SWINBURNE UNIVERSITY OF TECHNOLOGY

Issues with Network Address Translation for SCTP

David Hayes

dahayes@swin.edu.au

Centre for Advanced Internet Architectures (CAIA)
Swinburne University of Technology

Outline

Overview of SCTP NAT

Alias_sctp data structures
- Lookup tables
- Timer Q

NAT state inconsistencies

Lookup table conflicts

Alias_sctp performance

Alias_sctp Interface

Some practical observations

Conclusions

Overview of SCTP NAT

Why is SCTP NAT different?

- Multiple control chunks

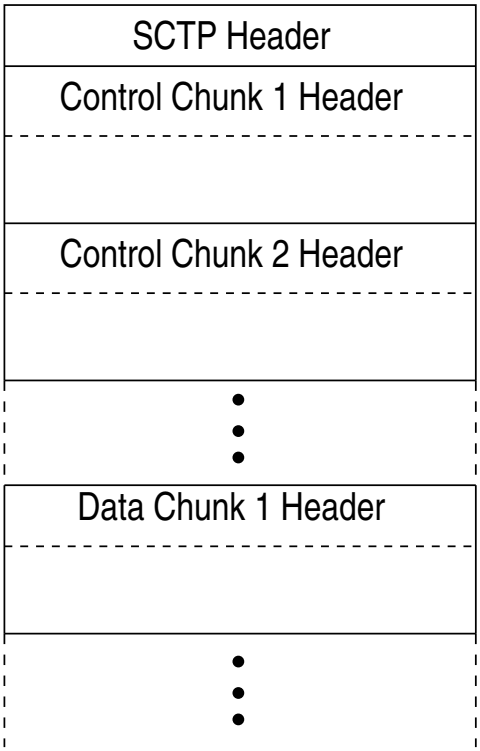

Overview of SCTP NAT

Why is SCTP NAT different?

- Multiple control chunks
  - And parameters
  - DoS – processing chunks and parameters

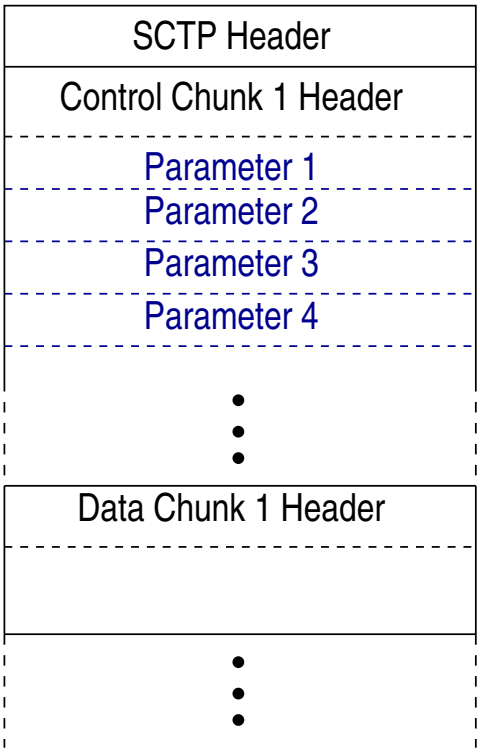

Overview of SCTP NAT

Why is SCTP NAT different?

- Multiple control chunks
  - And parameters
  - DoS – processing chunks and parameters
- Multi-homing

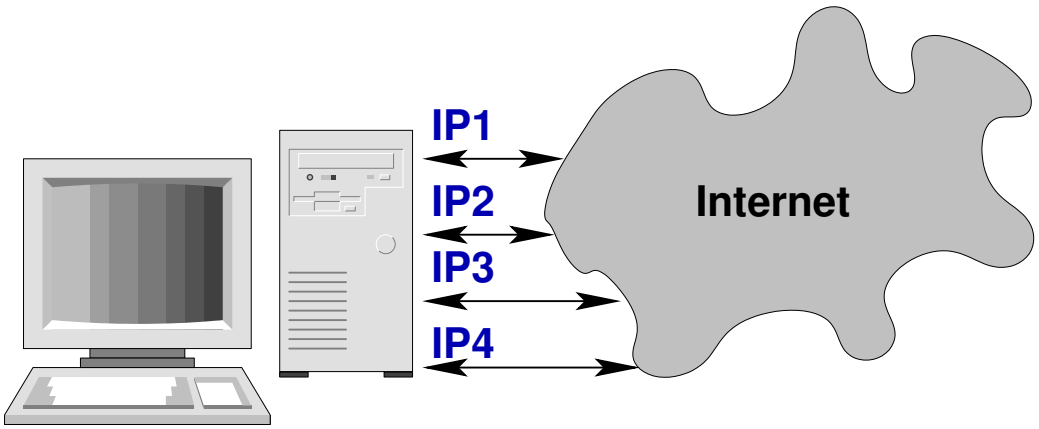

Overview of SCTP NAT

Why is SCTP NAT different?

- Multiple control chunks
  - And parameters
  - DoS – processing chunks and parameters
- Multi-homing
  - Can change on the fly
  - NAT state inconsistencies
  - DoS - Memory when tracking

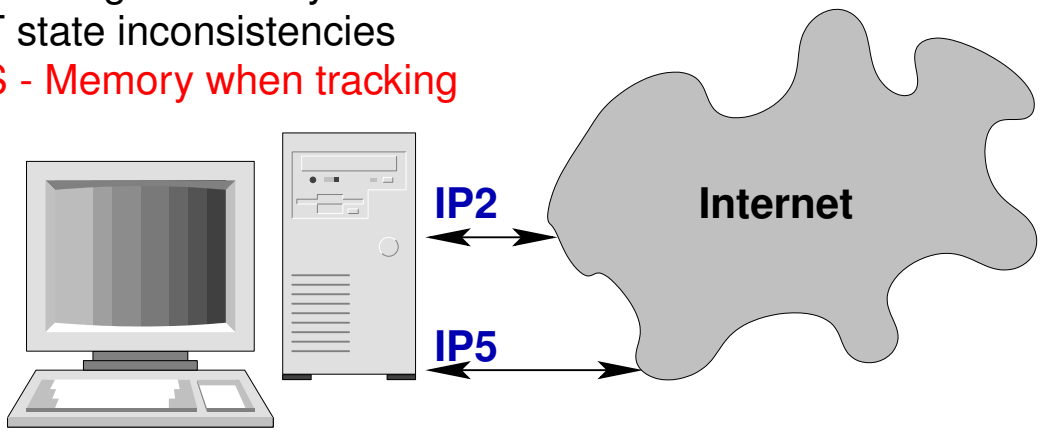

SWINBURNE UNIVERSITY OF TECHNOLOGY
Caia Seminar http://www.caia.swin.edu.au dahayes@swin.edu.au September 11, 2008 4

Overview of SCTP NAT

Gleaning association state information

- INIT ↔ INIT-Ack exchange

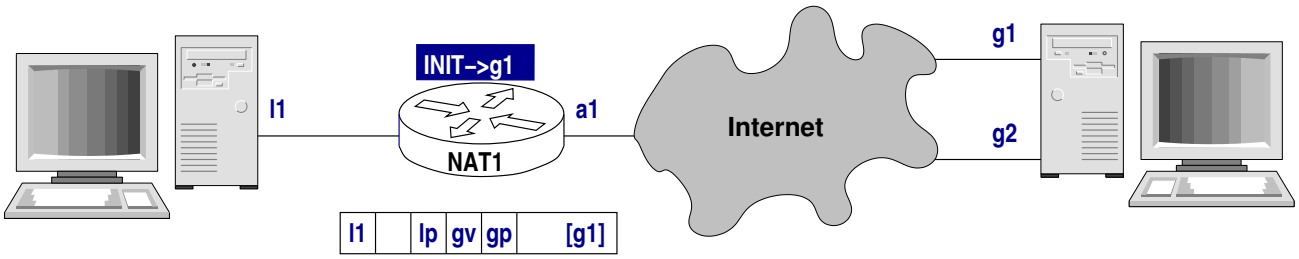

Local **Global**

Overview of SCTP NAT

Gleaning association state information

- INIT ↔ INIT-Ack exchange

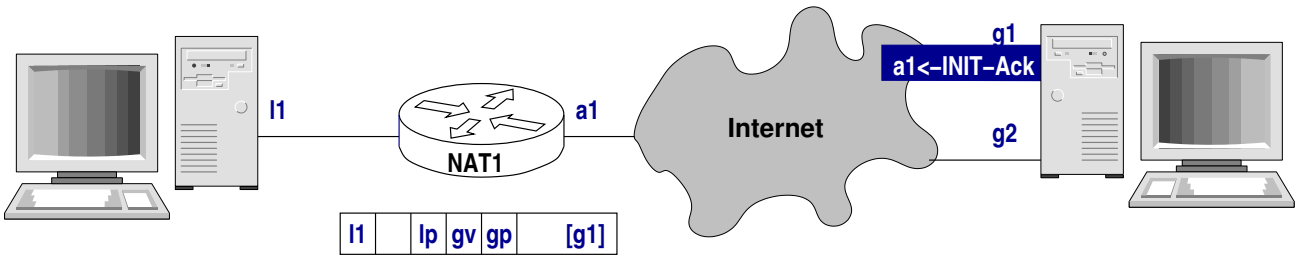

Local **Global**

SWINBURNE UNIVERSITY OF TECHNOLOGY
Caia Seminar http://www.caia.swin.edu.au dahayes@swin.edu.au September 11, 2008 5

Overview of SCTP NAT

Gleaning association state information

- INIT ↔ INIT-Ack exchange

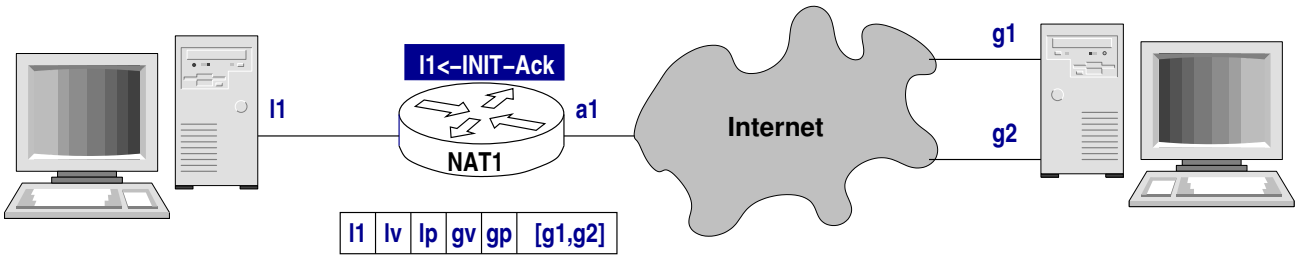

Local **Global**

Alias_sctp data structures

Lookup tables — Global and Local

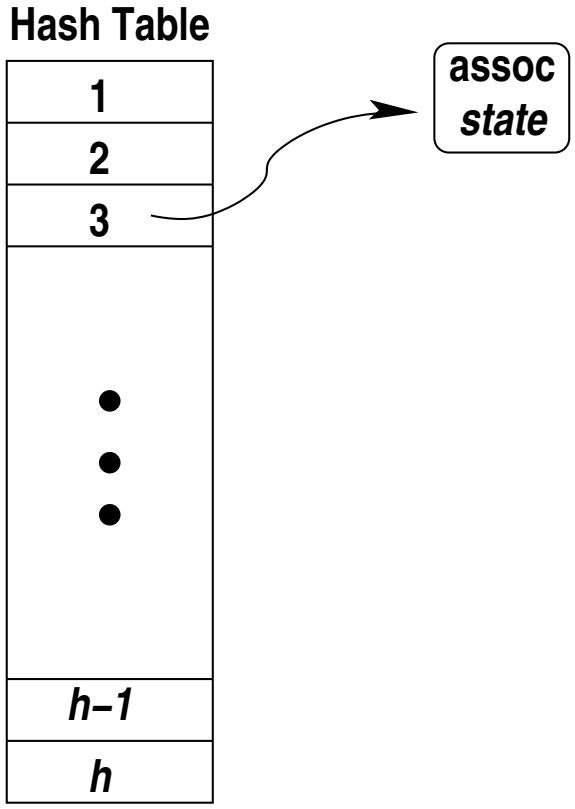

Association State

- local – vtag, port, address
- global – vtag, port, [address list]
- state – [Idle, · · · , UP, · · · , Closing]
- expiration time
- pointers for:
  - local and global lookup tables
  - timer Q

Alias_sctp data structures

Lookup tables — Global and Local

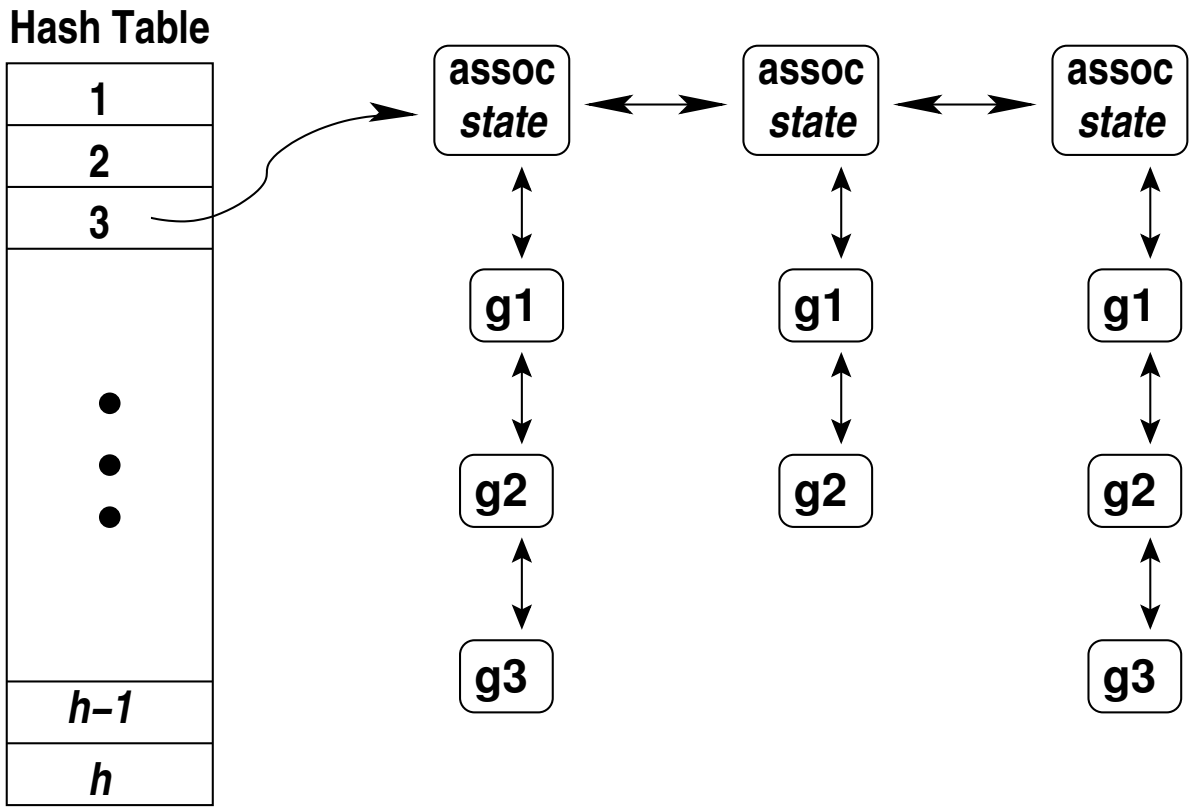

Alias_sctp data structures

Timer Queue

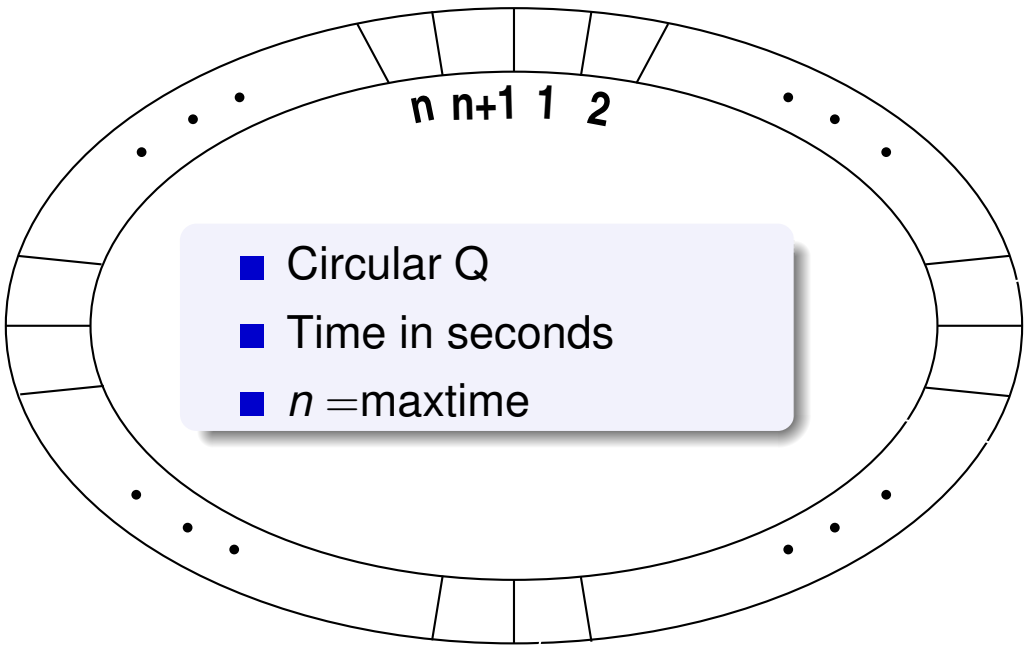

- Circular Q
- Time in seconds
- n =maxtime

Alias_sctp data structures
Timer Queue

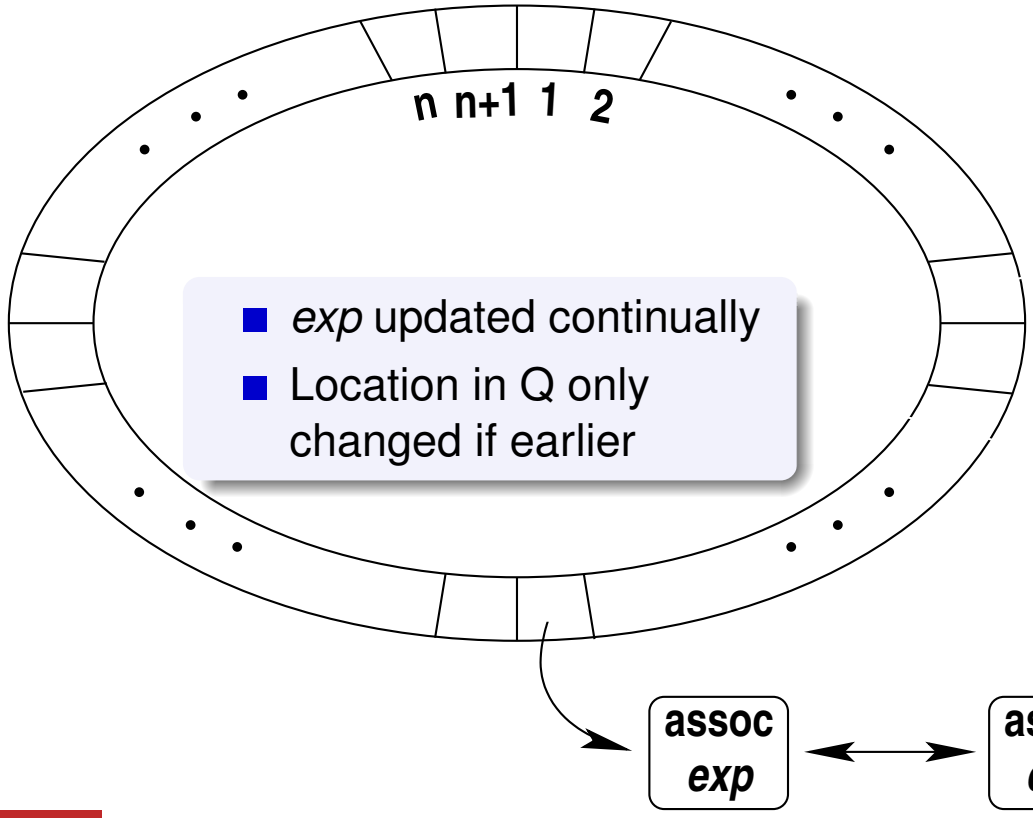

Alias_sctp data structures
Timer Queue

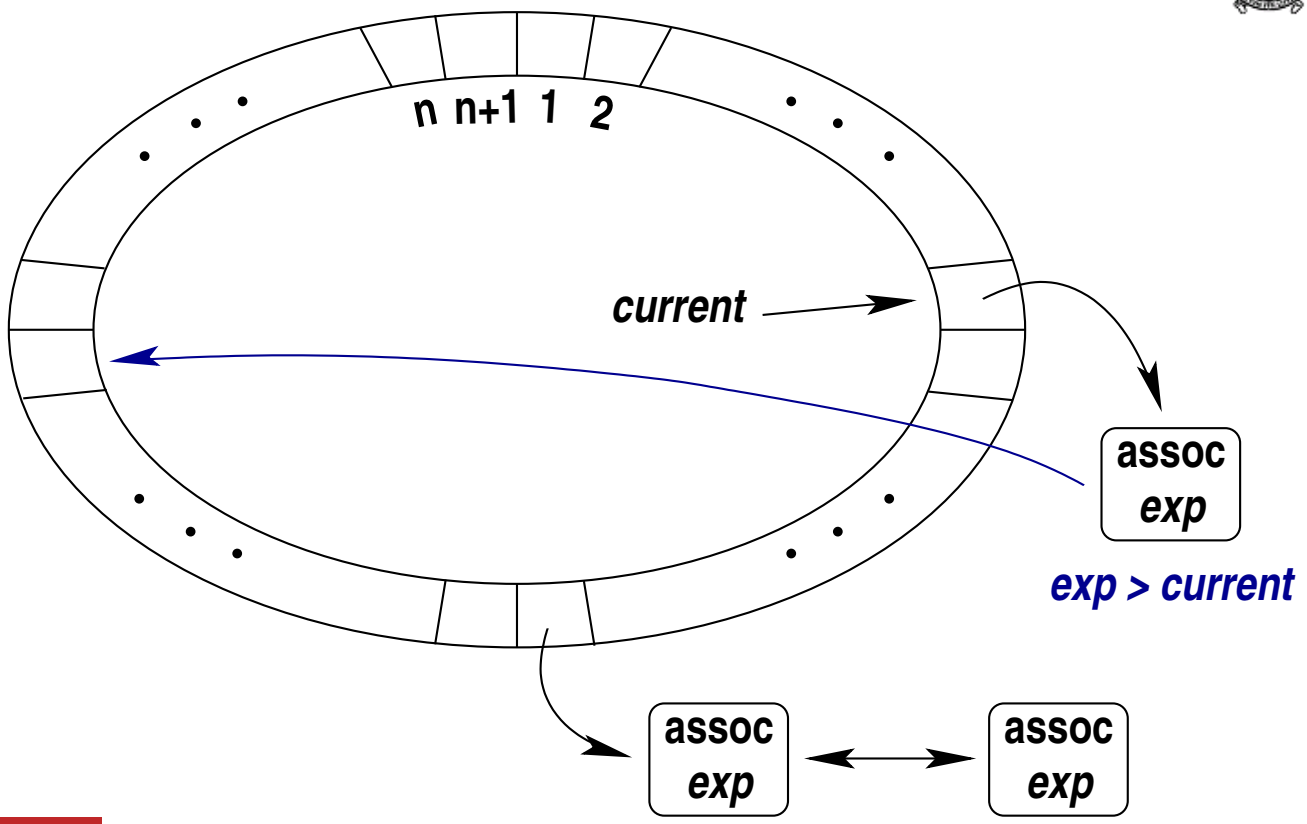

Caia Seminar http://www.caia.swin.edu.au dahayes@swin.edu.au September 11, 2008 10

NAT State Inconsistencies

When tracking global IP addresses

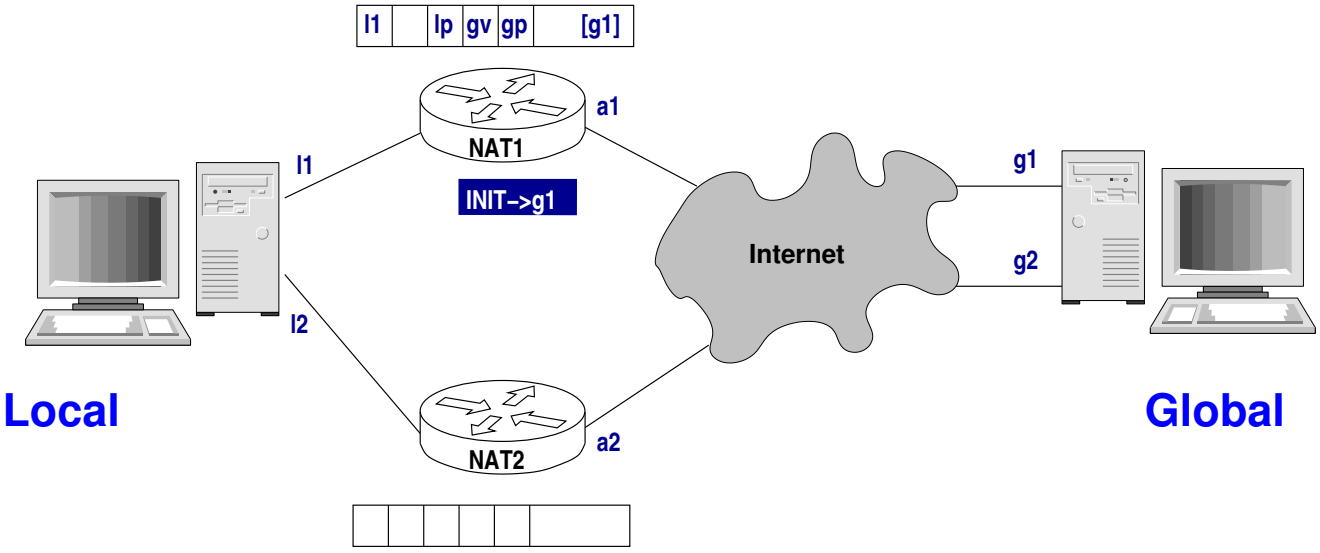

NAT State Inconsistencies

When tracking global IP addresses

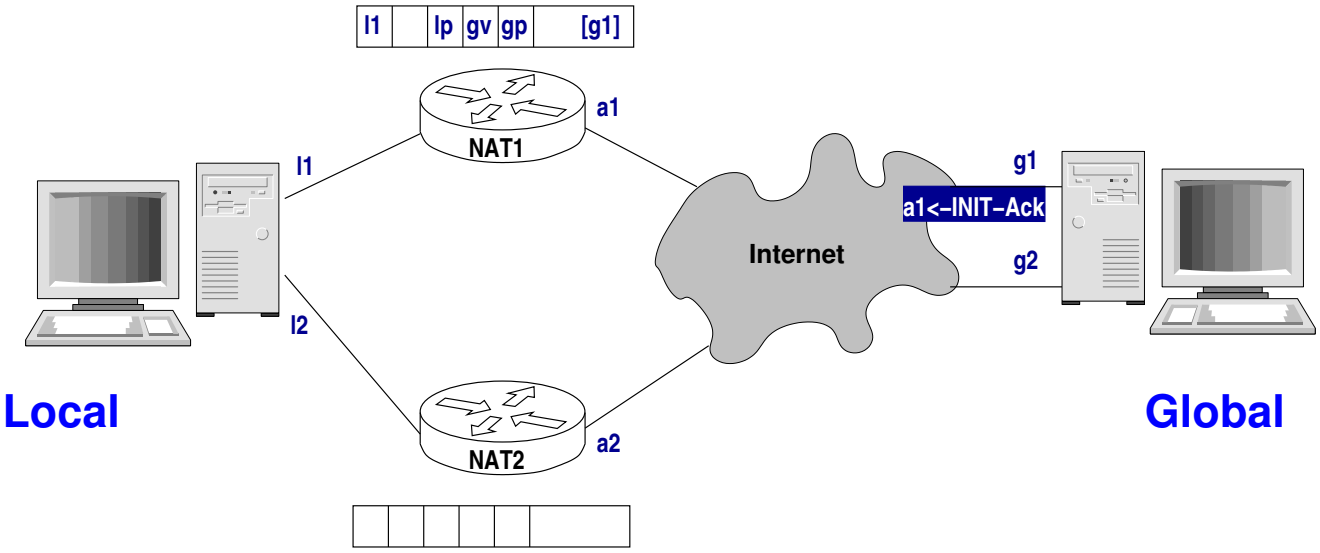

NAT State Inconsistencies

When tracking global IP addresses

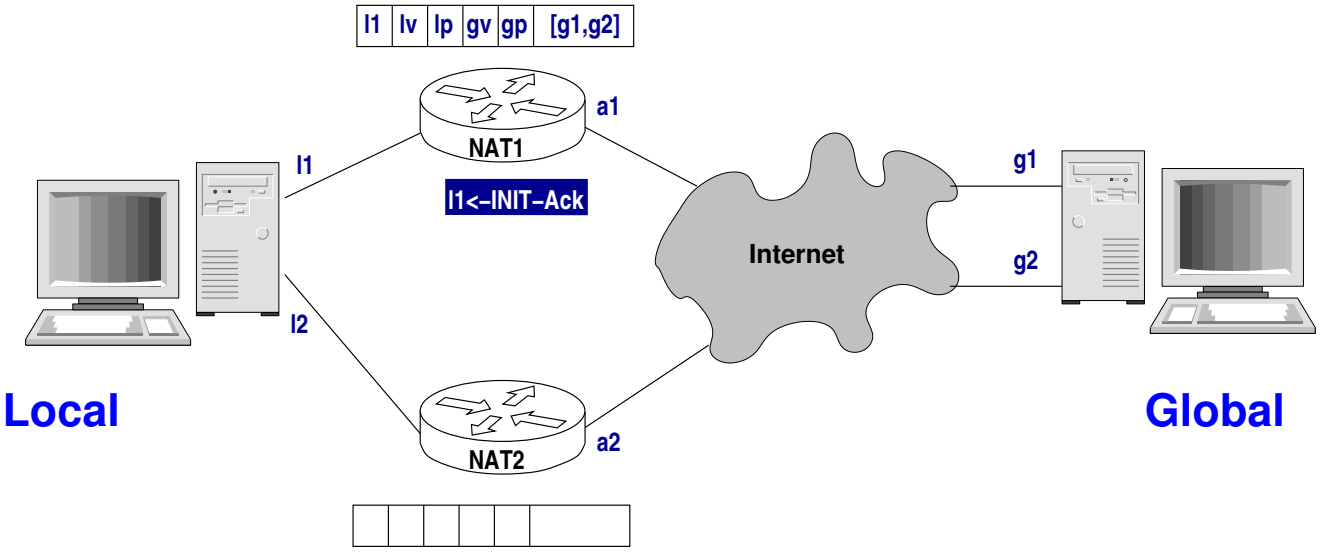

NAT State Inconsistencies

When tracking global IP addresses

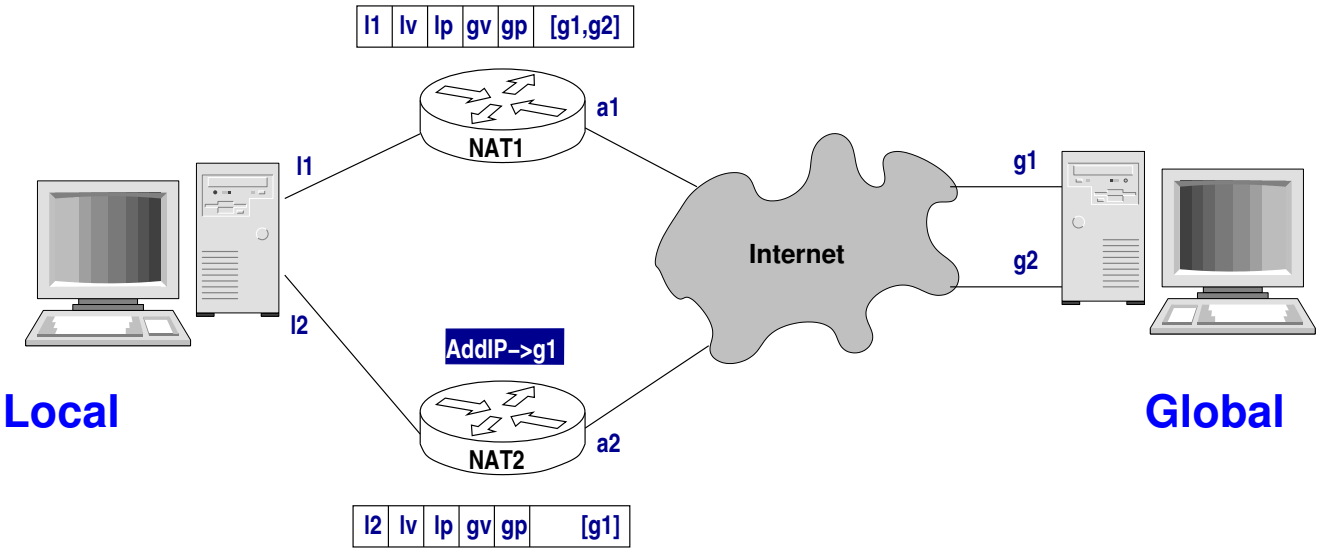

Caia Seminar http://www.caia.swin.edu.au dahayes@swin.edu.au September 11, 2008 12

NAT State Inconsistencies

When tracking global IP addresses

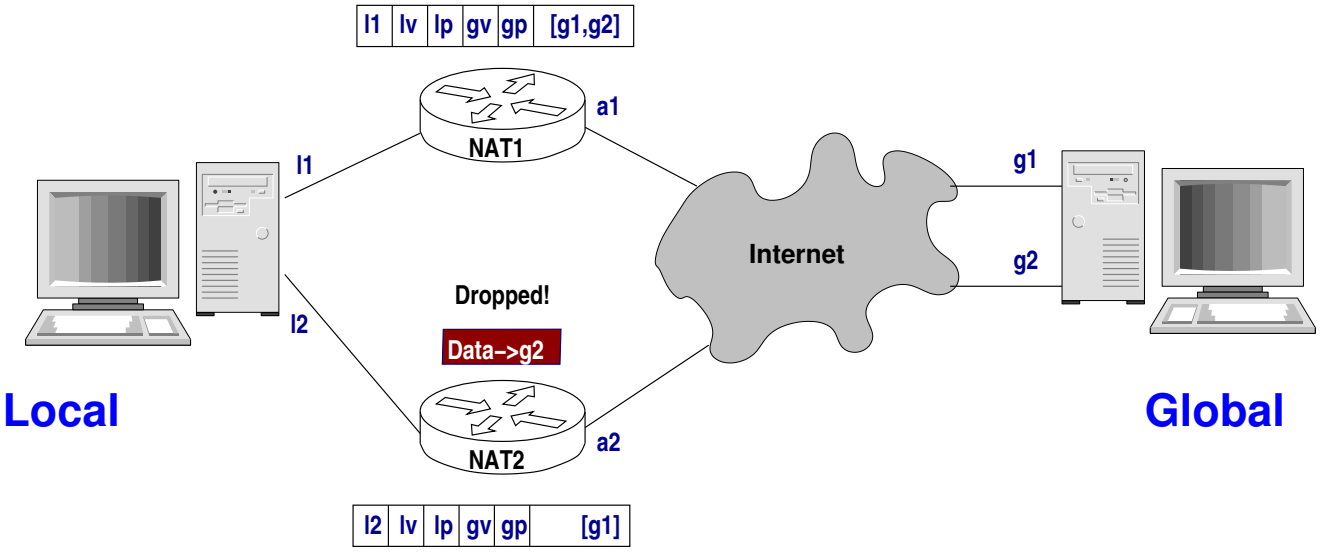

- There is no state inconsistency when global IP addresses are not tracked

Lookup table conflicts

General model of lookup table process

Is it necessary to track global IP addresses?

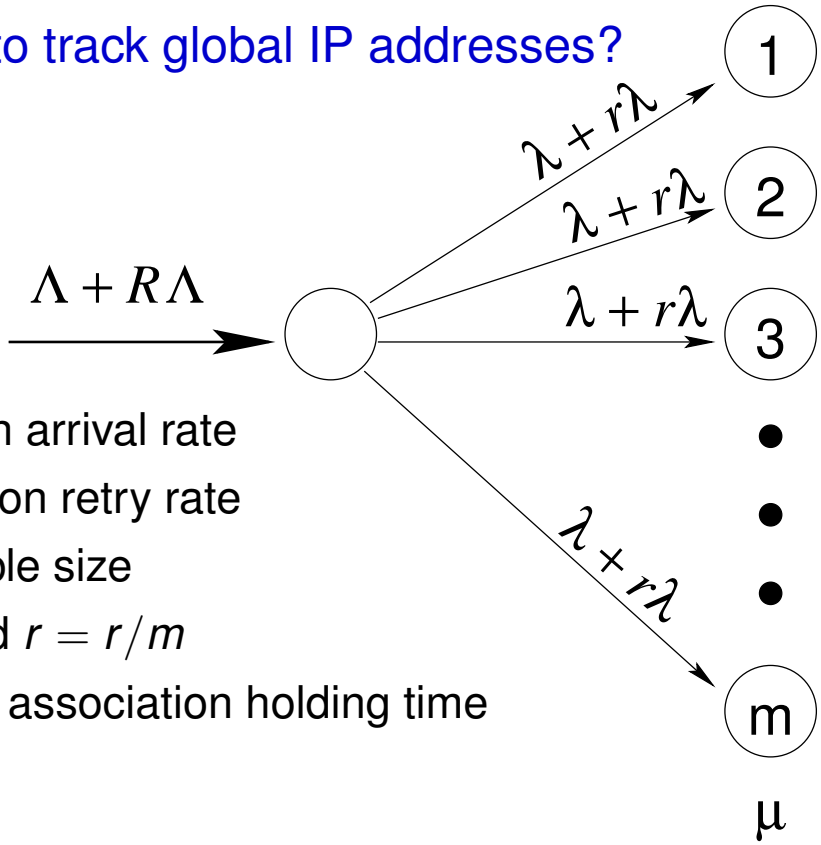

- Λ association arrival rate
- $R\Lambda$ association retry rate
- m lookup table size
- $\lambda = \Lambda/m$ and $r = r/m$
- $1/\mu$ average association holding time

Lookup table conflicts
Single element model

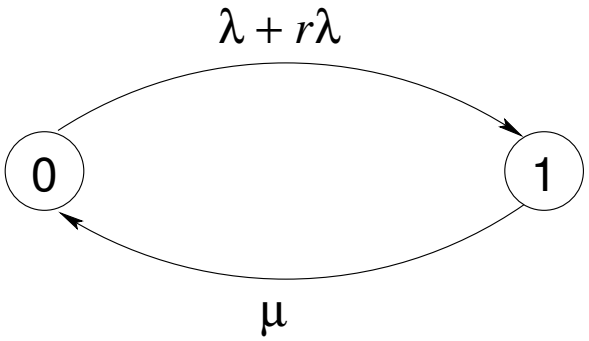

Solving the balance equations

$$P[\text{conflict}] = \mathrm{P}[1] = \left(1 + \left(\frac{\mu}{2\lambda}\right)^2\right)^{\frac{1}{2}} - \frac{\mu}{2\lambda}$$

Lookup table conflicts
Two scenarios

- No global IP address tracking

Vtag only

- All local hosts always use the same source and destination ports
- Conflict when two hosts coincidentally choose the same **vtag**
- $m = 2^{32} - 1$

vtag and port

- All local hosts always use the same destination port, but random source port.
- Conflict when two hosts coincidentally choose the same **vtag** and **port.**
- $m = (2^{32} - 1)(2^{16} - 1024)$

- Average association lasts 180 s $(1/\mu)$
- Association arrival rate varies

Lookup table conflicts

P[conflict] versus Λ for $1/\mu = 180\,s$

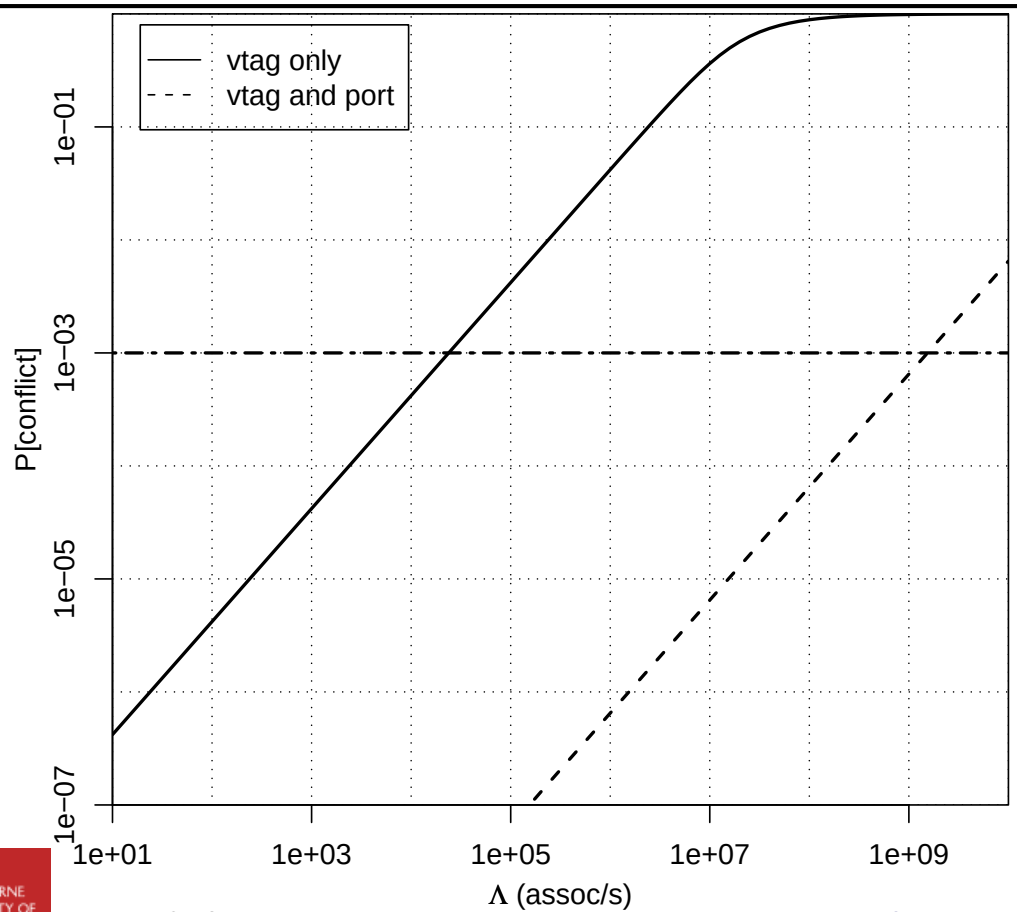

Alias_sctp performance

Experimental setup

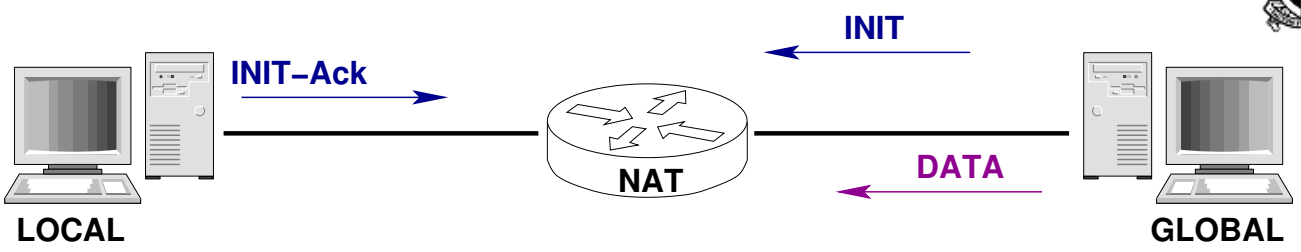

High association rate

- 1 500 assoc/s
- Assoc duration 10 s
- 4 addresses in INIT

High data rate

- 2000 associations
- 200×10^3 pkts/s
- 5 global addresses

- With and without global IP address tracking
- Baseline
- Measuring kernel process CPU time

Alias_sctp performance

Results — High association arrival rate

kernel cpu time	**Total**		`alias_sctp` (−base) (*s*)
	Mean (*s*)	Variance (s^2)	
tracking	177.46	0.062	33.56
not tracking	165.52	0.067	21.62
baseline	143.90	0.058	—
increase due to tracking	7.2 %	—	55.2 %

- Tracking significantly increases in alias_sctp's work
- Overall 7 %

Alias_sctp performance

Results — High data rate

kernel cpu time	**Total**		`alias_sctp` (−base) (*s*)
	Mean (*s*)	Variance (s^2)	
tracking	231.63	0.068	8.48
not tracking	229.61	0.063	6.46
baseline	223.15	0.130	—
increase due to tracking	0.9 %	—	31.3 %

- Not really significant

Alias_sctp performance
Memory usage

State space memory usage

$$M = 8h + 60n + 12\sum_{i=1}^{n} g_i \text{ bytes}$$

Memory usage	Tracking	Not tracking
High assoc rate	1.8×10^6 bytes	9.0×10^5 bytes
High packet rate	2.4×10^5 bytes	1.2×10^5 bytes

Alias_sctp Interface

- ipfw command
- sysctl variables (net.inet.ip.alias.sctp)
  - track_global_addresses
  - param_proc_limit
  - chunk_proc_limit
  - initialising_chunk_proc_limit
  - accept_global_ootb_addip
  - error_on_ootb
  - hashtable_size
  - holddown_timer
  - shutdown_timer
  - up_timer
  - init_timer
  - log_level d

Some practical observations

- freesbie – FreeBSD bootable disk
- cscope – Searchable C source tree database
- svn
- python
  - impacket

Conclusions

Alias_sctp module

- FreeBSD ipfw2 kernel based SCTP NAT
- extension to libalias kernel module
- ipfw and sysctl

SCTP NAT

- Multiple control streams
  - Limit number of parameters and chunks processed
- Global address tracking
  - Not necessary in any practical implementation
  - Memory usage significant
  - Disable by default
    - Prevents NAT state problems
  - Limit total number of addresses if enabled