Quality of Service for Ad Hoc Networks

MeshNetworks Enabled Architecture (MEA™) networks are designed to meet the data transport requirements for a wide array of services, including bulk data, priority data, video data, and interactive voice streams. The MEA system combines industry standard traffic classification and proprietary user prioritization systems with unique queuing and media access mechanisms. This creates a distributed Quality of Service solution that makes the most effective and efficient use of the bandwidth provided by the wireless mesh. MEA technolgy puts intelligent traffic handling into every node of the network, including mobile client devices. The result is a highly scalable and manageable solution for meeting data stream requirements across the entire mesh network.

Introduction

Packet-switched data networks are increasingly being pressed into service as transports for data, voice and video data streams that had once been the province of circuit-switched networks. These data streams may place widely variant but strict requirements for reliability, bandwidth, and latency on the transport. In some networks, these requirements vary not only with the type of data being transported, but also with the particular sender or receiver of the data. In designing a network to meet these requirements, it might be appropriate to simply build the data network to meet the strictest requirements for all traffic at the anticipated peak loading, regardless of the specific requirements of the individual data streams. This brute force approach is not always effective, particularly in mobile wireless networks where the power and spectrum needed to implement a sufficiently rapid and reliable network are limited and costly. Furthermore, a brute force approach will not fail gracefully if or when the demand exceeds the capacity for which the system was designed.

The need is apparent for a system capable of providing different qualities of service for different types of data by prioritizing, pruning, and shaping the data flows to best meet the individual streams' requirements. The system must maintain service for the most important streams, even when the demands exceed the network's capacity. To accomplish this, data packets must first be classified so that the network components can understand how to treat each packet. Then, the network must provide mechanisms to differentiate service among the different classes of data.

Quality of Service Solutions

Traditional solutions to this problem, known collectively as Quality of Service (QoS) solutions, use packet tagging to classify packets. The tagging can be done by the applications themselves or by application-aware gateways. Some systems may also use prioritized lists of nodes or ports to classify packets based on source of transmission or destination. They then differentiate streams by regulating packet flows into and out of a single point of control, such as a router, switch, base station, or wireless access point.

Traffic between nodes on a single segment of the network is generally not regulated.

Virtually any QoS solution should support leveraging of the concepts and protocols behind traditional packet classification systems. However, the traditional service differentiation mechanism does not fit mobile ad hoc networks. Ad hoc networks are designed to bypass points of contention or failure and to provide efficient peer-to-peer connectivity, so they lack a single point of control through which traffic can be managed. Data packets between peers in an ad hoc network, like packets between nodes on the same segment of a switch, may never be passed through the infrastructure devices where QoS policies would ordinarily be enforced. In addition, the lack of a central point of control makes traditional time-division access mechanisms, as used by Time Division Multiple Access (TDMA) and Global System for Mobile Communications (GSM) systems, impractical. Regulation of peer-to-peer data is important for ad hoc wireless networks because the peer nodes might contend with the infrastructure and other nearby nodes for access to the shared spectrum and thus interfere with any higher-priority data.

To address these issues, the MeshNetworks Enabled Architecture (MEA™) mobile ad hoc networking system incorporates support for industry-standard coding of packet classes as well as a proprietary per node priority system. Utilizing these classifications, each node in the MEA network is aware of exactly how each packet should be treated relative to other packets in the system, both within a node and across the entire network. The classifications control both the behaviors of an adaptive media access mechanism and a weighted queuing mechanism in every node to intelligently meet the demands of network traffic. These mechanisms are uniquely adapted for use on a dynamic mesh network.

The MEA QoS system can be augmented with standard QoS-aware mechanisms and application gateways in the core wired network's switches and routers. This works to enable the management of traffic on the wired portion of the network, although those services are not strictly required to differentiate qualities of service in the mesh.

Classification of Data Packets

The MEA system uses two methods to classify data packets in order to provide them the requested levels of service. The first is the use of the industry standard Differentiated Services Code Point (DSCP). The second is a proprietary system based on the assigned priority of the client node itself.

Differentiated Services Code Point

The DSCP is defined in the Internet Engineering Task Force Request for Comment (RFC) 2474 as a part of the Differentiated Services (DS, or "DiffServe") field. This field is defined for both the IPv4 and IPv6 headers. The DSCP is a dimensionless number used to index a particular class of traffic and thereby a particular transport behavior at each hop. These behaviors are known as Per-Hop Behaviors (PHBs).

Some basic code points are defined in RFC 2474 for backwards compatibility with the overlapping "Precedence" field in the IPv4 header, known collectively as Class Selector Code Points. Extensions to these basic code points, called Assured Forwarding and Expedited Forwarding classes, are defined in RFC 2597 and RFC 2598. Other groups of code points are reserved for experimental or site-local use.

Most often, the indexed per-hop behaviors are described (and sometimes, implemented) simply in terms of queue prioritization of packets at each hop. However, the behaviors may incorporate other features such as traffic pruning to condition the traffic flowing from the node.

For example, voice packets may be assigned a DSCP tag indicating the need for low-latency service. The transport can then intelligently decide to discard stale packets or prioritize new packets to provide the lowest latency possible. This may involve compromising performance at some levels, such as sacrificing raw bandwidth or reliability when the demand exceeds capacity.

High priority data transfers might be tagged to indicate the need for high reliability. The transport can use this knowledge to increase buffering or provide additional retries, perhaps increasing the latency of the stream but serving the data reliably. Likewise, bulk data packets can be assigned "best effort" service so that the transport layer can intersperse them between packets needing lower latency or higher bandwidth.

The MEA system extracts the DS field from IP packets as they enter the wireless mesh, either at the access point or at the client device. The DSCP classification then serves as one of the inputs for selecting an appropriate behavior in the transport mechanisms.

User Priority

The MEA QoS system also augments the standards-based DSCP traffic classification system with a proprietary feature to implement hierarchical service on a per-user basis. While DSCP can be used to classify the traffic by payload, the User Priority feature classifies traffic by the user that originated it, tagging packets with an assigned priority level. Using this tag, data packets from high priority nodes can be granted preferential service over packets from lower priority nodes.

The User Priority at each node is a managed variable with tight control granted to the network administrator or operator. The user is free to select a priority level through applications using the MeshAPI software interface. The operator maintains the control to set the bounds on each node, limiting the user priority to a specific range or setting it to a specific value.

The User Priority feature also underpins a secondary feature known as Emergency Mode. The Emergency Mode feature is also accessed by applications using the MeshAPI and can be used to temporarily grant the user elevated priority for use in emergency situations. However, this ability can also be tightly managed and granted by the operator as needed.

The current User Priority is assigned to every packet transmitted by the user of the node. This priority tag is preserved in each packet as it passes through the mesh until it reaches the final destination or the wired network.

Service Differentiation Mechanisms

The MEA network leverages intelligent and adaptive transport mechanisms in every node in the mesh, including client devices, to provide QoS behaviors. At each hop, the DSCP and User Priority classifications serve as inputs to dictate transport behaviors using these mechanisms. In this way, the mesh nodes cooperate to provide appropriate transport service to packets, even across multiple hops through the mesh.

Access Point Traffic Shaping

The MEA Intelligent Access Point (IAP) queues and shapes traffic as it is relayed into the wireless mesh from the wire in accordance with the DSCP coding in each packet. As many deployment scenarios have the vast majority of packets entering the mesh at the access points, this mechanism is particularly effective for shaping the traffic in a manner that optimizes the use of bandwidth at what is the main point of contention on most networks.

Media Access Queuing, Pruning, and Timing

The MEA transport layer uses a custom implementation of a Weighted Fair Queuing algorithm (WF2Q+) to split bandwidth between different classes of packets. Packets are assigned to priority queues based on both the DSCP and User Priority tags as dictated by a policy map. This queuing algorithm attempts to meet the demands of the highest priority traffic while preventing outright starvation of the lowest classes of traffic. At every hop along a route, packets are queued for transmission using this algorithm, multiplexing locally generated packets with packets routed on behalf of other users.

The queuing system is also aware of latency and reliability constraints for different classes of packets and will prune or queue traffic in accordance with these constraints. The pruning is based on packet age using sub-microsecond precision, with age limits dictated by a second policy map. This map takes the User Priority and DSCP of the packet as inputs.

While the queuing and pruning mechanisms regulate traffic flow within a single node, the queues feed a dynamically weighted media access timing system. The timing system then helps to balance traffic among the many nodes within transmit range of each other and among the many nodes along a route. Nodes routing high priority traffic on behalf of other nodes will gain high priority access to the media when transmitting the high priority traffic, even if the local user is currently using a lower priority. Also, nodes in an area advertise the aggregate indications of the depth and priority of the traffic in their queues. This enables nearby nodes to actively throttle traffic in busy neighborhoods or through busy devices. This throttling reduces packet collisions on the medium. Since collisions on the medium are wasteful of bandwidth, reducing the collisions improves throughput for the whole mesh system.

Management of Services

Management of the different aspects of the QoS feature set is performed by the MeshManager™ software through a standard SNMP interface. Network management packets are automatically tagged with the highest User Priority tag and the Network Management DSCP tag, so devices can be managed even when the network is heavily loaded.

Network managers can limit the DSCP codes set by the user's applications. User packets tagged with unauthorized DSCP codes are dropped at the sender's node.

Network managers can limit the maximum User Priority the user can select through MeshAPI, or change the default and current User Priority on a node-by-node basis. Separately, managers can enable, disable, activate, and deactivate the Emergency Mode feature on each node.

Users can set the local User Priority level for their own traffic within the bounds allowed by the network manager. Users can also activate and deactivate Emergency Mode if permitted by the network manager.

Future Efforts

Future efforts on the QoS features of the MEA network will center on more advanced management features.

One planned management enhancement is the ability to manage the DSCP and User Priority policy maps that control the behaviors of the transport mechanisms. These policies will be independently controllable at each node and allow the network manager to modify existing behaviors or define new ones based on DSCP and/or User Priority settings.