Position Summary: Anypoint Communication Protocol

Ken Yocum, Jeff Chase, and Amin Vahdat
Department of Computer Science, Duke University
{grant,chase,vahdat}@cs.duke.edu

It is increasingly common to use redirecting intermediary switches to *virtualize* network service protocols. Request redirection enables an intermediary to represent a dynamic set of servers as a unified service to the client. Services virtualized in this way include HTTP (using L4-L7 switches), NFS, and block storage protocols.

Virtualization using intermediaries is a powerful technique for building scalable cluster-based services while insulating clients from the details of server structure. However, intermediaries are controversial and difficult to implement in part because transport protocols are not designed to support them. For example, intermediaries compromise the end-to-end guarantees of point-to-point transports such as TCP. Current service intermediaries are constrained to either route requests at a connection granularity (L4-L7 switches for HTTP), use weak transports such as UDP, or terminate connections at the intermediary. These limitations compromise performance and generality. In particular, independent routing of requests is necessary for any content-based routing policy, but we know of no efficient intermediary scheme that supports independent routing for multiple requests arriving on the same transport connection. The challenges are increasingly evident as designers attempt to build intermediaries for commercially important protocols such as HTTP 1.1 and iSCSI.

These difficulties motivate consideration of new transport protocols with more decentralized notions of what constitutes a connection "endpoint". We are developing such a transport called the Anypoint Communication Protocol (ACP). ACP clients establish connections to abstract services, represented at the network edge by Anypoint intermediaries. The intermediary is an intelligent network switch that acts as an extension of the service; it encapsulates a service-specific policy for distributing requests among servers in the *active set* for each service. The switch routes incoming requests on each ACP connection to any active server at the discretion of the service routing policy, hence the name "Anypoint".

The ACP transport is similar to SCTP and TCP in that it provides reliable, sequenced delivery with congestion control. However, ACP defines some protocol properties as *end-to-edge* rather than *end-to-end*. A critical respect in which ACP is end-to-edge is that it does not define the delivery order for requests routed to different servers, or for responses returned from different servers. Ordering constraints and server coordination are the responsibility of the service protocol and its routing policy.

An Anypoint intermediary orchestrates the movement of requests and responses at the transport layer. To this end, ACP frames service protocol requests and responses at the transport layer in a manner similar to SCTP. Transport-level framing allows an Anypoint switch to identify frames from the network stream in a general way. The switch applies the service-specific routing policy to each inbound frame, and merges outbound frames into a single stream to the client.

While ACP is fundamentally similar to other reliable Internet transports, a central design challenge is that ACP connection endpoint state and functions are distributed between the intermediary and the end server nodes. ACP is designed to enable fast, space-efficient protocol intermediaries with minimal buffering. Acknowledgment generation, buffering of unacknowledged frames, and retransmission are the responsibility of the end nodes, thus reliable delivery is guaranteed end-to-end rather than end-to-edge. The Anypoint switch maintains a mapping between sequence number spaces seen by the client and end server nodes for each connection, for a bounded number of unacknowledged frames. The switch also coordinates congestion state across the active set of participants in each ACP connection. The congestion scheme assumes that the bottleneck transit link is between the switch and the client, or that the ACP stream may be throttled to the bandwidth to the slowest end server selected by the routing policy.

The Anypoint abstraction and ACP protocol enable virtualization using intermediaries for a general class of wide-area network services based on request/response communication over persistent transport connections. Potential applications include scalable IP-based network storage protocols and next-generation Web services.