g) storing at the second data acquisition device historical performance data based on the values of the first and second performance data.

53. The method of claim 51, further comprising the steps of:

(f) collecting, at the second data acquisition device, third performance data related to performance of the first transmission circuit;

5 (g) collecting, at the first data acquisition device, fourth performance data related to performance of the first transmission circuit;

(h) inserting a second inter-device message into data traffic bound for the first site that is offered by the second site to the switching network over the first data transmission circuit; and

10 (i) using the second inter-device message to delimit a set of the data traffic for which the third and fourth performance data are collected;

wherein step (e) includes generating service level analysis data of at least one network performance metric from the first, second, third and fourth performance data.

54. The method of claim 53, further comprising the steps of:

15 (k) transmitting to the first data acquisition device a value of the third performance data; and

(l) storing at the first data acquisition device historical performance data based on the values of the third and fourth performance data.

55. The method of claim 53, wherein:

20 step (a) includes maintaining as the first performance data an indication of an amount of data traffic offered by the first site in the set delimited by the first inter-device message;

25 step (b) includes maintaining as the second performance data an indication of an amount of data traffic, in the set delimited by the first inter-device message, that is delivered by the switching network to the second site over the first data transmission circuit;

step (f) includes maintaining as the third performance data an indication of an amount of data traffic offered by the second site in the set delimited by the second inter-device message; and

step (g) includes maintaining as the fourth performance data an indication of an amount of data traffic, in the set delimited by the second inter-device message, that is delivered by the switching network to the first site over the first data transmission circuit.

5 56. The method of claim 53, wherein step (e) includes determining data delivery performance of the first data transmission circuit from the first, second, third and fourth performance data.

57. The method of claim 53, wherein:

10 step (a) includes maintaining as the first performance data a count OFF_A indicating a number of protocol data units (PDUs) bound for the second site that are offered by the first site to the switching network over the first data transmission circuit;

15 step (b) includes maintaining as the second performance data a count DEL_B indicating a number of PDUs originating from the first site that are delivered by the switching network to the second site over the first data transmission circuit;

 step (f) includes maintaining as the third performance data a count OFF_B indicating a number of PDUs bound for the first site that are offered by the second site to the switching network over the first data transmission circuit; and

20 step (g) includes maintaining as the fourth performance data a count DEL_A indicating a number of PDUs originating from the second site that are delivered by the switching network to the first site over the first data transmission circuit.

58. The method of claim 57, wherein:

25 step (c) includes inserting a plurality of first inter-device messages into data traffic originating from the first site and bound for the second site, such that successive first inter-device messages define first data collection intervals corresponding to successive sets of data traffic bound for the second site;

 step (a) includes determining the count OFF_A for each of the first data collection intervals;

30 step (b) includes determining the count DEL_B for each of the first data collection intervals;

step (h) includes inserting a plurality of second inter-device messages into data traffic originating from the second site and bound for the first site, such that successive second inter-device messages define second data collection intervals corresponding to successive sets of data traffic bound for the first site;

5 step (f) includes determining the count OFF_B for each of the first data collection intervals; and

step (g) includes determining the count DEL_A for each of the first data collection intervals.

59. The method of claim 58, further comprising the steps of:

10 (j) determining, for each of the first data collection intervals, a count ΔDEL_B indicating a number of PDUs originating from the first site that are delivered by the switching network to the second site over the first data transmission circuit during the first data collection interval, and a count ΔOFF_A indicating a number of PDUs bound for the second site that are offered by the first site to the switching network over the
15 first data transmission circuit during the second data collection interval; and

(k) determining, for each of the second data collection intervals, a count ΔDEL_A indicating a number of PDUs originating from the second site that are delivered by the switching network to the first site over the first data transmission circuit during the second data collection interval, and a count ΔOFF_B indicating a
20 number of PDUs bound for the first site that are offered by the second site to the switching network over the first data transmission circuit during the second data collection interval.

60. The method of claim 59, wherein step (e) includes determining data delivery performance of the first data transmission circuit as a function of the counts
25 ΔDEL_A , ΔOFF_B , ΔDEL_B and ΔOFF_A .

61. The method of claim 60, wherein step (e) includes calculating a data delivery ratio of a number of PDUs delivered on the first data transmission circuit by the switching network to a number of PDUs offered to the switching network on the first transmission circuit as $(\Delta DEL_A + \Delta DEL_B) / (\Delta OFF_A + \Delta OFF_B)$.

62. The method of claim 59, wherein:

steps (a), (f), (g) and (b) respectively include determining the counts OFF_A , OFF_B , DEL_A and DEL_B as running totals counted from a time of counter initialization or a time of counter rollover; and

5 steps (j) and (k) include determining the counts ΔOFF_A , ΔOFF_B , ΔDEL_A and ΔDEL_B by subtracting counts from previous data collection intervals from the counts OFF_A , OFF_B , DEL_A and DEL_B .

63. The method of claim 58, wherein said first and second data acquisition devices exchange a sequence of inter-device messages over the first transmission
10 circuit, the sequence of inter-device messages including said first and second inter-device messages, the method further comprising the step of:

(j) transmitting at least one of the first and third performance data between said first and second data acquisition devices in a message in the sequence of inter-device messages.

15 64. The method of claim 63, further comprising the step of:

(k) determining a round trip delay of data transmitted over the first data transmission circuit using timestamp data transmitted in the sequence of inter-device messages.

65. The method of claim 63, further comprising the step of:

20 (k) generating performance data indicating an amount of time that said first data transmission circuit was available during a predetermined time period.

66. The method of claim 63, wherein:

25 step (c) includes transmitting the first inter-device message to said second data acquisition device at a transmit time T_1 , the first inter-device message being received at said second data acquisition device at time T_2 ;

step (a) includes determining the count OFF_A as of time T_1 ;

step (b) includes determining the count DEL_B as of time T_2 ;

step (h) includes transmitting the second inter-device message to said first data acquisition device at a transmit time T_3 , the second inter-device message being received at said first data acquisition device at time T_4 ;

step (f) includes determining the count OFF_B as of time T_3 ;

5 step (g) includes determining the count DEL_A as of time T_4 ;

the method further comprising the steps of:

(k) transmitting a third inter-device message to said second data acquisition device containing the count OFF_A ; and

10 (l) transmitting a fourth inter-device message to said first data acquisition device containing the count OFF_B .

67. The method of claim 66, wherein:

step (h) includes transmitting timestamp information indicating a value of the receive time T_2 within the second inter-device message;

15 step (k) includes transmitting timestamp information indicating values of the transmit time T_1 of the first inter-device message and the receive time T_4 of the second inter-device message in the third inter-device message;

step (l) includes transmitting timestamp information indicating a value of the transmit time T_3 of the second message in the fourth inter-device message;

the method further comprising the step of:

20 (m) calculating the round trip delay as: $(T_4 - T_3) + (T_2 - T_1)$.

68. The method of claim 53, wherein the switching network is an interworked network employing plural, different data transmission protocols;

25 step (c) includes constructing the first inter-device message such that the first inter-device message is capable of being encapsulated in a single protocol data unit of each of the plural data transmission protocols; and

step (h) includes constructing the second inter-device message such that the second inter-device message is capable of being encapsulated in a single protocol data unit of each of the plural data transmission protocols.

69. The method of claim 68, wherein:

the interworked network includes an asynchronous transfer mode (ATM) network employing an ATM data transmission protocol and a frame relay network employing a frame relay data transmission protocol, the first site is a site on one of the ATM network and the frame relay network, and the second site is a site on the other of the ATM network and the frame relay network, wherein

step (c) includes constructing the first inter-device message such that the first inter-device message is capable of being encapsulated in a single ATM cell and a single frame relay frame; and

step (h) includes constructing the second inter-device message such that the second inter-device message is capable of being encapsulated in a single ATM cell and a single frame relay frame.

70. A method of measuring data delivery performance of a data transmission circuit forming a transmission path between a first site and a second site through a switching network over which the first and second sites communicate, the method comprising the steps of:

(a) maintaining at a first site end of the data transmission circuit a count OFF_A of a number of protocol data units (PDUs) bound for the second site that are offered by the first site to the switching network over the data transmission circuit;

(b) maintaining at the first site end a count DEL_A of a number of PDUs originating from the second site that are delivered by the switching network to the first site over the data transmission circuit;

(c) maintaining at a second site end of the data transmission circuit a count OFF_B of a number of PDUs bound for the first site that are offered by the second site to the switching network over the data transmission circuit;

(d) maintaining at the second site end a count DEL_B of a number of PDUs originating from the first site that are delivered by the switching network to the second site over the data transmission circuit; and

(e) computing a measure of PDU delivery performance of the data transmission circuit from the counts OFF_A , OFF_B , DEL_A and DEL_B .

71. The method of claim 70, wherein step (e) includes calculating a data delivery ratio as the ratio of the number of PDUs delivered by the switching network

to the first and second sites over the data transmission circuit to the number of PDUs offered to the switching network by the first and second sites for transmission over the data transmission circuit.

72. The method of claim 70, further comprising the steps of:

5 (f) periodically transmitting the count OFF_A from the first site end to the second site end over the data transmission circuit;

(g) periodically transmitting the count OFF_B from the second site end to the first site end over the data transmission circuit;

10 (h) periodically storing at the first site end of the data transmission circuit a count ΔOFF_B of a number of PDUs bound for the first site that were offered by the second site to the switching network over the data transmission circuit during a data collection interval, and a count ΔDEL_A of a number of PDUs originating from the second site that were delivered by the switching network to the first site over the data transmission circuit during the data collection interval; and

15 (i) periodically storing at the second site end of the data transmission circuit a count ΔOFF_A of a number of PDUs bound for the second site that were offered by the first site to the switching network over the data transmission circuit during a data collection interval and a count ΔDEL_B of a number of PDUs originating from the first site that were delivered by the switching network to the second site over the data transmission circuit during the data collection interval;

20 wherein step (e) includes computing the measure of PDU delivery performance from the counts periodically stored in steps (h) and (i).

73. The method of claim 72, wherein step (e) includes calculating a data delivery ratio as the ratio of the number of PDUs delivered by the switching network to the first and second sites over the data transmission circuit to the number of PDUs offered to the switching network by the first and second sites for transmission over the data transmission circuit as $(\Delta DEL_A + \Delta DEL_B) / (\Delta OFF_A + \Delta OFF_B)$.

74. A method of calculating a round trip delay of data transmitted over a data transmission circuit forming a path through a packetized switching network between first and second sites, comprising the steps of:

30

(a) transmitting a first message from a first data acquisition device associated with the first site to a second data acquisition device associated with the second site at time T_1 , the first message being received at the second data acquisition device at time T_2 ;

5 (b) transmitting a second message from the second data acquisition device to the first data acquisition device at time T_3 , the second message being received at the first data acquisition device at time T_4 ;

10 (c) transmitting a third message from the first data acquisition device to the second data acquisition device containing timestamp information indicating values of the transmit time T_1 of the first message and the receive time T_4 of the second message; and

(d) computing the round trip delay using the values of times T_1 , T_2 , T_3 and T_4 .

75. The method of claim 74, wherein step (d) includes computing the round trip delay as:

15 $(T_4 - T_3) + (T_2 - T_1)$.

76. The method of claim 74, wherein step (b) includes transmitting in the second message timestamp information containing a value of the receive time T_2 of the first message, the method further comprising the step of:

20 (e) transmitting from the second data acquisition device to the first data acquisition device a fourth message containing a value of the transmit time T_3 of the second message;

wherein step (d) is performed by both the first data acquisition device and the second data acquisition device.

25 77. The method of claim 76, wherein step (d) includes computing the round trip delay as:

$(T_4 - T_3) + (T_2 - T_1)$.

78. A data transmission structure for transmitting a sequence of inter-device messages over a data transmission circuit forming a path through a switching

network between first and second sites, the switching network employing at least one data transmission protocol, the structure comprising:

a message type field identifying the inter-device message as a particular message within the sequence; and

5 a count field for transmitting an indicator of a number of protocol data units (PDUs) offered for transmission to the switching network over the data transmission circuit, wherein, when said message type field contains a first value, said count field contains a count OFF_A indicating a number of PDUs bound for the second site that are offered by the first site to the switching network over the first transmission circuit,
10 and when said message type field contains a second value, said count field contains a count OFF_B indicating a number of PDUs bound for the first site that are offered by the second site to the switching network over the first transmission circuit;

wherein each inter-device message conforming to the data transmission structure is capable of being encapsulated in a single protocol data unit of said at
15 least one data transmission protocol.

79. The structure of claim 78, wherein the switching network is an interworked switching network employing plural, different data transmission protocols, wherein each inter-device message conforming to the data transmission structure is capable of being encapsulated in a single protocol data unit of each of
20 the plural data transmission protocols.

80. The structure of claim 79, wherein said plural data transmission protocols includes at least one of: Asynchronous Transfer Mode (ATM), Frame Relay, High Level Data Link Control (HDLC), X.25, and tunneled protocols.

81. The structure of claim 79, wherein:
25 the interworked network includes an asynchronous transfer mode (ATM) network employing an ATM data transmission protocol and a frame relay network employing a frame relay data transmission protocol;

the first site is a site on one of the ATM network and the frame relay network, and the second site is a site on the other of the ATM network and the frame relay
30 network; and

each inter-device message conforming to the data transmission structure is capable of being encapsulated in a single ATM cell and in a single frame relay frame.

82. The structure of claim 78, wherein:

- 5 the sequence of data measurement messages includes first and third messages transmitted over the data transmission circuit from the first site to the second site in accordance with the data transmission structure, and second and fourth messages transmitted over the data transmission circuit from the second site to the first site in accordance with the data transmission structure;
- 10 when said message type field identifies the inter-device message as the third message in the sequence, said count field contains the count OFF_A ; and
- when said message type field identifies the inter-device message as the fourth message in the sequence, the count field contains the count OFF_B .

83. The structure of claim 78, further comprising:

- 15 a first timestamp field for transmitting timestamp data indicating a time at which a previous message in the sequence was received; and
- a second timestamp field for transmitting timestamp data indicating a time at which a previous message in the sequence was transmitted.

84. The structure of claim 83, wherein:

- 20 the sequence of data measurement messages includes first and third messages transmitted over the data transmission circuit from the first site to the second site in accordance with the data transmission structure, and second and fourth messages transmitted over the data transmission circuit from the second site to the first site in accordance with the data transmission structure;
- 25 when said message type field identifies the inter-device message as the second message in the sequence, said first timestamp field indicates a receive time of the first message;
- when said message type field identifies the inter-device message as the third message in the sequence, said first timestamp field indicates a receive time of the

second message, said second timestamp field indicates a transmit time of the first message, and said count field contains the count OFF_A ; and

when said message type field identifies the inter-device message as the fourth message in the sequence, said second timestamp field indicates a transmit
5 time of the second message, and the count field contains the count OFF_B .

85. A method of transmitting a sequence of inter-device messages over a data transmission circuit forming a path through a switching network between first and second sites, the switching network employing at least one data transmission protocol, the inter-device messages containing data relating to transmission
10 performance of the data transmission circuit, the method comprising the steps of:

(a) forming each of the inter-device messages in the sequence in accordance with an inter-device message transmission structure, such that each of the inter-device messages is capable of being encapsulated in a single protocol data unit of said at least one data transmission protocol of the switching network, wherein each
15 inter-device message includes: a message type field identifying the inter-device message as a particular message within the sequence; and a count field for transmitting an indicator of a number of protocol data units (PDUs) offered for transmission to the switching network over the data transmission circuit, wherein, when the message type field contains a first value, the count field contains a count
20 OFF_A indicating a number of PDUs bound for the second site that are offered by the first site to the switching network over the first transmission circuit, and when the message type field contains a second value, the count field contains a count OFF_B indicating a number of PDUs bound for the first site that are offered by the second site to the switching network over the first transmission circuit;

25 (b) inserting each of the inter-device messages into data traffic to be transmitted on the data transmission circuit between the first and second sites; and

(c) extracting the inter-device messages from the data traffic transmitted on the data transmission circuit.

86. The method of claim 85, wherein the switching network is an interworked
30 switching network employing plural, different data transmission protocols, and wherein step (a) includes forming each of the inter-device messages, such that each

of the inter-device messages is capable of being encapsulated in a single protocol data unit of each of the plural data transmission protocols of the interworked network.

87. The method of claim 86, wherein said plural data transmission protocols
5 includes at least one of: Asynchronous Transfer Mode (ATM), Frame Relay, High Level Data Link Control (HDLC), X.25, and tunneled protocols.

88. The method of claim 86, wherein:
the interworked network includes an asynchronous transfer mode (ATM)
network employing an ATM data transmission protocol and a frame relay network
10 employing a frame relay data transmission protocol;

the first site is a site on one of the ATM network and the frame relay network,
and the second site is a site on the other of the ATM network and the frame relay
network; and

step (a) includes forming each inter-device message to be capable of being
15 encapsulated in a single ATM cell and in a single frame relay frame.

89. The method of claim 85, wherein:
the sequence of data measurement messages includes first and third
messages transmitted over the data transmission circuit from the first site to the
second site in accordance with the inter-device message transmission structure, and
20 second and fourth messages transmitted over the data transmission circuit from the
second site to the first site in accordance with the inter-device message transmission
structure; and

step (a) includes forming the inter-device messages such that: when said
message type field identifies the inter-device message as the third message in the
25 sequence, said count field contains the count OFF_A ; and when said message type
field identifies the inter-device message as the fourth message in the sequence, the
count field contains the count OFF_B .

90. The method of claim 85, wherein step (a) includes forming each of the
inter-device messages to include:

a first timestamp field for transmitting timestamp data indicating a time at which a previous message in the sequence was received; and

a second timestamp field for transmitting timestamp data indicating a time at which a previous message in the sequence was transmitted.

5 91. The method of claim 90, wherein:

the sequence of data measurement messages includes first and third messages transmitted over the data transmission circuit from the first site to the second site in accordance with the inter-device message transmission structure, and second and fourth messages transmitted over the data transmission circuit from the
10 second site to the first site in accordance with the inter-device message transmission structure; and

step (a) includes forming the inter-device messages such that: when said message type field identifies the inter-device message as the second message in the sequence, said first timestamp field indicates a receive time of the first message;
15 when said message type field identifies the inter-device message as the third message in the sequence, said first timestamp field indicates a receive time of the second message, said second timestamp field indicates a transmit time of the first message, and said count field contains the count OFF_A ; and when said message type field identifies the inter-device message as the fourth message in the
20 sequence, said second timestamp field indicates a transmit time of the second message, and the count field contains the count OFF_B .




FIG.1




FIG.2

3 / 5





FIG. 3A

4 / 5

**FIG.3B**

Byte 0	Byte 1	Byte 2	Byte 3
Version #	Message Type	Sequence #	
SlapUpTime(s)			
Last RX timestamp (T_2 in M_2, T_4 in M_3)			
Last TX timestamp (T_1 in M_3, T_3 in M_4)			
Non-burst offered PDUs (M_3, M_4)			

FIG.4