SWINBURNE UNIVERSITY OF TECHNOLOGY

Forwarding SCTP Traffic through a NAT

Jason But

jbut@swin.edu.au

Centre for Advanced Internet Architectures (CAIA)

Swinburne University of Technology

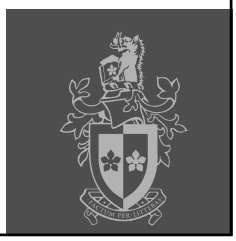

Outline

- SCTP Overview
- NAT and SCTP
- How to NAT SCTP Packets
- SONATA and alias_sctp
- Conclusion

SCTP – What is it?

- SCTP is a Transport Layer Protocol that supports
  - A selection of guaranteed/best-effort/mixed-mode delivery
  - Multi-homed hosts
  - Multiple streams within an association
  - Message oriented service
  - Selective Acknowledgements
  - CRC32 Checksum for better error detection
- Originally developed for voice signalling traffic but potential uses are more widespread

Setting up an Association

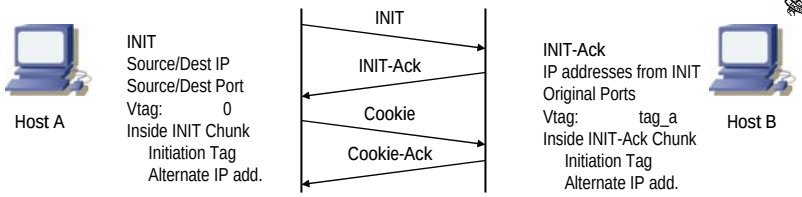

INIT

- Host B learns
  - Alternate addresses for A
  - SCTP Ports to use
  - vtag to include in all subsequent packets to A (Initiation Tag)
- Determine State Cookie

INIT-Ack

- Check source/ports/vtag
- Host A learns
  - Alternate addresses for B
  - vtag to include in all subsequent packets to B (Initiation Tag)
- Echo State Cookie
- Association is UP for Host A

Cookie

- Check source/ports/vtag
- Host B unpacks Cookie
  - Sets up kernel TCB for association
- Data may be appended to the Cookie chunk
- Association is UP for Host B

What is the deal with…

- The State Cookie
  - Designed to protect against DoS attacks
  - Server association up second, no resources allocated until Client-side is up
- Vtags
  - An extra piece of information to help protect against spoofing
  - Allows multiple associations to one host to use the same port pair as long as the vtags are unique

Association Maintenance

- Communications are not spread amongst all interfaces on a host
- Unused interfaces will periodically send a Heartbeat to alternate host and receive a Heartbeat-Ack
  - Tells remote host that alternate path is still available
  - Updates alternate path RTT
- Other interfaces can be added/removed using the SCTP ASCONF extensions
  - AddIP – Add a new interface to the association
  - RmIP – Remove an interface from the association

NAT and SCTP

- NAT allows us to share a single public IP addressing by mapping a private address space to a public IP address
  - Typically used within home-user environments
- Important because
  - SCTP will never see widespread use until it is usable from behind a NAT
  - Eg. BT already runs ~1,000,000 concurrent SCTP SIP sessions over a private network

Extend Existing Implementation for SCTP

- Not feasible
- SCTP Checksum – CRC32 is computationally intensive to recalculate
  - Difficult for consumer broadband equipment
  - Also difficult for corporate NATs dealing with high traffic volume
- SCTP Multi-homing
  - NAT has to manage packets for a single association from potentially multiple sources
  - Database management becomes more complex

Multi-homing Problems

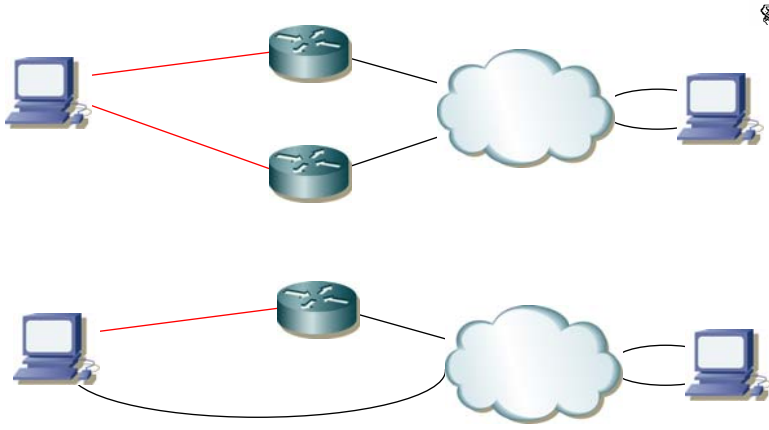

How to NAT SCTP Packets

- We can't modify port numbers as per traditional NAT
  - CRC32 Checksum is computationally intensive
- Need to track
  - Vtags within tuples – this allows retaining the port numbers as-is
  - Allow for multi-homed hosts both within and without the NAT – there are two ways to manage this
- Basic approach is outlined in
  - Soon to be released Internet Draft
  - Paper submitted for publication
    - M. Tüxen, I. Rüngeler, R. Stewart, E. Rathgeb, "Network Address Translation (NAT) for the Stream Control Transmission Protocol (SCTP)", Submitted to IEEE Networks Special Issue on Implication and Control of Middleboxes in the Internet.

SCTP Modifications to assist NAT

- Only Public IP addresses are allowed in an INIT, INIT-Ack or AddIP SCTP message
  - This ensures that remote hosts only learn about real IP addresses
  - Private addresses are learnt from the source IP address of the packet – modified by the NAT
- New AddIP format for private addresses
  - Includes both vtags within message so NAT can learn information

SCTP Modifications to assist NAT

- A new Abort message (AbortM) can be sent to the sender of an INIT or INIT-Ack
  - Will cause recipient to select a new vtag and re-transmit the packet
  - Allows notification of a collision in the NAT table and shortcut the timeout at the end-host
- A new Error message (ErrorM) can be sent to the sender of an SCTP packet
  - Indicates that the NAT is not aware of this association – causes recipient to initiate an AddIP message to update the NAT database
  - Allows for multi-homed hosts behind multiple NATs to learn about new associations and extra end-hosts within that association

Basic Approach

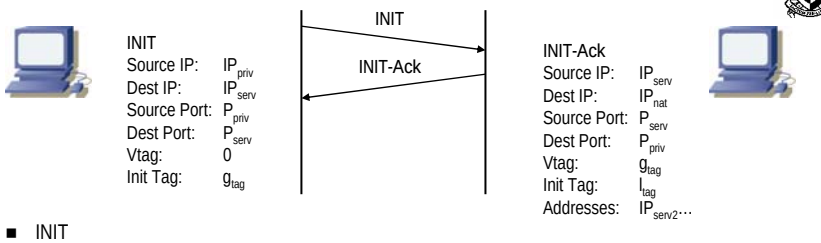

- INIT
  - Half populate NAT table entry
    - local_ip = IP_{priv}, local_port = P_{priv}, local_tag = 0, global_ip = IP_{serv}, global_port = P_{serv}, global_tag = g_{tag}
  - NAT IP_{priv} => IP_{nat}
- INIT-Ack
  - Check source IP, ports, vtag for matching entry
  - Complete NAT entry
    - local_tag = l_{tag}, extract all IP addresses in INIT-Ack and store in NAT table (IP_{serv2}…)
  - NAT IP_{nat} => IP_{priv}

Basic Approach

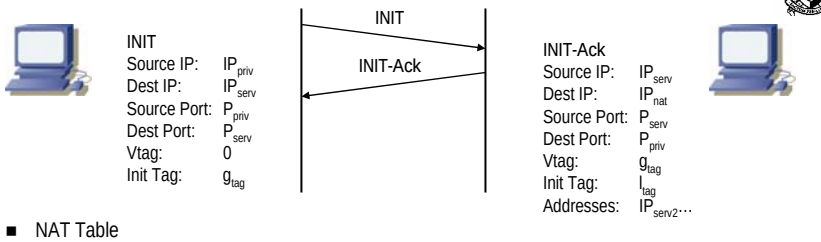

- NAT Table
  - local_ip = IP_{priv}, local_port = P_{priv}, local_tag = l_{tag}, global_ip = (IP_{serv} IP_{serv2}) global_port = P_{serv}, global_tag = g_{tag}
- All other packets
  - Outgoing
    - local_ip/Ports/local_tag match, global_ip in list. NAT local_ip and forward
    - No match, send ErrorM to private host
  - Incoming
    - Ports/global_tag match, global_ip in list. NAT to local_ip and forward
    - No match, drop silently

Multi-homed Private Hosts

- INIT goes out one private address to server
  - NAT_1 gets entry
- INIT-Ack comes back via same connection
  - NAT_1 brings up entry
- Private host sends AddIP(0.0.0.0) out alternate address
  - NAT_2 populates entry
- AddIP-Ack comes back via alternate address
  - NAT_2 confirms entry
- Both NATs are now aware of the association, ports and vtags
- Subsequent NATs may not know all IP addresses of global host

Do We Need to Track Global Addresses

- Proposed Internet Draft says yes
- Optional to assume that any global IP address where the port numbers and vtags match is part of the existing association
  - Don't store any global IP addresses in the NAT DB
  - Will result in more potential collisions in the NAT
  - Resolved via use of AbortM messages
- More efficient

Tracking Global Addresses – Advantages

- Less database collisions – more associations will get up without an AbortM being sent
  - Better for older SCTP stacks that do not recognise an AbortM
- Can drop forged packets *before* they enter the private network

Tracking Global Addresses - Disadvantages

- Resource Requirements
  - Parsing INIT, INIT-Ack, AddIP messages to extract public IP addresses
  - Store a more complex database entry (and lookup)
- More edge cases
  - NAT not knowing all addresses, increased ErrorM messages
  - More complex handling when different hosts contact a multi-homed host outside the NAT

Tracking Global Addresses – State

- Extended debate between myself and ID author about need to track Global IP Addresses
  - Hung up on potential number of collisions
  - Hung up on issue of older stacks
- Managed to extract an entry in new ID to state that tracking Global Addresses may not be required or optimal on low end devices
  - We believe this is true on high end devices as well

SONATA and alias_sctp

- Cisco funded project to implement a NAT for SCTP
  - http://caia.swin.edu.au/urp/sonata
  - Kernel patches for ipfw2 and libalias
  - Initial June 18 release – FreeBSD 8
  - Upcoming release – FreeBSD 7+8

alias_sctp – Version 0.1

- Single homed private hosts to multi-homed public hosts
- IP Address forwarding
- Configurable logging levels – Kernel compile time
- Tested with approx. 10,000 concurrent flows for periods of up to 72 hours
- Logging of AbortM/ErrorM generating situations

alias_sctp – Next Release

- Sending of AbortM/ErrorM packets
- Per port IP Address forwarding
- Dynamically configurable
  - ☐ Log levels
  - ☐ Hash Table size
  - ☐ Timeouts
- ASCONF AddIP – Support for multi-homed private hosts
- This will be a fully functional SCTP NAT implementation

alias_sctp – Future Releases

- Dynamically configurable support for tracking Global IP Addresses
- Code Optimisation
- Fully tested with testing documentation
  - Functionality Testing
  - Performance Testing

Conclusions

- SCTP is a new Transport Layer Protocol with no SCTP support
- An Internet Draft is being developed to outline how to NAT SCTP packets using addresses, port numbers **and** vtags in the flow identification tuple
- The SONATA project is aiming for a public release of code to implement the draft
  - BSD Licensed
  - Fully tested – functionality and performance
  - Optional Global IP Address Tracking
- Initial prototype verifies that NAT for SCTP is possible