New Mechanism for Datagram Service

Brian Batchelder

Hewlett-Packard

May 24, 1999

Definitions

- Connection-oriented vs. Connectionless
  - Connection-oriented: Persistent communication path between two endpoints.
  - Connectionless: Time-of-transfer communication path between two endpoints.

Definitions

- Stream vs. Message
  - Stream: Data is delivered as a stream of bytes w/o maintaining boundaries.
  - Message: Boundaries are maintained in delivered data.

Definitions

- Reliable vs. unreliable
  - Reliable: Data will neither be lost nor duplicated.
  - Unreliable: Data may be lost or duplicated.

Definitions

- Guaranteed Order vs. not guaranteed order
  - Guaranteed Order: Data will arrive in the order in which it was sent and it will not be duplicated
  - Guaranteed Order: Data may arrive out of order or it may be duplicated

Datagram Services

- Datagram services are typically unreliable, connectionless message services
- UDP is an example of an typical datagram service
- SNMP is an example of an application that uses a typical datagram service (UDP)

Datagram API

- socket()
  - socket_type = SOCK_DGRAM
- bind()
  - for servers
- sendto()
  - same as send() plus destination address
- recvfrom()
  - same as recv() plus returns source address

Connectionless

- We could emulate using connections, but:
  - setup and tear-down overhead
  - might use up available queues
  - how do we know when to tear-down queue?

Unreliable, Connectionless Datagrams using Queue 0

- New control: TransferDatagram
- Request only, no response
- New variable-length parameter: Datagram
- TranferDatagram(service ID, datagram)

I2T Reliability on Queue 0

- Initiator
  - Log in, if necessary
  - If I2T ORB is available, and datagram fits, queue it
  - Otherwise, bit-bucket
- Target
  - If buffer is available (recvfrom), deliver as much of the datagram as fits
  - Otherwise, bit-bucket

T2I Reliability on Queue 0

- Target
  - Target initiated login, if necessary
  - If T2I ORB is available, and datagram fits, write it and complete ORB
  - Otherwise, bit-bucket
- Initiator
  - If buffer is available (recvfrom), deliver as much of the datagram as fits
  - Otherwise, bit-bucket