

Internet Engineering Task Force
INTERNET DRAFT

A. Ghanwani
J. W. Pace
V. Srinivasan
IBM
3 June 1996

Integrated Services over Token Ring Networks
draft-ghanwani-istr-00.txt

Status of This Memo

This document is an Internet-Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months, and may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material, or to cite them other than as a ``working draft'' or ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the internet-drafts Shadow Directories on ds.internic.net (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

Abstract

The RSVP and Integrated Services working groups have been working towards the goal of an "Integrated Services Internet" where applications can request a Quality of Service (QoS) from the network in order to meet their end-to-end service requirements. This document reviews various issues related to supporting Integrated Service models over token ring networks. It also describes features of token ring networks that can be leveraged to provide QoS guarantees. It is intended as a starting point for building an architectural framework to support these services in token ring environments.

1. Introduction

The Internet is soon expected to carry significant amounts of multimedia traffic requiring end-to-end Quality of Service (QoS) guarantees. In anticipation of this transition from largely data-oriented traffic to multimedia traffic with QoS over the Internet, the IETF is studying several service definitions within the Integrated Services (INTSERV) working group. These service definitions are called Guaranteed Delay, Predictive, Controlled Delay, and Controlled Load. The QoS guarantees that these services provide to flows range from the very stringent Guaranteed Delay service to the minimal -- "better than best effort" -- Controlled Load service.

The Internet spans a wide range of subnetwork technologies, from dedicated, high bandwidth, router-to-router, point-to-point links to shared LANs and dialup SLIP/PPP connections. This range of subnetworks continues to expand with the advent of new technologies such as ATM. Consequently, to be able provide a multimedia session with end-to-end guaranteed QoS over the Internet, we must be able to support the Integrated Services on any subnetwork that the session might traverse. To this end, the IETF recently formed the Integrated Services over Specific Link Layers (ISSLL) proto-working-group to develop techniques and guidelines needed to support Internet Integrated Service capabilities on specific network technologies. The purpose of this document is to describe some of the issues related to supporting Integrated Services over token ring networks. Token ring networks possess many features such as source routing and frame priorities that differentiate them from other shared media LAN technologies. In this document, we will identify key features of token rings that may be exploited in order to provide Integrated Services and describe a possible framework for supporting QoS in a shared environment.

2. Source Routing

The token ring MAC frame format is shown in Figure 1.




Figure 1: Token ring Network Frame Format. SD is the starting delimiter; AC is the access control field; FC is the frame control field; ED is the ending delimiter; and FS is the frame status field.

Source routing is an important feature that differentiates token ring networks from other LAN types. In source routing, each frame carries information about the route it will follow through a multiple-segment bridged network. This information is specified in the optional Routing Information (RI) field in the frame. The RI field (RIF) can be up to 32 bytes. When the RIF is included by the station originating the frame, bridges processing the frame examine this field to determine how to forward the frame.

The RIF is made up of a 2 byte control field, followed by up to 30 bytes of Route Designator (RD) fields. The control field contains:

- the type of frame,
- length of the RIF,
- direction in which to process the RIF, and
- largest-size information field that can be transmitted between two communicating stations on a specific route.

Each segment in a given multiple-segment network is assigned a unique ring number, and each bridge is assigned a number which may or may not be unique. Together, a (ring number, bridge number) pair make up a route designator. The RD fields in the RIF are made up of these pairs that specify the frame's path through the network.

There are three token ring frame types that can be specified in the control portion of the RIF:

- Specifically Routed Frame (SRF): this indicates that the RD fields contain a specific route for the frame to traverse the network.
- All Routes Explorer (ARE): this indicates that the frame will be transmitted along every route in the network to the destination station. Frames transmitted as ARE will result in as many copies being delivered to the destination as there are different routes to the destination station.
- Spanning Tree Explorer (STE): this indicates that only certain designated bridges (the ones on the spanning tree) will relay the frame from one segment to another with the result that the frame will appear exactly once on every segment in the network.

When an ARE or STE frame is transmitted, each bridge that forwards the frame to another ring adds its bridge number and that ring's number to the frame's RIF. When such a frame reaches its destination, the routing information in the frame indicates the path taken by the frame through the network. This routing information can be used for subsequent communication between the stations, eliminating the need to broadcast frames and conserving network bandwidth. In this manner, source routing provides a dynamic mechanism to select routes between stations in a token ring network. This feature can be used in conjunction with a resource allocation mechanism, such as the one

described in the last section, to perform "QoS routing" of RSVP flows through multiple ring networks.

3. Token Priority

The token ring standard [1] provides a priority mechanism that can be used to control both the queuing of packets for transmission and the access of packets to the shared media. The priority mechanisms are implemented using bits within the Access Control (AC) and the Frame Control (FC) fields of a LLC frame. The first three bits of the AC field, the Token Priority bits, together with the last three bits of the AC field, the Reservation bits, regulate which stations get access to the ring. The last three bits of the FC field of an LLC frame, the User Priority bits, are obtained from the higher layer when it requests transmission of a packet. The requested user priority establishes the requested access priority. The user priority is conveyed end-to-end by the User Priority bits in the FC field. In all cases, B'000' is the lowest priority.

A token ring station is theoretically capable of separately queuing each of the eight levels of requested user priority and then transmitting frames in order of priority. A station sets Reservation bits according to the user priority of frames that are queued for transmission in the highest priority queue. This allows the access mechanism to ensure that the frame with the highest priority throughout the entire ring will be transmitted before any lower priority frame. Annex I to the IEEE 802.5 token ring standard ([1]) recommends that stations queue MAC frames and LAN management frames with a user priority of 7 and 4, respectively. For LLC frames, the annex recommends that time sensitive traffic (e.g., video playback) be queued with user priority 5, real time critical service traffic (e.g., interactive video) with user priority 6, and non-time critical traffic with user priority 0. Priorities 1-3 are left for use by other user traffic. To reduce frame jitter associated with high-priority traffic, the annex also recommends that only one frame be transmitted per token and that the maximum information field size be 4399 octets whenever delay-sensitive traffic is traversing the ring. Most existing implementations of token ring bridges forward all LLC frames with a default access priority of 4. Annex I recommends that bridges forward LLC frames that have a user priorities greater than 4 with a reservation equal to the user priority. (The draft IEEE Standard P802.1p ([2]) permits network management to control whether bridges forward frames with the outbound user priority set to the same value as received in the user priority field. It also describes the forwarding process for bridges that support traffic classes and multiple queues.) The capabilities provided by token ring's user and reservation priorities and by IEEE

802.1p can provide effective support for Integrated Services flows that request QoS using RSVP. These mechanisms can provide, with few or no addition to the token ring architecture, bandwidth guarantees and the network flow control necessary to support guarantees.

4. Multicast Addressing

RSVP is designed to inherently support multicast IP destination addresses. References [3] and [4] define methods for mapping multicast IP addresses (respectively) to multicast MAC addresses for use on all IEEE 802 LANs and to functional addresses for use on token ring networks. IEEE Standard 802.1p is also defining a method to support the filtering of multicast MAC addresses by LAN bridges and switches. With a fully implemented network of 802.1p compliant bridges, it therefore becomes possible to filter Integrated Services flows at both the IP layer (via IGMP) and at the MAC layer (via 802.1p). Multicast filtering can potentially reduce the bandwidth consumed by Integrated Services flows on parts of the spanning tree. This filtering is available on both token ring and IEEE 802.3 networks. In the absence of 802.1p filtering, source routing might be used on token ring networks to limit the bandwidth consumed on those segments of the network that have no receivers for a particular Integrated Services flow.

5. Client-Server QoS Framework

One method of exploiting the above features of token ring and of supporting QoS in a shared LAN is via a client-server mechanism. The server implements a "QoS Allocator" that is responsible for maintaining awareness of the resource usage of various LAN components (e.g. bridges, segments). End stations desiring to transmit data with QoS guarantees implement a "QoS Requestor," which is the client component responsible for making requests to reserve resources across the LAN. The Allocator determines whether a particular request for a quality of service allocation can be satisfied, and maintains knowledge of existing QoS allocations. (Note: this simple client-server framework has been informally referred to as a "Mother May I?" protocol.)

The Allocator may be implemented in a distributed fashion or as a single entity. In the former case, each token ring segment could have a separate Allocator, and all Allocators would periodically exchange information about available resources on all segments. If implemented as a single entity, there would be a single Allocator for all segments and bridges in the network. Having a centralized Allocator is simpler to implement, and avoids synchronization and

deadlock problems. However, this approach does not scale well with larger networks.

In order to support Integrated Services, the Requestor-Allocator framework must interact with higher layer QoS signalling mechanisms such as RSVP. Each RSVP router/host in the network must use the QoS Requestor to reserve resources for RSVP flows. When an RSVP RESV message containing a reservation request is received at a station (router/host) on the LAN, it must be translated into an appropriate request to the Allocator to reserve resources at Layer 2. The Allocator will process the request and send a reply to the Requestor with a success/failure indication. This reply is returned to the RSVP daemon which will take appropriate action such as forwarding the RESV message upstream, or sending a RESVERR to the originator of the RESV message.

The QoS Allocator can use the various features of token rings, such as source routing and token priorities, in order to accomplish resource reservations for RSVP flows. In the earlier sections, we have described some of these features.

6. References

- [1] IEEE Standards for Local and Metropolitan Area Networks: Token Ring Access Method and Physical Layer Specifications. IEEE Std 802.5-1995.
- [2] IEEE Standards for Local and Metropolitan Area Networks: Draft Standard for Traffic Class and Dynamic Multicast Filtering Services in Bridged Local Area Networks (Draft Supplement to 802.1D), P802.1p/D3, April 30, 1996.
- [3] S. Deering, Host Extensions for IP Multicasting, Std 5, [RFC 1112](#), August 1989.
- [4] S. Thomas, A Method for the Transmission of IPv6 Packets over Token Ring Networks, [draft-ietf-ipnwg-token-ring-01.txt](#), work in progress, Jan 22, 1996.

7. Authors' Address

Anoop Ghanwani
IBM Corporation
P. O. Box 12195
Research Triangle Park, NC 27709

Phone: +1-919-254-0260
Fax: +1-919-254-5410
E-mail: anoop@raleigh.ibm.com

Wayne Pace
IBM Corporation
P. O. Box 12195
Research Triangle Park, NC 27709

Phone: +1-919-254-4930
Fax: +1-919-254-5410
E-mail: pacew@raleigh.ibm.com

Vijay Srinivasan
IBM Corporation
P. O. Box 12195
Research Triangle Park, NC 27709

Phone: +1-919-254-2730
Fax: +1-919-254-5410
E-mail: vijay@raleigh.ibm.com

