Network Working Group
Request for Comments: 1940
Category: Informational

D. Estrin
USC
T. Li
Y. Rekhter
cisco Systems
K. Varadhan
D. Zappala
USC
May 1996

Source Demand Routing: Packet Format and Forwarding Specification (Version 1).

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

1. Overview

The purpose of SDRP is to support source-initiated selection of routes to complement the route selection provided by existing routing protocols for both inter-domain and intra-domain routes. This document refers to such source-initiated routes as "SDRP routes". This document describes the packet format and forwarding procedure for SDRP. It also describes procedures for ascertaining feasibility of SDRP routes. Other components not described here are routing information distribution and route computation. This portion of the protocol may initially be used with manually configured routes. The same packet format and processing will be usable with dynamic route information distribution and computation methods under development.

The packet forwarding protocol specified here makes minimal assumptions about the distribution and acquisition of routing information needed to construct the SDRP routes. These minimal assumptions are believed to be sufficient for the existing Internet. Future components of the SDRP protocol will extend capabilities in this area and others in a largely backward-compatible manner.

This version of the packet forwarding protocol sends all packets with the complete SDRP route in the SDRP header. Future versions will address route setup and other enhancements and optimizations.

2. Model of operations

An Internet can be viewed as a collection of routing domains interconnected by means of common subnetworks, and Border Routers (BRs) attached to these subnetworks. A routing domain itself may be composed of further subnetworks, routers interconnecting these subnetworks, and hosts. This document assumes that there is some type of routing present within the routing domain, but it does not assume that this intra-domain routing is coordinated or even consistent.

For the purposes of this discussion, a BR belongs to only one domain. A pair of BRs, each belonging to a different domain, but attached to a common subnetwork, form an inter-domain connection. By definition, packets that traverse multiple domains must traverse BRs of these domains. Note that a single physical router may act as multiple BRs for the purposes of this model.

A pair of domains is said to be adjacent if there is at least one pair of BRs, one in each domain, that form an inter-domain connection.

Each domain has a globally unique identifier, called a Domain Identifier (DI). All the BRs within a domain need to know the DI assigned to the domain. Management of the DI space is outside the scope of this document. This document assumes that Autonomous System (AS) numbers are used as DIs. A domain path (or simply path) refers to a list of DIs such as might be taken from a BGP AS path [1, 2, 3] or an IDRP RD path [4]. We refer to a route as the combination of a network address and domain paths. The network addresses are represented by NLRI (Network Layer Reachability Information) as described in [3].

This document assumes that the routing domains are congruent to the autonomous systems. Thus, within the content of this document, the terms autonomous system and routing domain can be used interchangeably.

An application residing at a source host inside a domain, communicates with a destination host at another domain. An intermediate router in the path from the source host to the destination host may decide to forward the packet using SDRP. It can do this by encapsulating the entire IP packet from the source host in an SDRP packet. The router that does this encapsulation is called the "encapsulating router."

2.1 SDRP routes

A component in an SDRP route is either a DI (AS number) or an IP address. Thus, an SDRP route is defined as a sequence of domains and routers, syntactically expressed as a sequence of DIs and IP addresses. Thus an SDRP route is a collection of source routed hops.

Each component of the SDRP route is called a "hop." The packet traverses each component of the SDRP route exactly once. When a router corresponding to one of the components of the SDRP route receives the packet from a router corresponding to the previous component of the SDRP route, the router will process the packet according to the SDRP forwarding rules in this packet. The next component of the SDRP route that this router will forward the packet to, is called the "next hop," with respect to this router and component of the SDRP route.

An SDRP hop can either be a "strict" source routed hop, or a "loose" source routed hop. A strict source route hop is one in which, if the next hop specified is a DI, refers to an immediately adjacent domain, and the packet will be forwarded directly to a route within the domain; if the next hop specified is an IP address, refers to an immediately adjacent router on a common subnetwork. Any other kind of a source route hop is a loose source route hop.

A route is a "strict source route" if the current hop being executed is processed as a strict source route hop. Likewise, a route is a "loose source route" if the current hop being executed is processed as a loose source route hop.

It is assumed that each BR participates in the intra-domain routing protocol(s) (IGPs) of the domain to which the BR belongs. Thus, a BR may forward a packet to any other BR in its own domain using intra-domain routing procedures. Forwarding a packet between two BRs that form an inter-domain connection requires neither intra-domain nor the inter-domain routing procedures (an inter-domain connection is a common Layer 2 subnetwork).

It is also assumed that all routers participate in the intra-domain routing protocol(s) (IGPs) of the domain to which they belong.

While SDRP does not require that all domains have a common network layer protocol, all the BRs in the domains along a given SDRP route are required to support a common network layer. This document specifies SDRP operations when that common network layer

protocol is IP ([5]).

While this document requires all the BRs to support IP, the document does not preclude a BR from additionally supporting other network layer protocols as well (e.g., CLNP, IPX, AppleTalk). If a BR supports multiple network layers, then for the purposes of this model, the BR must maintain multiple Forwarding Information Bases (FIBs), one per network layer.

2.2 SDRP encapsulation

Forwarding an IP packet along an SDRP route is accomplished by encapsulating the entire packet in an SDRP packet. An SDRP packet consists of the SDRP header followed by the SDRP data. The SDRP header carries the SDRP route constructed by the domain that originated the SDRP packet. The SDRP data carries the original packet that the source domain decided to forward via SDRP.

An SDRP packet is carried across domains as the data portion of an IP packet with protocol number 42.

This document refers to the IP header of a packet that carries an SDRP packet as the delivery IP header (or just the delivery header). This document refers to the packet carried as SDRP data s the payload packet, and the IP header of the payload packet is the payload header.

Thus, an SDRP Packet can be represented as follows:

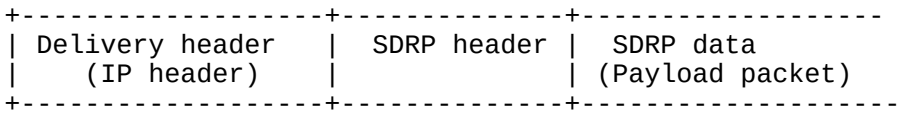

Each SDRP route may have an MTU associated with it. An MTU of an SDRP route is defined as the maximum length of the payload packet that can be carried without fragmentation of an SDRP packet. This means that the SDRP MTU as seen by the transport layer and applications above the transport layer is the actual link MTU less the length of the Delivery and SDRP headers. Procedures for MTU discovery are specified in Section 9.

2.3 D-FIB

It is assumed that a BR participates in either BGP or IDRP. A BR participating in SDRP augments its FIBs with a D-FIB that contains routes to domains. A route to a domain is a triplet <DI, Next-Hop, NLRI>, where DI depicts a destination domain, Next-Hop

depicts the IP address of the next-hop BR, and NLRI depicts the set of reachable destinations within the destination domain. D-FIBs are constructed based on the information obtained from either BGP, IDRP, or configuration information.

An SDRP packet is forwarded across multiple domains by utilizing the forwarding databases (both FIBs and D-FIBs) maintained by the BRs.

The operational status of SDRP routes is monitored via passive (Error Reporting) and active (Route Probing) mechanisms. The Error Reporting mechanism provides the originator of the SDRP route with a failure notification. The Probing mechanism provides the originator of the SDRP route with confirmation of a route's feasibility.

3. SDRP Packet format

The total length of an SDRP packet (header plus data) can be determined from the information carried in the delivery IP header. The length of the payload packet can be determined from the total length of an SDRP packet and the length of its SDRP Header.

The following describes the format of an SDRP packet.

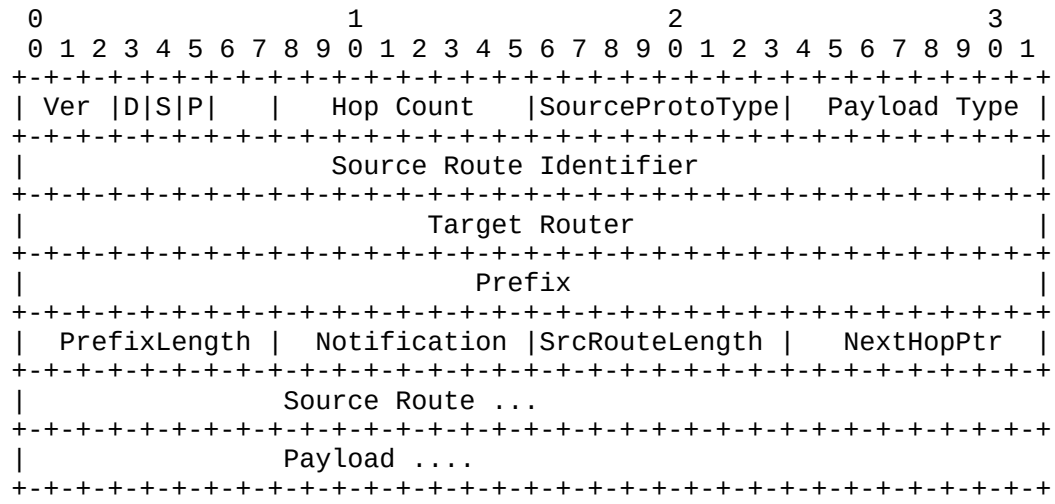

Version and Flags (1 octet)

The SDRP version number and control flags are coded in the first octet. Bit 0 is the most significant bit, bit 7 is the least significant bit.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.