Internet Engineering Task Force
INTERNET-DRAFT
draft-ietf-rsvp-partial-service-00.txt

RSVP WG
L. Breslau/S. Shenker
Xerox PARC

April 23, 1997
Expires: 10/23/97

Partial Service Deployment in the Integrated Services Architecture

Status of this Memo

Abstract

Specifications for providing enhanced qualities of service in the Internet have been defined in [2,4]. Technical and administrative concerns may prevent a network element from offering one or more of these services. In this document, we present a mechanism for dealing with heterogeneity in the set of services offered by different network elements. This approach enables end-to-end service to be composed of different per-hop services while not requiring end systems to be aware of the details of the service provided at each hop.

Introduction

The Integrated Services Working Group has produced specifications for new types of service[2,4]. These services will provide enhanced qualities of service to applications that use them. Such services will be most useful when they are provided at all network elements along a data distribution path. However, because these services impose stricter requirements on the network elements than traditional best-effort service, it may not be practical to provide these services on some subnet technologies. Furthermore, the ultimate decision to implement and deploy a particular service belongs to vendors and network service providers, respectively. A vendor may choose not to implement a service, or network service provider may choose not to offer a service (e.g., for administrative reasons), even if the underlying technology is able to support the service. Therefore, newly defined services will not always be available end-to-end along a data distribution path. In this document we describe a strategy for coping with this heterogeneity in the set of services offered in a network.

Replacement Services

The mechanism for addressing the problem of service heterogeneity is built on the concept of "replacement" services. When a network element does not offer a service, it can offer a replacement service for it. Under circumstances described below, when a network element receives a request for a service it does not offer, it treats the request as if it were for the replacement service. A replacement service can be one of the other real-time services, best-effort service, or a non-compliant implementation of the original service. Decisions about the use of replacement services are a local matter.

Replacement services are characterized as either "reliable" or "unreliable". A reliable replacement is one that is expected to meet the requirements of the service being replaced a large majority of the time (e.g., over 95%). No assurance is given about the resulting quality of an unreliable replacement. Since best effort service qualifies as an unreliable replacement in all circumstances, network elements are always able to at least offer unreliable replacements for every service.

When a network element offers a reliable replacement it MUST also export meaningful values for any characterization parameters required by the original service. (See [1] and [3] for a discussion of

characterization parameters.) These parameters must accurately characterize the replacement service, and when composed end-to-end with parameters from other network elements they must provide applications with a valid characterization of the end-to-end service. When a network element offers an unreliable replacement it MAY export values for any characterization parameters of the original service if it is able to accurately characterize the service. However, given that an unreliable replacement may be arbitrarily bad, the network element may instead set the value of all characterization parameters to the "invalid" values defined for those parameters.

We provide the following guidelines for determining whether a network element offers an actual service, a reliable replacement, or an unreliable replacement.

Offered Service: In order to claim to offer a service, a network element's implementation of the service must conform to the service specification. Specific conformance requirements are included in the service specification documents (e.g., [2,4]). Some knowledge about the environment in which an implementation is deployed can be used in making this determination. For example, an implementation of Controlled-Load in a router attached to an ethernet may be compliant if all routers and hosts attached to the network participate in a distributed admission control process to reserve resources on the ethernet. However, the same implementation may not be compliant if there are hosts attached to the ethernet that do not participate in the bandwidth allocation procedure. Knowledge of link bandwidth can also be used to determine service compliance. For example, a router may be able to offer a service on an interface using FIFO scheduling and no admission control if the bandwidth on the link exceeds the sum of the input bandwidths on all other links. On the other hand, knowledge about average levels of offered load cannot be used to claim compliance with the currently defined service specifications; the offered service must comply with the service specifications independent of ambient loading.

Reliable Replacement: A network element can claim to offer a reliable replacement if it does not offer the actual service but the service it provides is expected to adhere to the specification of the original service a large majority of the time. This determination can take local conditions, including expected load, into account. However, since the semantics of a reliable replacement are that it emulates very closely the actual service, applications using reliable replacements will expect to receive service that is in general not significantly different than the original service. Therefore, the reliable replacement label should be applied to service replacements with caution. Furthermore, since a reliable replacement will often depend on local conditions, and conditions may change over time,

network operators should monitor these conditions and continually reassess the suitability of reliable replacements. Finally, note that the ability to provide a reliable replacement may also depend on the availability of appropriate invocation parameters for the replacement service.

We offer two examples of reliable replacements. First, a router on a vastly underutilized point-to-point link, which rarely experiences persistent congestion, may offer best effort service as a reliable replacement for Controlled Load service. Second, a router attached to an ethernet that has an otherwise compliant implementation of Controlled Load but has no means to control the load generated by other stations attached to the ethernet may claim to offer a reliable replacement for Controlled Load if the ethernet is not in general highly loaded.

Unreliable Replacement: Since unreliable replacements make no assurances about the service they provide, any service qualifies as an unreliable replacement for any other service. Hence, when the actual service or a reliable replacement is not offered, an unreliable replacement can always be offered. For example, best effort service on a congested link qualifies as an unreliable replacement for any real-time service. Applications should have no expectations about the resulting service when they use an unreliable replacement. Finally, additional real-time services may be defined after a particular implementation is deployed. Hence, a network element may not even know about a requested service, so it cannot make an informed decision about the suitability of its offered services (other than best effort) to act as replacements. In such cases, the router can always use best effort as an unreliable replacement.

Characterizing Service Offerings

Each network element must export characterization parameters describing the various services and replacements that it offers. An example format and composition rules for these characterization parameters are described in the Appendix to this document.

Service Handling Flags

When characterization parameters are provided end-to-end by a setup protocol, an application will know before issuing a service request whether any replacement services will be substituted for its request.

Applications that would be dissatisfied with the level of assurance provided by the resulting service should refrain from issuing service requests when such substitutions would be made.

The Integrated Services architecture is intended to allow services to be invoked by more than one setup protocol or by network management functions. Therefore, we cannot assume that end-to-end characterizations of service offerings will always be available to the applications. When they are not, mechanisms are needed so that applications can express to the network elements their willingness to accept replacement services.

We propose that application preferences be expressed in a new object that we refer to as the Service Handling Flags, which is optionally included in service requests. These flags are associated with service requests in general, and not with specific services, so they are associated with service_name 0 (just like general characterization parameters). The parameter is a 16-bit value. The most significant bit is set to 1 if service replacements are *not* allowed and 0 if replacements are allowed. The remaining 15 bits are currently unused.

Hence, the default router action is to perform replacements when a requested service is not available. In environments where end-to-end characterizations are available (e.g., as with RSVP) the Service Handling Flags are not needed. When end-to-end characterizations are not available, an end system must include the Service Handling Flags with the most significant bit set in order to prevent the use of replacement services.

Security Considerations

Security considerations are not discussed in this memo.

Appendix -- Service Availability Characterization Parameters

The following is an example format for characterization parameters describing service availability.

An integrated services aware router exports a general characterization parameter for each service that it knows about indicating whether it offers the service, a reliable replacement for the service, or an unreliable replacement for the service. These parameters, when composed end-to-end, inform the endpoints about the

end-to-end availability of services.

The parameter_name for the local parameter is N. A single router can export multiple characterization parameters with parameter_name N, each corresponding to a different service. The specific service referenced by a particular parameter is identified by a field within the parameter itself.

Each local parameter is represented by a sequence of 4 16-bit unsigned integers in network byte order. The first is the service_name for the service referenced by the parameter. The service_name is defined within the service specification for each service (e.g., see [2,4]). The second has the value 1 if the router offers the indicated service and 0 if the router does not offer the service. The third field has the value 1 if the router offers a reliable replacement for the service and 0 if the router does not offer a reliable replacement for the service. The fourth field has the value 1 if the router offers an unreliable replacement for the service and 0 if the router does not offer an unreliable replacement for the service. A router must assign a value of 1 to exactly one of the latter three fields. Guidelines for offering reliable replacements and unreliable replacements are specified earlier in this document.

The parameter_name for the composed end-to-end parameter is N+1. The specific service referenced by a particular parameter is specified by a field inside the parameter itself.

Each composed parameter is represented by a sequence of 4 16-bit unsigned integers in network byte order. The first is the service_name for the service characterized by the parameter. The second is the number of routers that offer the service. The third is the number of routers that offer reliable replacements for the service. The fourth is the number of routers that offer unreliable replacements for the service.

The composition rule for composing a local parameter with a composed parameter is to add the i_th value of the local parameter to the i_th value of the composed parameter, for i = {2,3,4}. Two parameters can be composed only if the first field in each parameter (the service_name) is the same.

References

[1] S. Shenker and J. Wroclawski. "Specification of General

Characterization Parameters", Internet Draft, October 1996, <draft-ietf-intserv-charac-02.txt>.

[2] S. Shenker, C. Partridge and R. Guerin. "Specification of Guaranteed Quality of Service", Internet Draft, February 1997, <draft-ietf-intserv-guaranteed-svc-07.txt>.

[3] S. Shenker and J. Wroclawski. "Network Element Service Specification Template", Internet Draft, November 1996, <draft-ietf-intserv-svc-template-03.txt>.

[4] J. Wroclawski. "Specification of Controlled-Load Network Element Service", Internet Draft, November 1996, <draft-ietf-intserv-ctrl-load-svc-04.txt>.

Authors' Addresses:

Lee Breslau
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304-1314
breslau@parc.xerox.com
415-812-4402
415-812-4471 (FAX)

Scott Shenker
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304-1314
shenker@parc.xerox.com
415-812-4840
415-812-4471 (FAX)