

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
27 December 2007 (27.12.2007)

PCT

(10) International Publication Number
WO 2007/149165 A2

(51) International Patent Classification:
H04J 1/16 (2006.01)

(21) International Application Number:
PCT/US2007/011650

(22) International Filing Date: 15 May 2007 (15.05.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/454,206 16 June 2006 (16.06.2006) US

(71) Applicant (for all designated States except US): **HARRIS CORPORATION** [US/US]; 1025 W. Nasa Blvd., Ms A-11i, Melbourne, FL 32919 (US).

(72) Inventor: **SMITH, Donald, L.**; 584 Hawksbill Island Drive, Satellite Beach, FL 32937 (US).

(74) Agents: **YATSKO, Michael, S.** et al.; Blum II, Ronald S. Of Harris Corporation, 1025 W. Nasa Blvd., Ms A-11i, Melbourne, FL 32919 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND SYSTEM FOR OUTBOUND CONTENT- BASED QOS



(57) Abstract: Certain embodiments of the present invention provide a method (500) for communicating inbound network data to provide QoS. The method (500) includes receiving data from an application, prioritizing the data by assigning a priority to the data, and communicating the data over a network based at least in part on the priority of the data. The data priority is based at least in part on message content. Certain embodiments of the present invention provide a system (400) for communicating inbound networking data to provide QoS. The system (400) includes a data prioritization component (460) adapted to prioritize data by assigning a priority to the data and a data communications component (470) adapted to receive the data from an application and to communicate the data over a network based at least in part on the priority of the data. The data priority is based at least in part on message content.

METHOD AND SYSTEM FOR OUTBOUND CONTENT-BASED QoS

The present invention generally relates to communications networks. More particularly, the present invention relates to systems and methods for outbound content-based Quality of Service.

Communications networks are utilized in a variety of environments. Communications networks typically include two or more nodes connected by one or more links. Generally, a communications network is used to support communication between two or more participant nodes over the links and intermediate nodes in the communications network. There may be many kinds of nodes in the network. For example, a network may include nodes such as clients, servers, workstations, switches, and/or routers. Links may be, for example, modem connections over phone lines, wires, Ethernet links, Asynchronous Transfer Mode (ATM) circuits, satellite links, and/or fiber optic cables.

A communications network may actually be composed of one or more smaller communications networks. For example, the Internet is often described as network of interconnected computer networks. Each network may utilize a different architecture and/or topology. For example, one network may be a switched Ethernet network with a star topology and another network may be a Fiber-Distributed Data Interface (FDDI) ring.

Communications networks may carry a wide variety of data. For example, a network may carry bulk file transfers alongside data for interactive real-time conversations. The data sent on a network is often sent in packets, cells, or frames. Alternatively, data may be sent as a stream. In some instances, a stream or flow of data may actually be a sequence of packets. Networks such as the Internet provide general purpose data paths between a range of nodes and carrying a vast array of data with different requirements.

Communication over a network typically involves multiple levels of communication protocols. A protocol stack, also referred to as a networking stack or protocol suite, refers to a collection of protocols used for communication.

5 Each protocol may be focused on a particular type of capability or form of communication. For example, one protocol may be concerned with the electrical signals needed to communicate with devices connected by a copper wire. Other protocols may address ordering and reliable transmission
10 between two nodes separated by many intermediate nodes, for example.

Protocols in a protocol stack typically exist in a hierarchy. Often, protocols are classified into layers. One reference model for protocol layers is the Open Systems
15 Interconnection (OSI) model. The OSI reference model includes seven layers: a physical layer, data link layer, network layer, transport layer, session layer, presentation layer, and application layer. The physical layer is the "lowest" layer, while the application layer is the "highest" layer. Two well-
20 known transport layer protocols are the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). A well known network layer protocol is the Internet Protocol (IP).

At the transmitting node, data to be transmitted is passed down the layers of the protocol stack, from highest to
25 lowest. Conversely, at the receiving node, the data is passed up the layers, from lowest to highest. At each layer, the data may be manipulated by the protocol handling communication at that layer. For example, a transport layer protocol may add a header to the data that allows for ordering of packets
30 upon arrival at a destination node. Depending on the application, some layers may not be used, or even present, and data may just be passed through.

One kind of communications network is a tactical data network. A tactical data network may also be referred to

as a tactical communications network. A tactical data network may be utilized by units within an organization such as a military (e.g., army, navy, and/or air force). Nodes within a tactical data network may include, for example, individual
5 soldiers, aircraft, command units, satellites, and/or radios. A tactical data network may be used for communicating data such as voice, position telemetry, sensor data, and/or real-time video.

An example of how a tactical data network may be
10 employed is as follows. A logistics convoy may be in-route to provide supplies for a combat unit in the field. Both the convoy and the combat unit may be providing position telemetry to a command post over satellite radio links. An unmanned
15 aerial vehicle (UAV) may be patrolling along the road the convoy is taking and transmitting real-time video data to the command post over a satellite radio link also. At the command post, an analyst may be examining the video data while a controller is tasking the UAV to provide video for a specific
20 section of road. The analyst may then spot an improvised explosive device (IED) that the convoy is approaching and send out an order over a direct radio link to the convoy for it to halt and alerting the convoy to the presence of the IED.

The various networks that may exist within a tactical data network may have many different architectures
25 and characteristics. For example, a network in a command unit may include a gigabit Ethernet local area network (LAN) along with radio links to satellites and field units that operate with much lower throughput and higher latency. Field units may communicate both via satellite and via direct path radio
30 frequency (RF). Data may be sent point-to-point, multicast, or broadcast, depending on the nature of the data and/or the specific physical characteristics of the network. A network may include radios, for example, set up to relay data. In addition, a network may include a high frequency (HF) network

which allows long rang communication. A microwave network may also be used, for example. Due to the diversity of the types of links and nodes, among other reasons, tactical networks often have overly complex network addressing schemes and
5 routing tables. In addition, some networks, such as radio-based networks, may operate using bursts. That is, rather than continuously transmitting data, they send periodic bursts of data. This is useful because the radios are broadcasting on a particular channel that must be shared by all
10 participants, and only one radio may transmit at a time.

Tactical data networks are generally bandwidth-constrained. That is, there is typically more data to be communicated than bandwidth available at any given point in time. These constraints may be due to either the demand for
15 bandwidth exceeding the supply, and/or the available communications technology not supplying enough bandwidth to meet the user's needs, for example. For example, between some nodes, bandwidth may be on the order of kilobits/sec. In bandwidth-constrained tactical data networks, less important
20 data can clog the network, preventing more important data from getting through in a timely fashion, or even arriving at a receiving node at all. In addition, portions of the networks may include internal buffering to compensate for unreliable links. This may cause additional delays. Further, when the
25 buffers get full, data may be dropped.

In many instances the bandwidth available to a network cannot be increased. For example, the bandwidth available over a satellite communications link may be fixed and cannot effectively be increased without deploying another
30 satellite. In these situations, bandwidth must be managed rather than simply expanded to handle demand. In large systems, network bandwidth is a critical resource. It is desirable for applications to utilize bandwidth as efficiently as possible. In addition, it is desirable that applications

avoid "clogging the pipe," that is, overwhelming links with data, when bandwidth is limited. When bandwidth allocation changes, applications should preferably react. Bandwidth can change dynamically due to, for example, quality of service, jamming, signal obstruction, priority reallocation, and line-of-sight. Networks can be highly volatile and available bandwidth can change dramatically and without notice.

In addition to bandwidth constraints, tactical data networks may experience high latency. For example, a network involving communication over a satellite link may incur latency on the order of half a second or more. For some communications this may not be a problem, but for others, such as real-time, interactive communication (e.g., voice communications), it is highly desirable to minimize latency as much as possible.

Another characteristic common to many tactical data networks is data loss. Data may be lost due to a variety of reasons. For example, a node with data to send may be damaged or destroyed. As another example, a destination node may temporarily drop off of the network. This may occur because, for example, the node has moved out of range, the communication's link is obstructed, and/or the node is being jammed. Data may be lost because the destination node is not able to receive it and intermediate nodes lack sufficient capacity to buffer the data until the destination node becomes available. Additionally, intermediate nodes may not buffer the data at all, instead leaving it to the sending node to determine if the data ever actually arrived at the destination.

Often, applications in a tactical data network are unaware of and/or do not account for the particular characteristics of the network. For example, an application may simply assume it has as much bandwidth available to it as it needs. As another example, an application may assume that

data will not be lost in the network. Applications which do not take into consideration the specific characteristics of the underlying communications network may behave in ways that actually exacerbate problems. For example, an application may
5 continuously send a stream of data that could just as effectively be sent less frequently in larger bundles. The continuous stream may incur much greater overhead in, for example, a broadcast radio network that effectively starves other nodes from communicating, whereas less frequent bursts
10 would allow the shared bandwidth to be used more effectively.

Certain protocols do not work well over tactical data networks. For example, a protocol such as TCP may not function well over a radio-based tactical network because of the high loss rates and latency such a network may encounter.
15 TCP requires several forms of handshaking and acknowledgments to occur in order to send data. High latency and loss may result in TCP hitting time outs and not being able to send much, if any, meaningful data over such a network.

Information communicated with a tactical data
20 network often has various levels of priority with respect to other data in the network. For example, threat warning receivers in an aircraft may have higher priority than position telemetry information for troops on the ground miles away. As another example, orders from headquarters regarding
25 engagement may have higher priority than logistical communications behind friendly lines. The priority level may depend on the particular situation of the sender and/or receiver. For example, position telemetry data may be of much higher priority when a unit is actively engaged in combat as
30 compared to when the unit is merely following a standard patrol route. Similarly, real-time video data from an UAV may have higher priority when it is over the target area as opposed to when it is merely in-route.

There are several approaches to delivering data over a network. One approach, used by many communications networks, is a "best effort" approach. That is, data being communicated will be handled as well as the network can, given other demands, with regard to capacity, latency, reliability,
5 ordering, and errors. Thus, the network provides no guarantees that any given piece of data will reach its destination in a timely manner, or at all. Additionally, no guarantees are made that data will arrive in the order sent or
10 even without transmission errors changing one or more bits in the data.

Another approach is Quality of Service (QoS). QoS refers to one or more capabilities of a network to provide various forms of guarantees with regard to data that is
15 carried. For example, a network supporting QoS may guarantee a certain amount of bandwidth to a data stream. As another example, a network may guarantee that packets between two particular nodes have some maximum latency. Such a guarantee may be useful in the case of a voice communication where the
20 two nodes are two people having a conversation over the network. Delays in data delivery in such a case may result in irritating gaps in communication and/or dead silence, for example.

QoS may be viewed as the capability of a network to
25 provide better service to selected network traffic. The primary goal of QoS is to provide priority including dedicated bandwidth, controlled jitter and latency (required by some real-time and interactive traffic), and improved loss characteristics. Another important goal is making sure that
30 providing priority for one flow does not make other flows fail. That is, guarantees made for subsequent flows must not break the guarantees made to existing flows.

Current approaches to QoS often require every node in a network to support QoS, or, at the very least, for every

node in the network involved in a particular communication to support QoS. For example, in current systems, in order to provide a latency guarantee between two nodes, every node carrying the traffic between those two nodes must be aware of
5 and agree to honor, and be capable of honoring, the guarantee.

There are several approaches to providing QoS. One approach is Integrated Services, or "IntServ." IntServ provides a QoS system wherein every node in the network supports the services and those services are reserved when a
10 connection is set up. IntServ does not scale well because of the large amount of state information that must be maintained at every node and the overhead associated with setting up such connections.

Another approach to providing QoS is Differentiated
15 Services, or "DiffServ." DiffServ is a class of service model that enhances the best-effort services of a network such as the Internet. DiffServ differentiates traffic by user, service requirements, and other criteria. Then, DiffServ marks packets so that network nodes can provide different
20 levels of service via priority queuing or bandwidth allocation, or by choosing dedicated routes for specific traffic flows. Typically, a node has a variety of queues for each class of service. The node then selects the next packet to send from those queues based on the class categories.

Existing QoS solutions are often network specific
25 and each network type or architecture may require a different QoS configuration. Due to the mechanisms existing QoS solutions utilize, messages that look the same to current QoS systems may actually have different priorities based on
30 message content. However, data consumers may require access to high-priority data without being flooded by lower-priority data. Existing QoS systems cannot provide QoS based on message content at the transport layer.

As mentioned, existing QoS solutions require at least the nodes involved in a particular communication to support QoS. However, the nodes at the "edge" of network may be adapted to provide some improvement in QoS, even if they
5 are incapable of making total guarantees. Nodes are considered to be at the edge of the network if they are the participating nodes in a communication (i.e., the transmitting and/or receiving nodes) and/or if they are located at
chokepoints in the network. A chokepoint is a section of the
10 network where all traffic must pass to another portion. For example, a router or gateway from a LAN to a satellite link would be a choke point, since all traffic from the LAN to any nodes not on the LAN must pass through the gateway to the satellite link.

15 Thus, there is a need for systems and methods providing QoS in a tactical data network. There is a need for systems and methods for providing QoS on the edge of a tactical data network. Additionally, there is a need for adaptive, configurable QoS systems and methods in a tactical
20 data network.

Certain embodiments of the present invention provide a method for communicating inbound network data to provide QoS. The method includes receiving data from an application, prioritizing the data by assigning a priority to the data, and
25 communicating the data over a network based at least in part on the priority of the data. The priority of the data is based at least in part on message content.

Certain embodiments of the present invention provide a system for communicating inbound networking data to provide
30 QoS. The system includes a data prioritization component adapted to prioritize data by assigning a priority to the data and a data communications component adapted to receive the data from an application and to communicate the data over a network based at least in part on the priority of the data.

The priority of the data is based at least in part on message content.

Certain embodiments of the present invention provide a computer-readable medium. The computer-readable medium
5 includes a set of instructions for execution on a computer. The set of instructions includes a data prioritization routine configured to prioritize data by assigning a priority to the data and a data communications routine configured to receive the data from an application and to communicate the data over
10 a network based at least in part on the priority of the data. The priority of the data is based at least in part on message content.

Fig. 1 illustrates a tactical communications network environment operating with an embodiment of the present
15 invention.

Fig. 2 shows the positioning of the data communications system in the seven layer OSI network model in accordance with an embodiment of the present invention.

Fig. 3 depicts an example of multiple networks
20 facilitated using the data communications system in accordance with an embodiment of the present invention.

Fig. 4 illustrates a data communications environment operating according to an embodiment of the present invention.

Fig. 5 illustrates a flow chart for a method of
25 communicating data according to an embodiment of the present invention.

The foregoing summary, as well as the following detailed description of certain embodiments of the present invention, will be better understood when read in conjunction
30 with the appended drawings. For the purpose of illustrating the invention, certain embodiments are shown in the drawings. It should be understood, however, that the present invention is not limited to the arrangements and instrumentality shown in the attached drawings.

Fig. 1 illustrates a tactical communications network environment 100 operating with an embodiment of the present invention. The network environment 100 includes a plurality of communication nodes 110, one or more networks 120, one or more links 130 connecting the nodes and network(s), and one or more communication systems 150 facilitating communication over the components of the network environment 100. The following discussion assumes a network environment 100 including more than one network 120 and more than one link 130, but it should be understood that other environments are possible and anticipated.

Communication nodes 110 may be and/or include radios, transmitters, satellites, receivers, workstations, servers, and/or other computing or processing devices, for example.

Network(s) 120 may be hardware and/or software for transmitting data between nodes 110, for example. Network(s) 120 may include one or more nodes 110, for example.

Link(s) 130 may be wired and/or wireless connections to allow transmissions between nodes 110 and/or network(s) 120.

The communications system 150 may include software, firmware, and/or hardware used to facilitate data transmission among the nodes 110, networks 120, and links 130, for example. As illustrated in Fig. 1, communications system 150 may be implemented with respect to the nodes 110, network(s) 120, and/or links 130. In certain embodiments, every node 110 includes a communications system 150. In certain embodiments, one or more nodes 110 include a communications system 150. In certain embodiments, one or more nodes 110 may not include a communications system 150.

The communication system 150 provides dynamic management of data to help assure communications on a tactical communications network, such as the network environment 100.

As shown in Fig. 2, in certain embodiments, the system 150 operates as part of and/or at the top of the transport layer in the OSI seven layer protocol model. The system 150 may give precedence to higher priority data in the tactical
5 network passed to the transport layer, for example. The system 150 may be used to facilitate communications in a single network, such as a local area network (LAN) or wide area network (WAN), or across multiple networks. An example of a multiple network system is shown in Fig. 3. The system
10 150 may be used to manage available bandwidth rather than add additional bandwidth to the network, for example.

In certain embodiments, the system 150 is a software system, although the system 150 may include both hardware and software components in various embodiments. The system 150
15 may be network hardware independent, for example. That is, the system 150 may be adapted to function on a variety of hardware and software platforms. In certain embodiments, the system 150 operates on the edge of the network rather than on nodes in the interior of the network. However, the system 150
20 may operate in the interior of the network as well, such as at "choke points" in the network.

The system 150 may use rules and modes or profiles to perform throughput management functions such as optimizing available bandwidth, setting information priority, and
25 managing data links in the network. By "optimizing" bandwidth, it is meant that the presently described technology can be employed to increase an efficiency of bandwidth use to communicate data in one or more networks. Optimizing bandwidth usage may include removing functionally redundant
30 messages, message stream management or sequencing, and message compression, for example. Setting information priority may include differentiating message types at a finer granularity than Internet Protocol (IP) based techniques and sequencing messages onto a data stream via a selected rule-based

sequencing algorithm, for example. Data link management may include rule-based analysis of network measurements to affect changes in rules, modes, and/or data transports, for example. A mode or profile may include a set of rules related to the operational needs for a particular network state of health or condition. The system 150 provides dynamic, "on-the-fly" reconfiguration of modes, including defining and switching to new modes on the fly.

The communication system 150 may be configured to accommodate changing priorities and grades of service, for example, in a volatile, bandwidth-limited network. The system 150 may be configured to manage information for improved data flow to help increase response capabilities in the network and reduce communications latency. Additionally, the system 150 may provide interoperability via a flexible architecture that is upgradeable and scalable to improve availability, survivability, and reliability of communications. The system 150 supports a data communications architecture that may be autonomously adaptable to dynamically changing environments while using predefined and predictable system resources and bandwidth, for example.

In certain embodiments, the system 150 provides throughput management to bandwidth-constrained tactical communications networks while remaining transparent to applications using the network. The system 150 provides throughput management across multiple users and environments at reduced complexity to the network. As mentioned above, in certain embodiments, the system 150 runs on a host node in and/or at the top of layer four (the transport layer) of the OSI seven layer model and does not require specialized network hardware. The system 150 may operate transparently to the layer four interface. That is, an application may utilize a standard interface for the transport layer and be unaware of the operation of the system 150. For example, when an

application opens a socket, the system 150 may filter data at this point in the protocol stack. The system 150 achieves transparency by allowing applications to use, for example, the TCP/IP socket interface that is provided by an operating
5 system at a communication device on the network rather than an interface specific to the system 150. System 150 rules may be written in extensible markup language (XML) and/or provided via custom dynamic link libraries (DLLs), for example.

In certain embodiments, the system 150 provides
10 quality of service (QoS) on the edge of the network. The system's QoS capability offers content-based, rule-based data prioritization on the edge of the network, for example. Prioritization may include differentiation and/or sequencing, for example. The system 150 may differentiate messages into
15 queues based on user-configurable differentiation rules, for example. The messages are sequenced into a data stream in an order dictated by the user-configured sequencing rule (e.g., starvation, round robin, relative frequency, etc.). Using QoS on the edge, data messages that are indistinguishable by
20 traditional QoS approaches may be differentiated based on message content, for example. Rules may be implemented in XML, for example. In certain embodiments, to accommodate capabilities beyond XML and/or to support extremely low latency requirements, the system 150 allows dynamic link
25 libraries to be provided with custom code, for example.

Inbound and/or outbound data on the network may be customized via the system 150. Prioritization protects client applications from high-volume, low-priority data, for example. The system 150 helps to ensure that applications receive data
30 to support a particular operational scenario or constraint.

In certain embodiments, when a host is connected to a LAN that includes a router as an interface to a bandwidth-constrained tactical network, the system may operate in a configuration known as QoS by proxy. In this configuration,

packets that are bound for the local LAN bypass the system and immediately go to the LAN. The system applies QoS on the edge of the network to packets bound for the bandwidth-constrained tactical link.

5 In certain embodiments, the system 150 offers dynamic support for multiple operational scenarios and/or network environments via commanded profile switching. A profile may include a name or other identifier that allows the user or system to change to the named profile. A profile may
10 also include one or more identifiers, such as a functional redundancy rule identifier, a differentiation rule identifier, an archival interface identifier, a sequencing rule identifier, a pre-transmit interface identifier, a post-transmit interface identifier, a transport identifier, and/or
15 other identifier, for example. A functional redundancy rule identifier specifies a rule that detects functional redundancy, such as from stale data or substantially similar data, for example. A differentiation rule identifier specifies a rule that differentiates messages into queues for
20 processing, for example. An archival interface identifier specifies an interface to an archival system, for example. A sequencing rule identifier identifies a sequencing algorithm that controls samples of queue fronts and, therefore, the sequencing of the data on the data stream. A pre-transmit
25 interface identifier specifies the interface for pre-transmit processing, which provides for special processing such as encryption and compression, for example. A post-transmit interface identifier identifies an interface for post-transmit processing, which provides for processing such as de-
30 encryption and decompression, for example. A transport identifier specifies a network interface for the selected transport.

A profile may also include other information, such as queue sizing information, for example. Queue sizing

information identifiers a number of queues and amount of memory and secondary storage dedicated to each queue, for example.

In certain embodiments, the system 150 provides a
5 rules-based approach for optimizing bandwidth. For example, the system 150 may employ queue selection rules to differentiate messages into message queues so that messages may be assigned a priority and an appropriate relative frequency on the data stream. The system 150 may use
10 functional redundancy rules to manage functionally redundant messages. A message is functionally redundant if it is not different enough (as defined by the rule) from a previous message that has not yet been sent on the network, for example. That is, if a new message is provided that is not
15 sufficiently different from an older message that has already been scheduled to be sent, but has not yet been sent, the newer message may be dropped, since the older message will carry functionally equivalent information and is further ahead in the queue. In addition, functional redundancy may include
20 actual duplicate messages and newer messages that arrive before an older message has been sent. For example, a node may receive identical copies of a particular message due to characteristics of the underlying network, such as a message that was sent by two different paths for fault tolerance
25 reasons. As another example, a new message may contain data that supersedes an older message that has not yet been sent. In this situation, the system 150 may drop the older message and send only the new message. The system 150 may also include priority sequencing rules to determine a priority-
30 based message sequence of the data stream. Additionally, the system 150 may include transmission processing rules to provide pre-transmission and post-transmission special processing, such as compression and/or encryption.

In certain embodiments, the system 150 provides fault tolerance capability to help protect data integrity and reliability. For example, the system 150 may use user-defined queue selection rules to differentiate messages into queues.

5 The queues are sized according to a user-defined configuration, for example. The configuration specifies a maximum amount of memory a queue may consume, for example. Additionally, the configuration may allow the user to specify a location and amount of secondary storage that may be used

10 for queue overflow. After the memory in the queues is filled, messages may be queued in secondary storage. When the secondary storage is also full, the system 150 may remove the oldest message in the queue, logs an error message, and queues the newest message. If archiving is enabled for the

15 operational mode, then the de-queued message may be archived with an indicator that the message was not sent on the network.

Memory and secondary storage for queues in the system 150 may be configured on a per-link basis for a

20 specific application, for example. A longer time between periods of network availability may correspond to more memory and secondary storage to support network outages. The system 150 may be integrated with network modeling and simulation applications, for example, to help identify sizing to help

25 ensure that queues are sized appropriately and time between outages is sufficient to help achieve steady-state and help avoid eventual queue overflow.

Furthermore, in certain embodiments, the system 150 offers the capability to meter inbound ("shaping") and

30 outbound ("policing") data. Policing and shaping capabilities help address mismatches in timing in the network. Shaping helps to prevent network buffers from flooding with high-priority data queued up behind lower-priority data. Policing helps to prevent application data consumers from being overrun

by low-priority data. Policing and shaping are governed by two parameters: effective link speed and link proportion. The system 150 may form a data stream that is no more than the effective link speed multiplied by the link proportion, for example. The parameters may be modified dynamically as the network changes. The system may also provide access to detected link speed to support application level decisions on data metering. Information provided by the system 150 may be combined with other network operations information to help decide what link speed is appropriate for a given network scenario.

FIG. 4 illustrates a data communications environment 400 operating according to an embodiment of the present invention. The data communications environment 400 includes one or more nodes 410, one or more networks 420, and one or more links 430 connecting the nodes 410 and the networks 420, and the data communications system 450 facilitating communications over the other components of the data communications environment 400. The data communications environment 400 may be similar to the data communications environment 100 of FIG. 1, as described above.

The data communications system 450 may operate within the node 410, as shown in FIG. 4. Alternatively, the data communications system 450 may operate within the network 420 and/or between the node 410 and the network 420. The node 410 may include one or more applications 415, such as Application A and Application B, as shown in FIG. 4.

The data communications system 450 is adapted to receive, store, organize, prioritize, process, transmit, and/or communicate data. The data received, stored, organized, prioritized, processed, transmitted, and/or communicated by the data communications system 450 may include, for example, a block of data, such as a packet, cell, frame, and/or stream.

In certain embodiments of the present invention, the data communications system 450 may include a data prioritization component 460 and a data communications component 470, which are described below in more detail.

5 The data prioritization component 460 prioritizes data. In certain embodiments of the present invention, the data prioritization component 460 may prioritize data based at least in part on one or more prioritization rules, such as differentiation rules and/or sequencing rules. The
10 prioritization rules may be user defined. The prioritization rules may be written in XML and/or provided in one or more DLLs.

In certain embodiments of the present invention, the data prioritization component 460 may prioritize data based at
15 least in part on message content. For example, the data priority may be based at least in part on data type, such as video, audio, telemetry, and/or position data. As another example, the data priority may be based at least in part on data source. For example, communications from a general may
20 be assigned a higher priority than communications from a lower ranking officer.

In certain embodiments of the present invention, the data prioritization component 460 may prioritize data based at least in part on protocol information, such as source address
25 and/or transport protocol. In certain embodiments of the present invention, the data prioritization component 460 prioritize data based at least in part on mode.

In certain embodiments of the present invention, the data prioritization component 460 may prioritize data by
30 assigning a priority to the data. For example, position data and emitter data for a near threat may be associated with a priority of "HIGH," next to shoot data may be associated with a priority of "MED HIGH," top-ten shoot list data may be associated with a priority of "MED," emitter data for a threat

over one hundred miles away and situational awareness (SA) data from satellite communications (SATCOM) may be associated with a priority of "MED LOW," and general status data may be assigned a priority of "LOW."

5 As described above, data may be assigned and/or associated with a priority. For example, the data priority may include "HIGH," "MED HIGH," "MED," "MED LOW," or "LOW." As another example, the data priority may include "KEEP PILOT ALIVE," "KILL ENEMY," or "INFORMATIONAL."

10 In certain embodiments of the present invention, the data priority may be based at least in part on a type, category, and/or group of data. For example, types of data may include position data, emitter data for a near threat, next to shoot data, top-ten shoot list data, emitter data for
15 a threat over one hundred miles away, SA data from SATCOM, and/or general status data. Additionally, the data may be grouped into categories, such as "KEEP PILOT ALIVE," "KILL ENEMY," and/or "INFORMATIONAL." For example, "KEEP PILOT ALIVE" data, such as position data and emitter data for a near
20 threat, may relate to the health and safety of a pilot. As another example, "KILL ENEMY" data, such as next to shoot data, top-ten shoot list data, and emitter data for a threat over one hundred miles away, may relate to combat systems. As another example, "INFORMATIONAL" data, such as SA data from
25 SATCOM and general status data, may relate to non-combat systems.

 As described above, the data type, category, and/or group may be the same as and/or similar to the data priority. For example, "KEEP PILOT ALIVE" data, such as position data
30 and emitter data for a near threat, may be associated with a priority of "KEEP PILOT ALIVE," which is more important than "KILL ENEMY" data, such as next to shoot data, top-ten shoot list data, and emitter data for a threat over one hundred miles away, associated with a priority of "KILL ENEMY." As

another example, "KILL ENEMY" data, such as next to shoot data, top-ten shoot list data, and emitter data for a threat over one hundred miles away, may be associated with a priority of "KILL ENEMY," which is more important than "INFORMATIONAL" data, such as SA data from SATCOM and general status data, associated with a priority of "INFORMATIONAL."

In certain embodiments of the present invention, the data prioritization component 460 may include a differentiation component 462, a sequencing component 464, and a data organization component 466, which are described below in more detail.

The differentiation component 462 differentiates data. In certain embodiments of the present invention, the differentiation component 462 may differentiate data based at least in part on one or more differentiation rules, such as queue selection rules and/or functional redundancy rules. The differentiation rules may be user defined. The differentiation rules may be written in XML and/or provided in one or more DLLs.

In certain embodiments of the present invention, the differentiation component 462 may add data to the data organization component 466. For example, the differentiation component 462 may add data to the data organization component 466 based at least in part on one or more queue selection rules.

In certain embodiments of the present invention, the differentiation component 462 may remove and/or withhold data from the data organization component 46. For example, the differentiation component 462 may remove data from the data organization component 466 based at least in part on one or more functional redundancy rules.

The sequencing component 464 sequences data. In certain embodiments of the present invention, the sequencing component 464 may sequence data based at least in part on one

or more sequencing rules, such as such as starvation, round robin, and relative frequency. The sequencing rules may be user defined. The sequencing rules may be written in XML and/or provided in one or more DLLs.

5 In certain embodiments of the present invention, the sequencing component 464 may select and/or remove data from the data organization component 466. For example, the sequencing component 464 may remove data from the data organization component 46 based at least in part on the
10 sequencing rules.

 The data organization component 466 stores and/or organizes data. In certain embodiments of the present invention, the data organization component 466 may store and/or organize the data based at least in part on priority,
15 such as "KEEP PILOT ALIVE," "KILL ENEMY," and "INFORMATIONAL."

 In certain embodiments of the present invention, the data organization component 466 may include, for example, one or more queues, such as Q1, Q2, Q3, Q4, and Q5. For example, data associated with a priority of "HIGH" may be stored in Q1,
20 data associated with a priority of "MED HIGH" may be stored in Q2, data associated with a priority of "MED" may be stored in Q3, data associated with a priority of "MED LOW" may be stored in Q4, and data associated with a priority of "LOW" may be stored in Q5. Alternatively, the data organization component
25 466 may include, for example, one or more trees, tables, linked lists, and/or other data structures for storing and/or organizing data.

 The data communications component 470 communicates data. In certain embodiments of the present invention, the
30 data communications component 470 receives data, for example, from a node 410 and/or an application 415 running on the node 410, or over a network 420 and/or a link 430 connecting the node 410 to the network 420. In certain embodiments of the present invention, the data communications component 470

transmits data, for example, to a node 410 and/or an application 415 running on the node 410, or over a network 420 and/or a link connecting the node 410 to the network 420.

In certain embodiments of the present invention, the data communications component 470 communicates with the data prioritization component 460. More particularly, the data communications component 470 transmits data to the differentiation component 462 and receives data from the sequencing component 464. Alternatively, the data communications component 470 may communicate with the data organization component 466.

In certain embodiments of the present invention, the data prioritization component 460 may perform one or more of the functions of the data communications component 470.

In certain embodiments of the present invention, the data communications component 470 may communicate data based at least in part on data priority.

In operation, for example, data is received from one or more applications 415 by the data communications component 470. The received data is prioritized by the data prioritization component 460 based at least in part on message content and/or mode. The prioritized data is transmitted over a network 420 by the data communications component 470.

In certain embodiments of the present invention, the data communication system 450 may not receive all of the data. For example, some of the data may be stored in a buffer and the data communication system 450 may receive only header information and a pointer to the buffer. As another example, the data communication system 450 may be hooked into the protocol stack of an operating system and when an application passes data to the operating system through a transport layer interface (e.g., sockets), the operating system may then provide access to the data to the data communication system 450.

In certain embodiments of the present invention, the data communications system 450 may not drop data. That is, although the data may be lower priority, it is not dropped by the data communications system 450. Rather, the data may be
5 delayed for a period of time, potentially dependent on the amount of higher priority data that is received.

In certain embodiments of the present invention, the data communications system 450 is transparent to other applications. For example, the processing, organizing, and/or
10 prioritization performed by the data communications system 450 may be transparent to one or more nodes 410 or other applications or data sources. As another example, an application 415 running on the same system as the data communications system 450, or on a node 410 connected to the
15 data communications system 450, may be unaware of the prioritization of data performed by the data communications system 450.

In certain embodiments of the present invention, the data communications system 450 may provide QoS.

20 As discussed above, the components, elements, and/or functionality of the data communication system 450 may be implemented alone or in combination in various forms in hardware, firmware, and/or as a set of instructions in software, for example. Certain embodiments may be provided as
25 a set of instructions residing on a computer-readable medium, such as a memory, hard disk, DVD, or CD, for execution on a general purpose computer or other processing device.

FIG. 5 illustrates a flow diagram of a method 500 for communicating data according to an embodiment of the
30 present invention. The method 500 includes the following steps, which will be described below in more detail. At step 510, data is received. At step 520, the data is prioritized. At step 530, the data is communicated. The method 500 is described with reference to elements of the data

communications environment 400 of FIG. 4, but it should be understood that other implementations are possible.

At step 510, the data is received. The data may be received, for example, by the data communications system 450, as described above. As another example, the data may be received from a node 410 and/or an application 415 running on the node 410. As another example, the data may be received, for example, over a network 420 and/or a link connecting the node 410 and the network 420.

In certain embodiments of the present invention, the data may be received from an application 415.

At step 520, the data is prioritized. The data prioritized may be the data received at step 510, for example. The data may be prioritized, for example, by the data communications system 450 of FIG. 4, as described above. As another example, the data may be prioritized by the data prioritization component 460 of the data communications system 450 based at least in part on data prioritization rules.

In certain embodiments of the present invention, the data may be prioritized based at least in part on one or more prioritization rules. In certain embodiments of the present invention, the data may be prioritized based at least in part on message content. In certain embodiments of the present invention, the priority of the data may be based at least in part on mode. In certain embodiments of the present invention, the data may be prioritized based at least in part on protocol information. In certain embodiments of the present invention, the data may be prioritized by assigning a priority to the data.

At step 530, the data is communicated. The data communicated may be the data received at step 510, for example. The data communicated may be the data prioritized at step 520, for example. The data may be communicated, for example, by the data communications system 450, as described

above. As another example, the data may be communicated to a node 410 and/or an application 415 running on the node 410. As another example, the data may be communicated over a network 420 and/or a link connecting the node 410 and the
5 network 420.

In certain embodiments of the present invention, the data may be communicated over a network 420 based at least in part on the priority of the data. The priority of the data may be the data priority determined at step 520, for example.

10 One or more of the steps of the method 500 may be implemented alone or in combination in hardware, firmware, and/or as a set of instructions in software, for example. Certain embodiments may be provided as a set of instructions residing on a computer-readable medium, such as a memory, hard
15 disk, DVD, or CD, for execution on a general purpose computer or other processing device.

Certain embodiments of the present invention may omit one or more of these steps and/or perform the steps in a different order than the order listed. For example, some
20 steps may not be performed in certain embodiments of the present invention. As a further example, certain steps may be performed in a different temporal order, including simultaneously, than listed above.

In one embodiment of the present invention, a method
25 for communicating inbound network data to provide QoS includes receiving data from an application, prioritizing the data by assigning a priority to the data, and communicating the data over a network based at least in part on the priority of the data. The priority of the data is based at least in part on
30 message content.

In one embodiment of the present invention, a system for communicating inbound networking data to provide QoS includes a data prioritization component adapted to prioritize data by assigning a priority to the data and a data

communications component adapted to receive the data from an application and to communicate the data over the network based at least in part on the priority of the data. The priority of the data is based at least in part on message content.

5 In one embodiment of the present invention, a computer-readable medium includes a set of instructions for execution on a computer. The set of instructions includes a data prioritization routine configured to prioritize data by assigning a priority to the data and a data communications
10 routine configured to receive the data from an application and to communicate the data over a network based at least in part on the priority of the data. The priority of the data is based at least in part on message content.

Certain embodiments of the present invention provide
15 a method for inbound content-based QoS. The method includes receiving TCP and/or UDP addressed network data over a network, prioritizing the network data using a priority algorithm, processing the network data for redundancy, and using an extraction algorithm to forward the network data to
20 one or more applications. The extraction algorithm may be similar to the sequencing algorithm, as described above. For example, the extraction algorithm may be based at least in part on starvation, relative frequency, or a combination of starvation and relative frequency. Starvation refers to
25 servicing the highest priority queue, unless it is empty, and then servicing lower priority queues. Starvation may be advantageous because the highest priority data never waits for lower priority data. However, starvation may be
disadvantageous because if there is enough of the highest
30 priority data, lower priority queues will never be serviced. Relative frequency is similar to starvation, except that there is a cap on the number of times that a queue gets serviced before the next queue is to be serviced. Relative frequency may be advantageous because all of the queues are serviced

However, relative frequency may be disadvantageous because the highest priority data may sometimes wait for lower priority data. A combination of starvation and relative frequency allows a user to select a subset of queues to be processed via
5 starvation and another subset of queues to be processed via relative frequency. In certain embodiments, the extraction algorithm may be configured by a user.

Thus, certain embodiments of the present invention provide systems and methods for outbound content-based QoS.
10 Certain embodiments provide a technical effect of outbound content-based QoS.

CLAIMS

1. A method for communicating outbound network data to provide quality of service, the method including:
5 receiving data from an application;
prioritizing the data by assigning a priority to the data, wherein the priority of the data is based at least in part on message content; and
communicating the data over a network based at least in
10 part on the priority of the data.
2. The method of claim 1, wherein the data is prioritized based at least in part on a user defined rule.
- 15 3. The method of claim 1, wherein the prioritizing step includes differentiating the data.
4. The method of claim 1, wherein the prioritizing step includes sequencing the data.
20
5. The method of claim 1, wherein the prioritizing step is transparent to an application program.
- 25 6. The method of claim 1, wherein the data is prioritized to provide quality of service.
7. A system for communicating outbound networking data to provide quality of service, the system including:
a data prioritization component adapted to prioritize
30 data by assigning a priority to the data, wherein the priority of the data is based at least in part on message content; and
a data communications component adapted to receive the data from an application and to communicate the data over a network based at least in part on the priority of the data.


8. The system of claim 7, further including a differentiation component adapted to differentiate the data.

5 9. The system of claim 7, further including a sequencing component adapted to sequence the data.

10 10. The system of claim 7, wherein the data prioritization component includes a data organization component adapted to organize the data based at least in part on the priority of the data.


1/5

FIG. 1




2/5

FIG. 2




3/5

FIG. 3



4/6

FIG. 4



5/5

FIG. 5

