



Stream Control Transport Protocol

Taarik Hassenmahomed

System and Network Engineering

February 4, 2009

What is SCTP

- connectie georiënteerd transport protocol
- werkt i.p.v. met byte stromen, met berichten
- eigen packet structuur
 - standaard packet header
 - data of control chunks

Bits	Bits 0 - 7	8 - 15	16 - 23	24 - 31
+0	Source port		Destination port	
32	Verification tag			
64	Checksum			
96	Chunk 1 type	Chunk 1 flags	Chunk 1 length	
128	Chunk 1 data			
...	...			
...	Chunk N type	Chunk N flags	Chunk N length	
...	Chunk N data			



SCTP vs TCP vs UDP

Features	SCTP	TCP	UDP
Datagrammen	yes	no	yes
Connection-oriented	yes	yes	no
Ordered data delivery	yes	yes	no
Unordered data delivery	yes	no	yes
Reach-ability check	yes	opt	no
Selective ACKs	yes	opt	no
Authentication	opt	opt	no
Multistreaming	yes	no	no
Multihoming	yes	no	no
Dynamic Multihoming	opt	no	no



Onderzoeks kader

- Simpele client server programma's
 - Multistreaming
 - Multihoming
 - Dynamic Multihoming



Onderzoeks vragen (1)

- Is het mogelijk om met SCTP twee verbindingen tegelijk te gebruiken voor load sharing met behulp van multihoming en multistreaming.



Onderzoeks vragen (2)

- Maakt het uit als een van de verbindingen een grotere latency heeft vanwege een langere roundtrip tijd
- Hoe verhoudt de performance van SCTP en TCP bij het gebruik van meerdere streams



Programmeer variant

- Sctp Kernel Implementation for FreeBSD 7.0
- Sctp Sockets API in C
- Sockets associaties ipv socket connecties
 - one2one
 - one2many



Multistreaming, hoe werkt het? (1)

CLIENT	SERVER
Connected to [192.168.138.2]	one-to-many: Waiting for associations ...



Multistreaming, hoe werkt het?(2)

CLIENT

Connected to [192.168.138.2]

There are 2 remote addresses and they are:

192.168.138.2

172.17.2.136

There are 2 local addresses and they are:

192.168.138.18

172.17.2.135

COMMUNICATION UP

assoc=0x28ced7cd, instr=3 outstr=3

SCTP ADDR CONFIRMED

SERVER

one-to-many: Waiting for associations ...

There are 2 remote addresses and they are:

192.168.138.18

172.17.2.135

There are 2 local addresses and they are:

192.168.138.2

172.17.2.136

COMMUNICATION UP

assoc=0x1fec53a3, instr=5 outstr=5

SCTP ADDR CONFIRMED



Multistreaming, hoe werkt het?(3)

CLIENT

```
Connected to [192.168.138.2]
There are 2 remote addresses and they are:
192.168.138.2
172.17.2.136
There are 2 local addresses and they are:
192.168.138.18
172.17.2.135
COMMUNICATION UP
assoc=0x28ced7cd, instr=3 outstr=3
SCTP ADDR CONFIRMED
<S>: [1]Send1 over stream 1
From str:1 sqn:0 [assoc:0x28ced7cd]:
[1]Send1 over stream 1
```

SERVER

```
one-to-many: Waiting for associations ...
There are 2 remote addresses and they are:
192.168.138.18
172.17.2.135
There are 2 local addresses and they are:
192.168.138.2
172.17.2.136
COMMUNICATION UP
assoc=0x1fec53a3, instr=5 outstr=5
SCTP ADDR CONFIRMED
Received on str:1 sqn:0 [assoc:0x1fec53a3]
Send on str:1 sqn:0 [assoc:0x1fec53a3]:
[1]Send1 over stream 1
```



Multistreaming, hoe werkt het?(4)

CLIENT

```
Connected to [192.168.138.2]
There are 2 remote addresses and they are:
192.168.138.2
172.17.2.136
There are 2 local addresses and they are:
192.168.138.18
172.17.2.135
COMMUNICATION UP
assoc=0x28ced7cd, instr=3 outstr=3
SCTP ADDR CONFIRMED
<S>: [1]Send1 over stream 1
From str:1 sqn:0 [assoc:0x28ced7cd]:
[1]Send1 over stream 1
<S>: [2]Send2 over stream 2
From str:2 sqn:0 [assoc:0x28ced7cd]:
[2]Send2 over stream 2
```

SERVER

```
one-to-many: Waiting for associations ...
There are 2 remote addresses and they are:
192.168.138.18
172.17.2.135
There are 2 local addresses and they are:
192.168.138.2
172.17.2.136
COMMUNICATION UP
assoc=0x1fec53a3, instr=5 outstr=5
SCTP ADDR CONFIRMED
Received on str:1 sqn:0 [assoc:0x1fec53a3]
Send on str:1 sqn:0 [assoc:0x1fec53a3]:
[1]Send1 over stream 1
Received on str:2 sqn:0 [assoc:0x1fec53a3]
Send on str:2 sqn:0 [assoc:0x1fec53a3]:
Received on str:2 sqn:1 [assoc:0x1fec53a3]
```



Multistreaming, hoe werkt het?(5)

CLIENT

```
Connected to [192.168.138.2]
There are 2 remote addresses and they are:
192.168.138.2
172.17.2.136
There are 2 local addresses and they are:
192.168.138.18
172.17.2.135
COMMUNICATION UP
assoc=0x28ced7cd, instr=3 outstr=3
SCTP ADDR CONFIRMED
<S>: [1]Send1 over stream 1
From str:1 sqn:0 [assoc:0x28ced7cd]:
[1]Send1 over stream 1
<S>: [2]Send2 over stream 2
From str:2 sqn:0 [assoc:0x28ced7cd]:
[2]Send2 over stream 2
<S>: [2]Send3 over stream 2
From str:2 sqn:1 [assoc:0x28ced7cd]:
[2]Send3 over stream 2
```

SERVER

```
one-to-many: Waiting for associations ...
There are 2 remote addresses and they are:
192.168.138.18
172.17.2.135
There are 2 local addresses and they are:
192.168.138.2
172.17.2.136
COMMUNICATION UP
assoc=0x1fec53a3, instr=5 outstr=5
SCTP ADDR CONFIRMED
Received on str:1 sqn:0 [assoc:0x1fec53a3]
Send on str:1 sqn:0 [assoc:0x1fec53a3]:
[1]Send1 over stream 1
Received on str:2 sqn:0 [assoc:0x1fec53a3]
Send on str:2 sqn:0 [assoc:0x1fec53a3]:
Received on str:2 sqn:1 [assoc:0x1fec53a3]
Send on str:2 sqn:1 [assoc:0x1fec53a3]:
Send on str:2 sqn:1 [assoc:0x1fec53a3]:
[2]Send3 over stream 2
```



Multistreaming, hoe werkt het?(6)

CLIENT

Connected to [192.168.138.2]
There are 2 remote addresses and they are:
192.168.138.2
172.17.2.136
There are 2 local addresses and they are:
192.168.138.18
172.17.2.135

COMMUNICATION UP

assoc=0x28ced7cd, instr=3 outstr=3

SCTP ADDR CONFIRMED

<S>: [1]Send1 over stream 1
From str:1 sqn:0 [assoc:0x28ced7cd]:
[1]Send1 over stream 1
<S>: [2]Send2 over stream 2
From str:2 sqn:0 [assoc:0x28ced7cd]:
[2]Send2 over stream 2
<S>: [2]Send3 over stream 2
From str:2 sqn:1 [assoc:0x28ced7cd]:
[2]Send3 over stream 2
<S>: exit

SHUTDOWN COMPLETE

assoc=0x28ced7cd, instr=3 outstr=3
Over !!

SERVER

one-to-many: Waiting for associations ...
There are 2 remote addresses and they are:
192.168.138.18
172.17.2.135
There are 2 local addresses and they are:
192.168.138.2
172.17.2.136

COMMUNICATION UP

assoc=0x1fec53a3, instr=5 outstr=5

SCTP ADDR CONFIRMED

Received on str:1 sqn:0 [assoc:0x1fec53a3]
Send on str:1 sqn:0 [assoc:0x1fec53a3]:
[1]Send1 over stream 1
Received on str:2 sqn:0 [assoc:0x1fec53a3]
Send on str:2 sqn:0 [assoc:0x1fec53a3]:
Received on str:2 sqn:1 [assoc:0x1fec53a3]
Send on str:2 sqn:1 [assoc:0x1fec53a3]:
Send on str:2 sqn:1 [assoc:0x1fec53a3]:
[2]Send3 over stream 2
SCTP SHUTDOWN EVENT: assoc=0x1fec53a3

SHUTDOWN COMPLETE

assoc=0x1fec53a3, instr=5 outstr=3
Waiting for associations ...



Multistreaming:Wireshark

- DATA chunk(ordered, complete segment, TSN: 3278131290, SID: 1, SSN: 0, PPID: 0, payload length: 9 bytes)
 - TSN: 3278131290
 - Stream Identifier: 0x0001
 - Stream sequence number: 0




Multistreaming:Wireshark



- DATA chunk(ordered, complete segment, TSN: 3278131290, SID: 1, SSN: 0, PPID: 0, payload length: 9 bytes)
 - TSN: 3278131290
 - Stream Identifier: 0x0001
 - Stream sequence number: 0
- DATA chunk(ordered, complete segment, TSN: 3278131291, SID: 2, SSN: 0, PPID: 0, payload length: 9 bytes)
- DATA chunk(ordered, complete segment, TSN: 3278131292, SID: 2, SSN: 1, PPID: 0, payload length: 9 bytes)




Multistreaming: Head-of-Line



(a) A single TCP connection





(b) One TCP connection per file



(c) A single SCTP association with several Streams

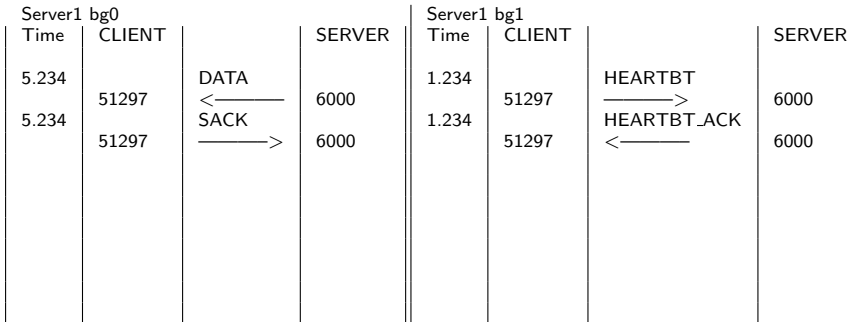


Multihoming Topologie



Multihoming Voorbeeld (1.1)

- blokkeren van uitgaande pakketten van de primaire IP verbinding vanaf de client na 20 seconden



Multihoming Voorbeeld (1.2)

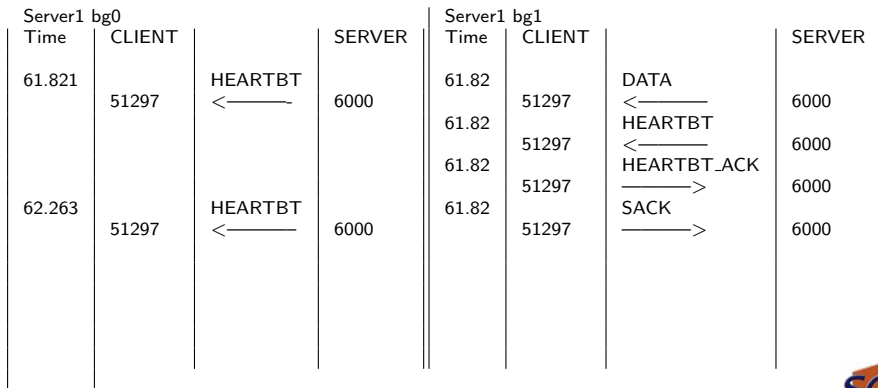
- blokkeren van uitgaande pakketten van de primaire IP verbinding vanaf de client na 20 seconden

Server1 bg0				Server1 bg1			
Time	CLIENT		SERVER	Time	CLIENT		SERVER
5.234	51297	DATA ←————	6000	1.234	51297	HEARTBT ————→	6000
5.234	51297	SACK ————→	6000	1.234	51297	HEARTBT_ACK ←————	6000
22.483	51297	DATA ←————	6000	22.481	51297	DATA ————→	6000
24.482	51297	DATA ←————	6000	22.482	51297	SACK ————→	6000
36.246	51297	HEARTBT ←————	6000	38.956	51297	SACK ————→	6000
38.957	51297	DATA ←————	6000	40.956	51297	DATA ←————	6000



Multihoming Voorbeeld (2.1)

- deblokken van uitgaande pakketten van de primaire IP verbinding vanaf de client na 62 seconden



Multihoming Voorbeeld (2.2)

- deblokkeren van uitgaande pakketten van de primaire IP verbinding vanaf de client na 62 seconden

Server1 bg0				Server1 bg1			
Time	CLIENT		SERVER	Time	CLIENT		SERVER
61.821	51297	HEARTBT ←	6000	61.82	51297	DATA ←	6000
				61.82	51297	HEARTBT ←	6000
				61.82	51297	HEARTBT_ACK →	6000
62.263	51297	HEARTBT ←	6000	61.82	51297	SACK →	6000
69.263	51297	HEARTBT_ACK →	6000	69.29	51297	HEARTBT →	6000
75.642	51297	DATA ←	6000	69.29	51297	HEARTBT_ACK ←	6000
75.642	51297	SACK DATA →	6000	102.41	51297	HEARTBT →	6000



Dynamic multihoming

- SCTP Socket API functions nodig
 - SctpBindx
 - SctpConnectx
 - SctpSetPeerPrimaryAddr



Conclusie

- Met de huidige versie van SCTP is gelijktijdig gebruik van twee verbindingen voor load sharing niet mogelijk.
 - maar 1 verbinding actief
 - gebruik `SCTP_SET_PEER_PRIMARY_ADDR`
- Bij verbindingen met een langere roundtrip tijd verstuurt SCTP pakketten met een langzame snelheid, net als TCP.
- Er zijn betere tools nodig voor het testen van de prestaties SCTP
 - netperf ondersteund SCTP, maar test geen multistreaming.

