Building Fault-tolerant Site-to-Site VPNs with Cisco ASA

Oleg Tipisov
Customer Support Engineer, Cisco TAC

Aug 29, 2013. Revision 1.0
Cisco Public

Cisco Support Community – Expert Series Webcasts in Russian

Сегодня на семинаре эксперт Cisco Олег Типисов расскажет об основных понятиях и аспектах построения отказоустойчивых Site-to-Site VPN на ASA

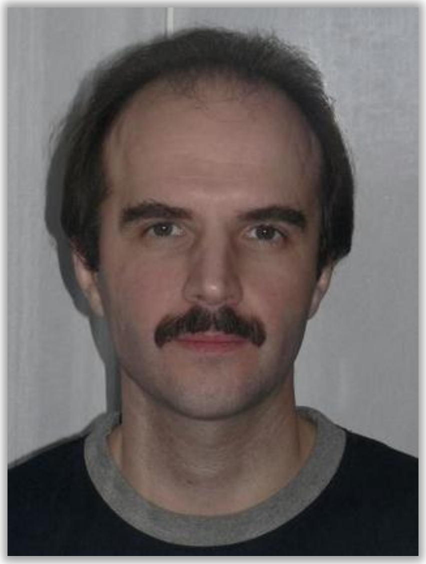

Олег Типисов

Инженер центра технической поддержки Cisco TAC в Москве

Спасибо, что посетили наш семинар сегодня

Сегодняшняя презентация включает опросы аудитории

Пожалуйста, участвуйте!

Скачать презентацию вы можете по ссылке:

https://supportforums.cisco.com/docs/DOC-36123

Присылайте Ваши вопросы!

Используйте Q&A панель, чтобы послать вопрос. Наши эксперты ответят на них

Опрос #1
Каков уровень ваших знаний о VPN

- Я знаю, что это такое, но работать с ними не приходилось
- Я использую VPN как пользователь
- Я являюсь администратором VPN в компании
- Мне приходилось проектировать и администрировать различные VPN на различных платформах

Introduction

- This presentation is about ASA Site-to-Site VPNs
- In TAC we see spike of customer cases where customers try to configure redundant Site-to-Site VPNs over multiple ISPs
- We will not discuss Remote Access VPNs
- Other platforms are also beyond the scope of this presentation
- Students are expected to understand ASA CLI including ASA VPN configuration commands

Agenda

- Failure types and common topologies
- Failover
- Ingredients of ISP redundancy
- Scenario: Dual ASA – Dual ISP
- Scenario: Single ASA – Dual ISP
- Connections creation and teardown
- OSPF over tunnels
- Conclusion

Failure Types and Common Topologies

Failure Types

- Device failure

  Can be mitigated by failover

- Interface failure

  Can also be mitigated by failover

- ISP failure

  Can be mitigated by dynamic routing or static routing + IP SLA + tracking

Common Topologies

- Failover

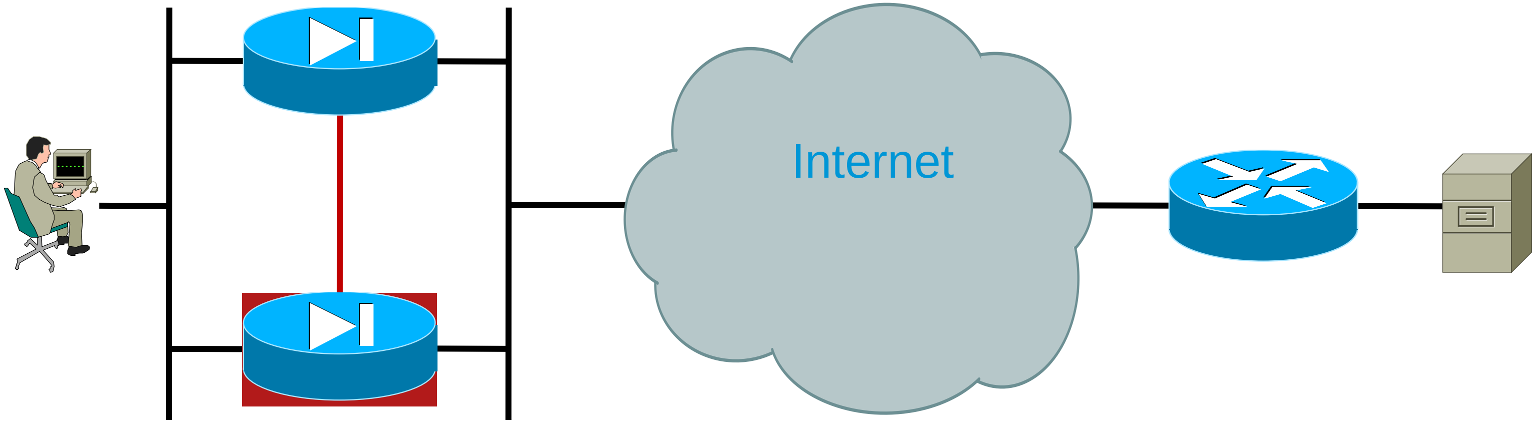

Common Topologies

- Dual ASA - Dual ISP topology

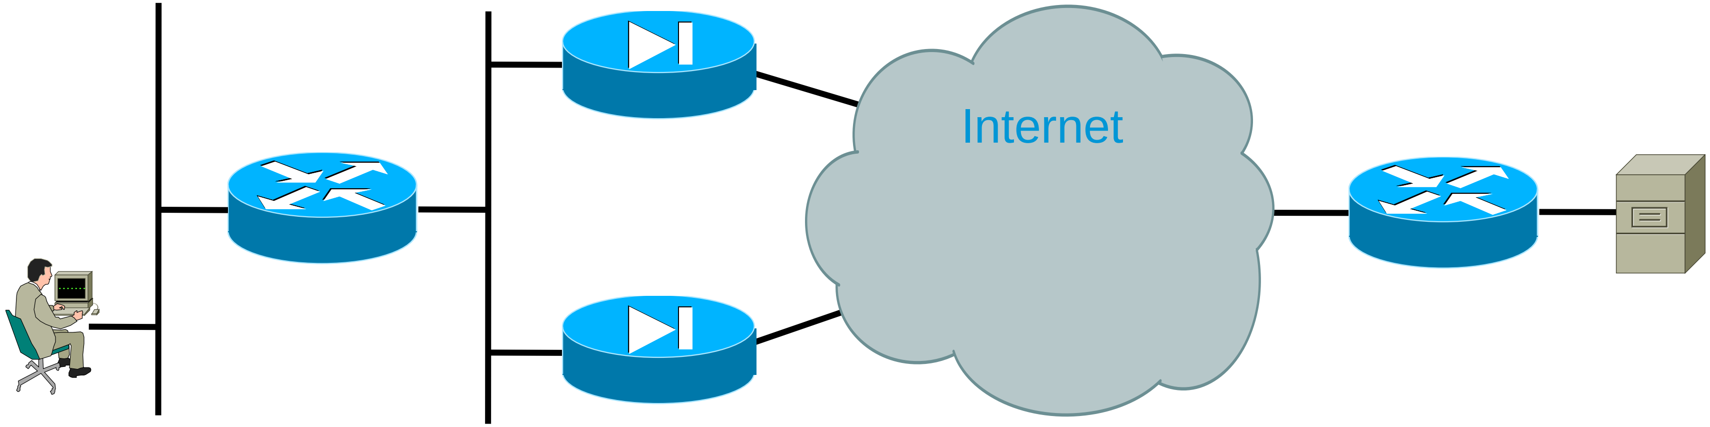

Common Topologies

- Single ASA - Dual ISP topology

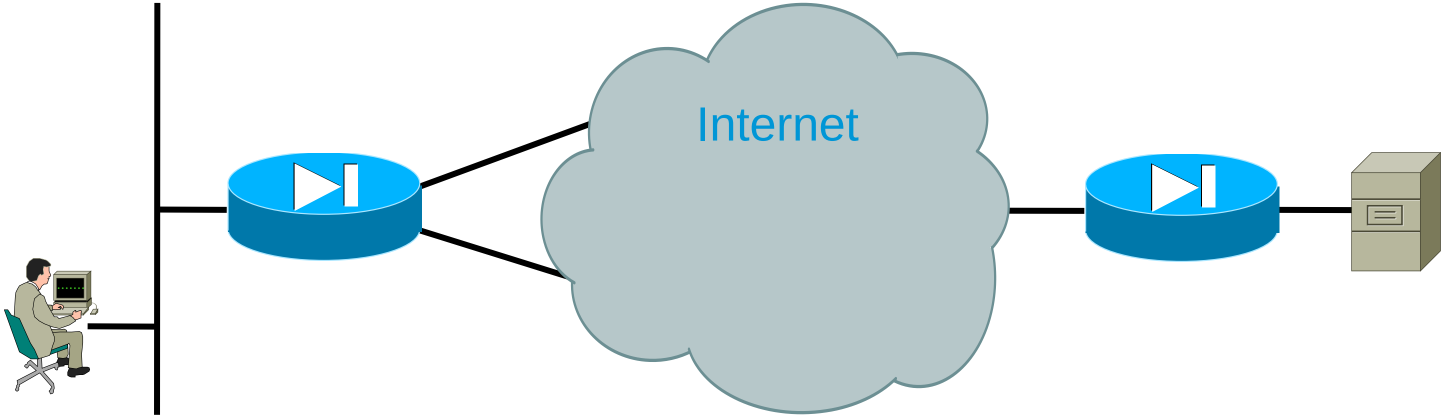

Failover

Failover

- Failover prevents traffic disruption from device or interface failures
- Failover can detect and react to failures in sub-second time
- Connection information and VPN tunnels are replicated between devices
- Configuration is also replicated

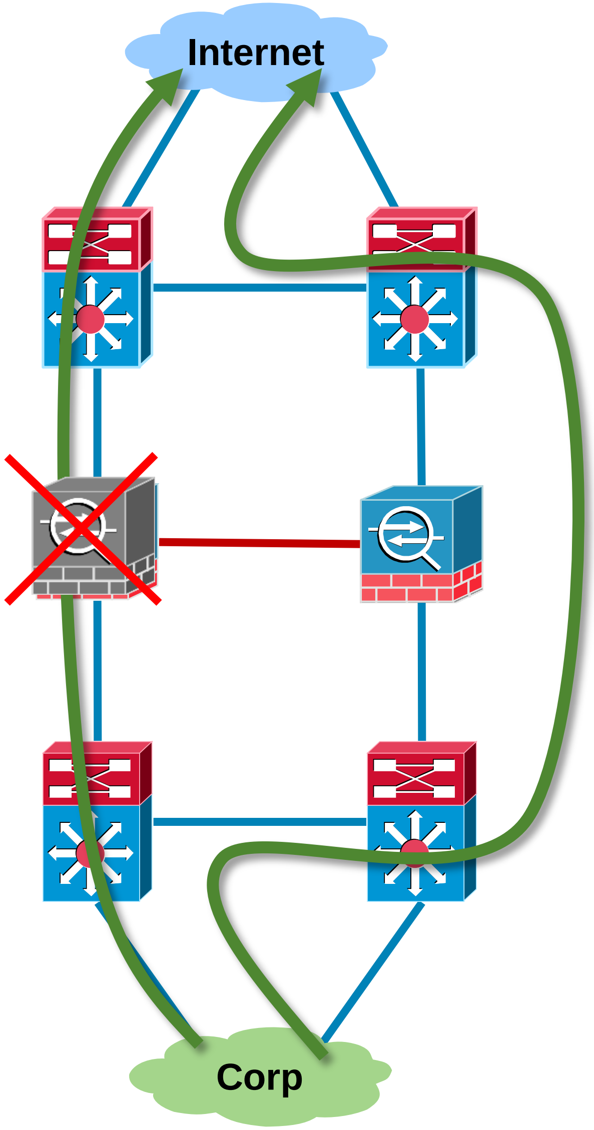

Failover Terminology

- A/S vs. A/A Failover
- Primary & Secondary units
- Active & Standby units
- Failover link & State link

- Firewalls swap MACs and IPs when failover occurs

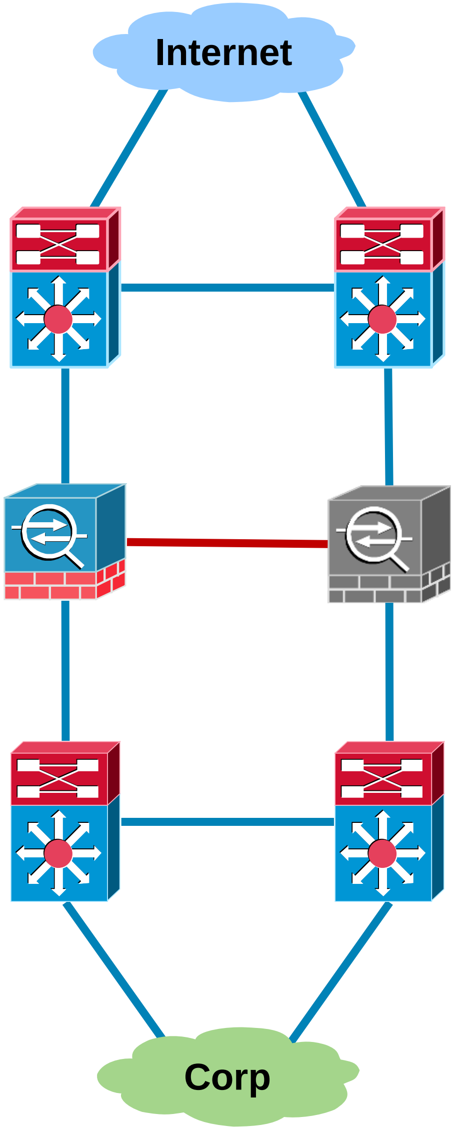

ASA 9.0 Enhancements

- Dynamic routing protocols in multiple context mode: EIGRP, OSPFv2
- L2L VPNs in multiple context mode
- This means that fault-tolerant Site-to-Site VPNs are now possible in multiple context mode

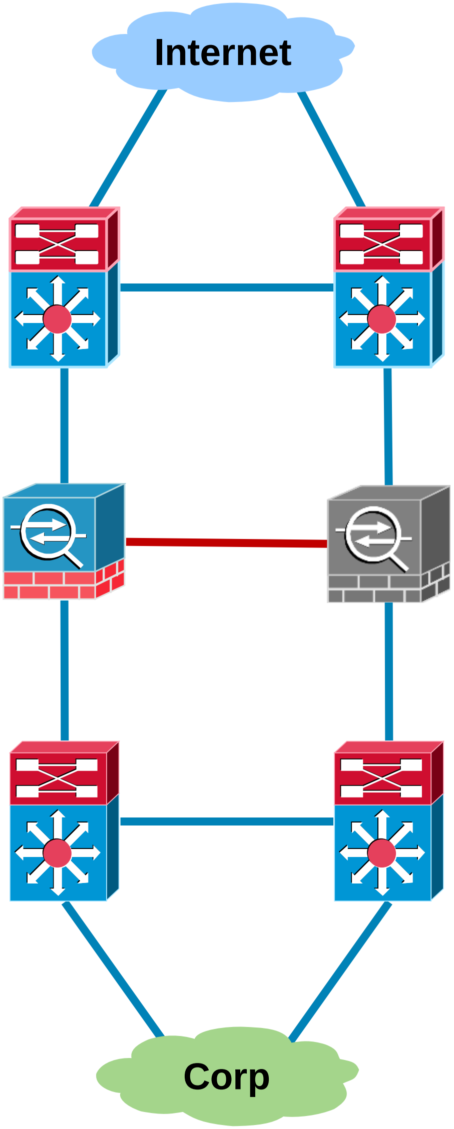

ASA 9.0 VPN Support in Multiple Mode

- IKEv1 & IKEv2
- L2L only (RA not supported at this time)
- Static Site-to-Site VPN/static crypto maps
- Dynamic Site-to-Site VPN/dynamic crypto maps
- Backup and Multi-peer support for IKEv1
- IPv6 Site-to-Site VPN
- Next Generation Encryption for IKEv2
- SNMP support
- Can be used with A/S or A/A failover

VPN in Multiple Mode Typical Problem

- ASA cannot initiate or accept tunnels with the message:

```
%ASA-5-713239: Group = 213.1.1.1, IP = 213.1.1.1, Tunnel Rejected: The
maximum tunnel count allowed has been reached
```

- This is caused by the fact that contexts do not have VPN licenses by default

VPN in Multiple Mode Typical Problem

- We can see in the context…

```
ASA/act/VPN/pri# show vpn-sessiondb license-summary
…
---------------------------------------------------------------------------
VPN Licenses and Configured Limits Summary
---------------------------------------------------------------------------
                                     Status : Installed :    Burst :  Limit
                                  -----------------------------------------
Other VPN (Available by Default) :  ENABLED :         0 :        0 :   NONE
VPN-3DES-AES                     :  ENABLED
VPN-DES                          :  ENABLED
---------------------------------------------------------------------------


---------------------------------------------------------------------------
VPN Licenses Usage Summary
---------------------------------------------------------------------------
                          Local : Shared :   All  :   Peak :  Eff.  :
                         In Use : In Use : In Use : In Use :  Limit : Usage
                       ----------------------------------------------------
Other VPN              :                 :      0 :      0 :      0 :    0%
  Site-to-Site VPN     :                 :      0 :      0          :    0%
---------------------------------------------------------------------------
```

VPN in Multiple Mode Typical Problem

- … while 2500 tunnels license is seen in system execution space:

```
ASA/act/pri# show vpn-sessiondb license-summary
---------------------------------------------------------------------------
VPN Licenses and Configured Limits Summary
---------------------------------------------------------------------------
                                     Status : Capacity : Installed :  Limit
                                  -----------------------------------------
…
Other VPN (Available by Default) :  ENABLED :     2500 :      2500 :   NONE
…
---------------------------------------------------------------------------


---------------------------------------------------------------------------
VPN Licenses Usage Summary
---------------------------------------------------------------------------
                          Local : Shared :   All  :   Peak :  Eff.  :
                         In Use : In Use : In Use : In Use :  Limit : Usage
                       ----------------------------------------------------
Other VPN              :                 :      0 :      0 :   2500 :    0%
  Site-to-Site VPN     :                 :      0 :      0          :    0%
---------------------------------------------------------------------------
```

VPN in Multiple Mode Typical Problem

- This can be corrected by configuring in system execution space:

```
class test
  limit-resource VPN Other 80.0%

context VPN
  member test
  allocate-interface GigabitEthernet0/0.98
  allocate-interface GigabitEthernet0/0.103
  config-url disk0:/VPN.cfg
  join-failover-group 2

ASA/act/pri# show resource usage context VPN
Resource               Current        Peak      Limit        Denied Context
Syslogs [rate]               0          78  unlimited             0 VPN
Conns                        0           1  unlimited             0 VPN
Hosts                        0           2  unlimited             0 VPN
Conns [rate]                 0           1  unlimited             0 VPN
Inspects [rate]              0           2  unlimited             0 VPN
Routes                       3           3  unlimited             0 VPN
Other VPN Sessions           4           4       2000             0 VPN
```

VPN in Multiple Mode Troubleshooting

- VPN configuration in context remains the same as before
- Use the following commands for troubleshooting:

```
ASA/act/VPN/pri# show vpn-sessiondb license-summary
…
---------------------------------------------------------------------------
VPN Licenses Usage Summary
---------------------------------------------------------------------------
                          Local : Shared :   All  :   Peak :  Eff.  :
                         In Use : In Use : In Use : In Use :  Limit : Usage
                       ----------------------------------------------------
Other VPN              :                 :      1 :      1 :   2000 :    0%
  Site-to-Site VPN     :                 :      1 :      1          :    0%
---------------------------------------------------------------------------
```

VPN in Multiple Mode Troubleshooting

- VPN configuration in context remains the same as before
- Use the following commands for troubleshooting:

```
ASA/act/VPN/pri# show vpn-sessiondb summary
---------------------------------------------------------------------------
VPN Session Summary
---------------------------------------------------------------------------
                               Active : Cumulative : Peak Concur : Inactive
                             ----------------------------------------------
Site-to-Site VPN             :      1 :          1 :           1
  IKEv1 IPsec                :      1 :          1 :           1
---------------------------------------------------------------------------
Total Active and Inactive    :      1             Total Cumulative :      1
---------------------------------------------------------------------------
```

VPN in Multiple Mode Troubleshooting

- VPN configuration in context remains the same as before
- Use the following commands for troubleshooting:

```
ASA/act/VPN/pri# show failover
…
Stateful Failover Logical Update Statistics
…
        Stateful Obj    xmit        xerr        rcv         rerr
        VPN IKEv1 SA    1           0           0           0
        VPN IKEv1 P2    1           0           0           0
        VPN IKEv2 SA    0           0           0           0
        VPN IKEv2 P2    0           0           0           0
        VPN CTCP upd    0           0           0           0
        VPN SDI upd     0           0           0           0
        VPN DHCP upd    0           0           0           0
```

VPN in Multiple Mode Troubleshooting

- VPN configuration in context remains the same as before
- Use the following commands for troubleshooting:

```
ASA/stby/VPN/sec# show crypto ikev1 sa

IKEv1 SAs:

  Standby SA: 1
    Rekey SA: 0 (A tunnel will report 1 Standby and 1 Rekey SA during rekey)
Total IKE SA: 1

1   IKE Peer: 213.1.1.1
    Type    : L2L             Role    : initiator
    Rekey   : no              State   : MM_STANDBY

ASA/stby/VPN/sec# show crypto ipsec sa
…
ASA/stby/VPN/sec# show vpn-sessiondb detail l2l
…
```

Stateful Failover for VPN

- VPN failover is stateful. Two pings are lost when active ASA is rebooted

```
failover polltime unit 1 holdtime 3
```

```
Aug 21 2013 10:14:30 system : %ASA-1-103001: (Secondary) No response from
other firewall (reason code = 4).
Aug 21 2013 10:14:30 system : %ASA-1-104001: (Secondary) Switching to ACTIVE -
HELLO not heard from mate.
Aug 21 2013 10:14:30 system : %ASA-1-104001: (Secondary_group_1) Switching to
ACTIVE - HELLO not heard from mate.
Aug 21 2013 10:14:30 system : %ASA-1-104001: (Secondary_group_2) Switching to
ACTIVE - HELLO not heard from mate.
Aug 21 2013 10:14:30: %ASA-6-720039: (VPN-Secondary) VPN failover client is
transitioning to active state
Aug 21 2013 10:14:30: %ASA-6-713905: IKE port 10000 for IPSec UDP already
reserved on interface outside
Aug 21 2013 10:14:30: %ASA-6-713905: IKE port 10000 for IPSec UDP already
reserved on interface inside
Aug 21 2013 10:14:30: %ASA-5-713120: Group = 213.1.1.1, IP = 213.1.1.1, PHASE
2 COMPLETED (msgid=b4ce0471)
```

Опрос #2

Опрос #2
Ваше мнение о реализации VPN на ASA

- Отлично работает, все устраивает
- Работает хорошо, но необходимо добавить возможности, имеющиеся на маршрутизаторах: GRE, DMVPN и т.д.
- Работает, но приходилось сталкиваться с багами. Необходимо также улучшить функциональность: GRE, DMVPN и т.д.
- На мой взгляд ASA вообще не предназначена для Site-to-Site VPN и исправить это нельзя, поскольку это не роутер. Поэтому использую маршрутизаторы Cisco

Ingredients of ISP Redundancy

Ingredients of ISP Redundancy

- Logical interfaces, GRE/VTI are not supported on ASA. Crypto maps is the only way to configure VPNs on ASA. Static crypto maps are not very scalable
- Dynamic routing over tunnels is possible with OSPF, but does not work with routers and has many other limitations
- Static routing requires IP SLA and tracking, which also have known limitations. BGP is not yet implemented
- ASA supports
  - Static crypto maps & dynamic crypto maps
  - Dead Peer Detection (DPD)
  - Reverse Route Injection (RRI)
  - Redistribution of RRI routes to dynamic routing protocols

Scenario: Dual ASA – Dual ISP

Topology

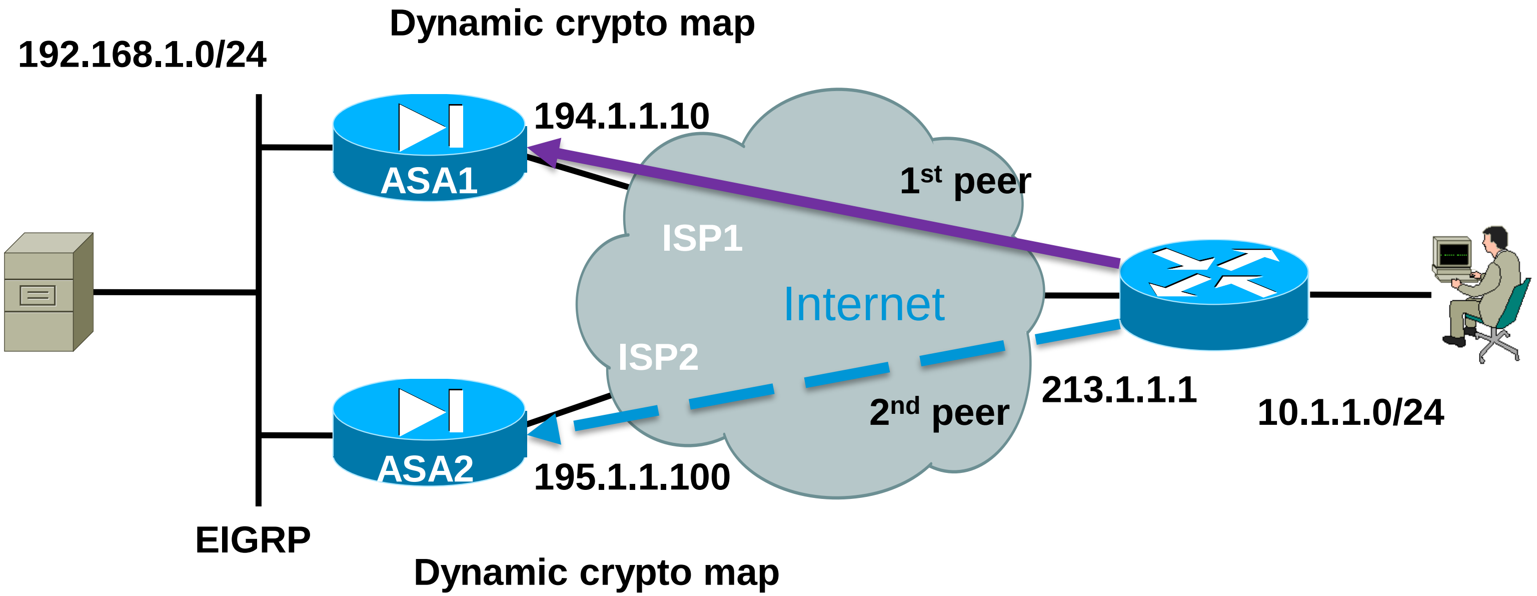
Dynamic crypto map
192.168.1.0/24
194.1.1.10
ASA1
1st peer
ISP1
Internet
ISP2
213.1.1.1
10.1.1.0/24
2nd peer
ASA2
195.1.1.100
EIGRP
Dynamic crypto map

ASA1 Configuration (1/2)

```
interface GigabitEthernet0/0.98
 nameif outside
 security-level 0
 ip address 194.1.1.10 255.255.255.0

interface GigabitEthernet0/0.103
 nameif inside
 security-level 100
 ip address 192.168.1.10 255.255.255.0

prefix-list TO-EIGRP seq 5 permit 0.0.0.0/0 ge 1

route-map TO-EIGRP permit 10
 match ip address prefix-list TO-EIGRP
 set metric 10000 100 255 1 1500

router eigrp 1
 no auto-summary
 network 192.168.1.0 255.255.255.0
 network 194.1.1.0 255.255.255.0
 passive-interface outside
 redistribute static route-map TO-EIGRP

route outside 0.0.0.0 0.0.0.0 194.1.1.2 1
```

ASA1 Configuration (2/2)

```
crypto ikev1 enable outside

crypto ikev1 policy 10
 authentication pre-share
 encryption aes-256
 hash sha
 group 5
 lifetime 86400

access-list 100 extended permit ip 192.168.1.0 255.255.255.0 10.1.1.0
255.255.255.0

crypto ipsec ikev1 transform-set SET1 esp-aes-256 esp-sha-hmac

crypto dynamic-map DMAP1 10 match address 100
crypto dynamic-map DMAP1 10 set ikev1 transform-set SET1
crypto dynamic-map DMAP1 10 set reverse-route
crypto map MAP1 10 ipsec-isakmp dynamic DMAP1
crypto map MAP1 interface outside

tunnel-group 213.1.1.1 type ipsec-l2l
tunnel-group 213.1.1.1 ipsec-attributes
 ikev1 pre-shared-key *****
 isakmp keepalive threshold 10 retry 3
```

Router Configuration (1/2)

```
interface FastEthernet0/0
 ip address 213.1.1.1 255.255.255.0
 crypto map MAP1

interface FastEthernet0/1
 ip address 10.1.1.1 255.255.255.0

ip route 0.0.0.0 0.0.0.0 213.1.1.2

ip sla 1
 icmp-echo 192.168.1.2 source-ip 10.1.1.1
 timeout 1000
 threshold 500
 frequency 10

ip sla schedule 1 life forever start-time now
```

Router Configuration (2/2)

```
crypto logging session

crypto isakmp policy 10
 encr aes 256
 authentication pre-share
 group 5

crypto isakmp key cisco address 0.0.0.0 0.0.0.0

crypto isakmp keepalive 20 5 periodic

access-list 100 permit ip 10.1.1.0 0.0.0.255 192.168.1.0 0.0.0.255

crypto ipsec transform-set SET1 esp-aes 256 esp-sha-hmac

crypto map MAP1 10 ipsec-isakmp
 set peer 194.1.1.10
 set peer 195.1.1.100
 set transform-set SET1
 match address 100
```

Notes

- Dynamic crypto maps cannot initiate tunnels (but scale better)
- IP SLA (or user traffic) is needed on router to bring tunnel up
- Redistribution of RRI routes to EIGRP guarantees correct routing
- DPD should be enabled on both sides to achieve switchover
- ASA should have lower DPD timers, than peer router for this topology

Possible Problems

- It's impossible to configure source IP address for IP SLA on ASA (if another ASA is used as a remote peer, instead of a router)
  - CSCtn29607 sla monitoring should allow users to configure source ip address
  - Inside host or other device should be used to initiate tunnels in this case
- Static crypto maps cannot be used on head office ASA1 and ASA2 in this topology due to ASA RRI implementation
  - CSCsx67450 ENH: ASA needs same RRI functionality as IOS
  - This is not a problem, because dynamic crypto maps is a good idea for this topology anyway

Basic Troubleshooting

- Verify that tunnel is established to ASA1

```
ASA1/VPN# show vpn-sessiondb detail l2l

Connection   : 213.1.1.1
Index        : 1                      IP Addr      : 213.1.1.1
Protocol     : IKEv1 IPsec
Encryption   : IKEv1: (1)AES256  IPsec: (1)AES256
Hashing      : IKEv1: (1)SHA1  IPsec: (1)SHA1
Bytes Tx     : 1344                   Bytes Rx     : 1344
Login Time   : 00:36:10 UTC Thu Aug 22 2013
Duration     : 0h:03m:38s
IKEv1 Tunnels: 1
IPsec Tunnels: 1

IKEv1:
  Tunnel ID    : 1.1
  UDP Src Port : 500                    UDP Dst Port : 500
  IKE Neg Mode : Main                   Auth Mode    : preSharedKeys
  Encryption   : AES256                 Hashing      : SHA1
  Rekey Int (T): 86400 Seconds          Rekey Left(T): 86182 Seconds
  D/H Group    : 5
  Filter Name  :
  IPv6 Filter  :
```

Basic Troubleshooting

- Verify that tunnel is established to ASA1

```
ASA1/VPN# show vpn-sessiondb detail l2l
…
IPsec:
  Tunnel ID    : 1.2
  Local Addr   : 192.168.1.0/255.255.255.0/0/0
  Remote Addr  : 10.1.1.0/255.255.255.0/0/0
  Encryption   : AES256                 Hashing      : SHA1
  Encapsulation: Tunnel
  Rekey Int (T): 3600 Seconds           Rekey Left(T): 3382 Seconds
  Rekey Int (D): 4608000 K-Bytes        Rekey Left(D): 4607999 K-Bytes
  Idle Time Out: 30 Minutes             Idle TO Left : 29 Minutes
  Bytes Tx     : 1344                   Bytes Rx     : 1344
  Pkts Tx      : 21                     Pkts Rx      : 21
```

Basic Troubleshooting

- Verify that RRI installed route to ASA1 routing table…

```
ASA1/VPN# show route
…
S     10.1.1.0 255.255.255.0 [1/0] via 194.1.1.2, outside
C     192.168.1.0 255.255.255.0 is directly connected, inside
D     195.1.1.0 255.255.255.0 [90/2560] via 192.168.1.100, 0:31:23, inside
C     194.1.1.0 255.255.255.0 is directly connected, outside
S*    0.0.0.0 0.0.0.0 [1/0] via 194.1.1.2, outside
```

- …and the route was received by ASA2 via EIGRP

```
ASA2/VPN# show route
…
D EX 10.1.1.0 255.255.255.0 [170/281600] via 192.168.1.10, 0:04:44, inside
C     192.168.1.0 255.255.255.0 is directly connected, inside
C     195.1.1.0 255.255.255.0 is directly connected, outside
D     194.1.1.0 255.255.255.0 [90/2560] via 192.168.1.10, 0:32:13, inside
S*    0.0.0.0 0.0.0.0 [1/0] via 195.1.1.1, outside
```

Switchover (ISP1 failure)

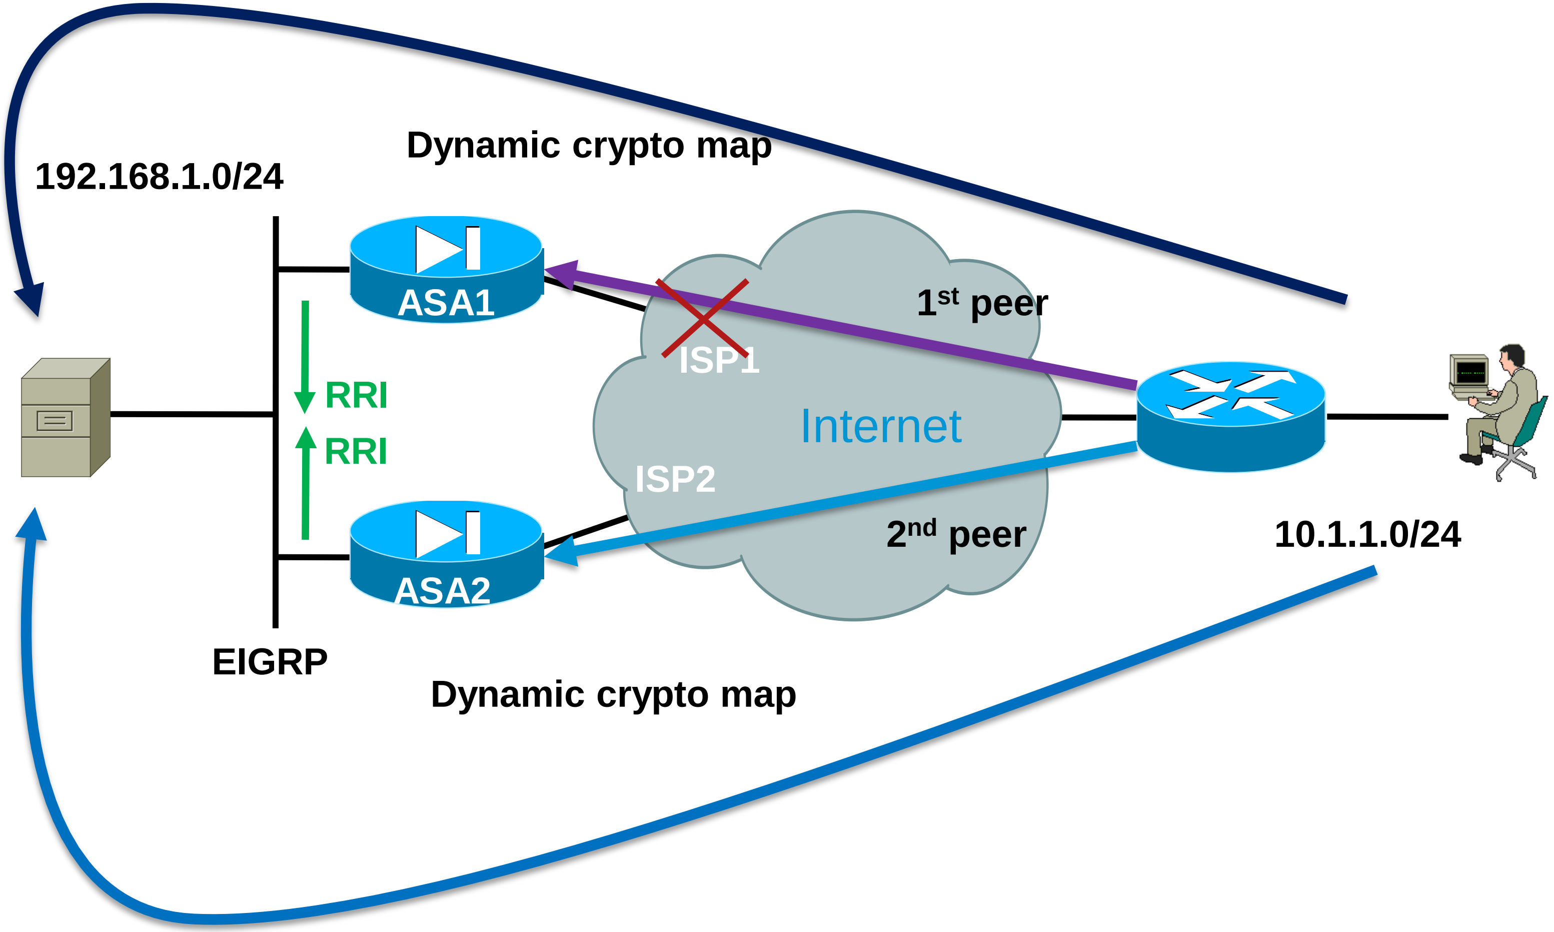

Switchover (ISP1 failure)

- What happens in case of ISP1 failure
- Remote IOS router needs 25-45 seconds to understand that its peer is unreachable (20 + 5 * 5)

```
crypto isakmp keepalive 20 5 periodic

Aug 22 12:08:10.187: ISAKMP:(2027):Sending NOTIFY DPD/R_U_THERE protocol 1
Aug 22 12:08:15.187: ISAKMP:(2027):DPD incrementing error counter (1/5)
Aug 22 12:08:20.187: ISAKMP:(2027):DPD incrementing error counter (2/5)
Aug 22 12:08:25.187: ISAKMP:(2027):DPD incrementing error counter (3/5)
Aug 22 12:08:30.187: ISAKMP:(2027):DPD incrementing error counter (4/5)
Aug 22 12:08:35.187: ISAKMP:(2027):DPD incrementing error counter (5/5)
Aug 22 12:08:35.187: ISAKMP:(2027):peer 194.1.1.10 not responding!

Aug 22 12:08:35.187: %CRYPTO-5-SESSION_STATUS: Crypto tunnel is DOWN.  Peer
194.1.1.10:500       Id: 194.1.1.10
```

Switchover (ISP1 failure)

- ASA needs 9-19 seconds to understand that peer is unreachable (10 + 3 * 3)

```
tunnel-group 213.1.1.1 ipsec-attributes
 isakmp keepalive threshold 10 retry 3

Aug 22 2013 12:08:20: %ASA-7-715036: Group = 213.1.1.1, IP = 213.1.1.1,
Sending keep-alive of type DPD R-U-THERE (seq number 0x707ae8c3)

Aug 22 2013 12:08:23: %ASA-7-715036: Group = 213.1.1.1, IP = 213.1.1.1,
Sending keep-alive of type DPD R-U-THERE (seq number 0x707ae8c4)

Aug 22 2013 12:08:26: %ASA-7-715036: Group = 213.1.1.1, IP = 213.1.1.1,
Sending keep-alive of type DPD R-U-THERE (seq number 0x707ae8c5)

Aug 22 2013 12:08:29: %ASA-3-713123: Group = 213.1.1.1, IP = 213.1.1.1, IKE
lost contact with remote peer, deleting connection (keepalive type: DPD)
```

Switchover (ISP1 failure)

- When IOS router understands that its peer is unreachable it immediately tries to establish new tunnel to 2nd VPN peer

```
Aug 22 12:08:35.187: %CRYPTO-5-SESSION_STATUS: Crypto tunnel is DOWN.  Peer
194.1.1.10:500       Id: 194.1.1.10
...
Aug 22 12:08:41.107: ISAKMP: Created a peer struct for 195.1.1.100, peer port
500
...
Aug 22 12:08:41.255: %CRYPTO-5-SESSION_STATUS: Crypto tunnel is UP  .  Peer
195.1.1.100:500       Id: 195.1.1.100
```

Switchover (ISP1 failure)

- If “default” option is used (“set peer 194.1.1.10 default”) it always tries to establish new tunnel to its default peer first, which slows down convergence

- Hence, “default” option is not recommended for this scenario

```
crypto map MAP1 10 ipsec-isakmp
 set peer 194.1.1.10
 set peer 195.1.1.100
 set transform-set SET1
 match address 100
```

- Note: switchback doesn't occur when ISP1 recovers

Scenario: Single ASA – Dual ISP

Topology

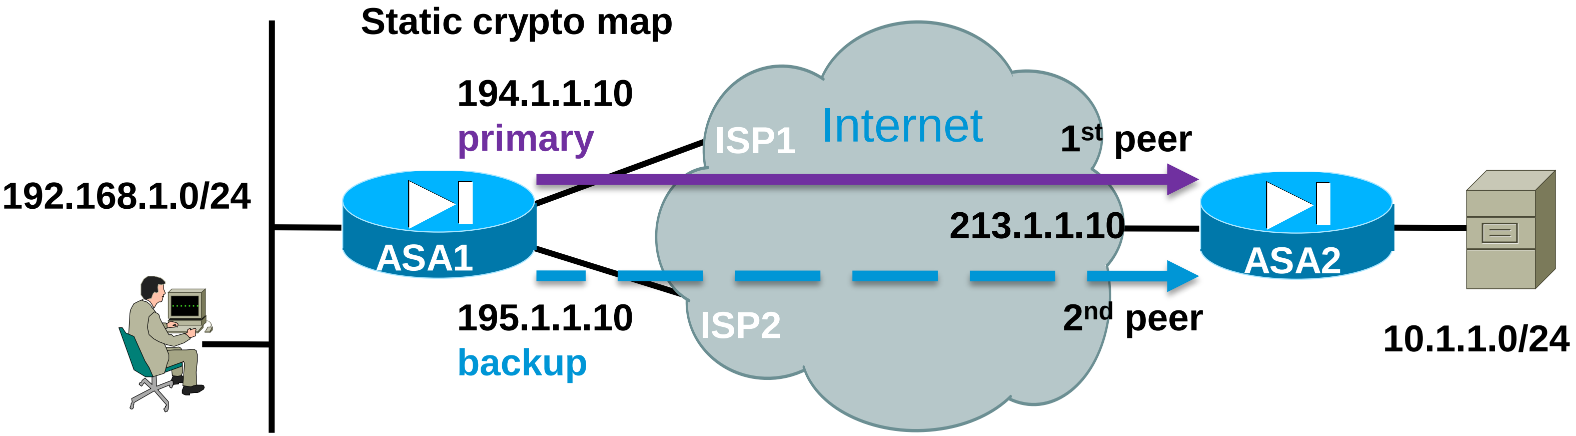
Static crypto map
194.1.1.10
primary
192.168.1.0/24
ASA1
195.1.1.10
backup
ISP1
Internet
ISP2
1st peer
213.1.1.10
2nd peer
ASA2
10.1.1.0/24

ASA1 Configuration (1/3)

```
interface GigabitEthernet0/0.98
 vlan 98
 nameif primary
 security-level 0
 ip address 194.1.1.10 255.255.255.0

interface GigabitEthernet0/0.99
 vlan 99
 nameif backup
 security-level 0
 ip address 195.1.1.10 255.255.255.0

interface GigabitEthernet0/0.103
 vlan 103
 nameif inside
 security-level 100
 ip address 192.168.1.10 255.255.255.0
```

ASA1 Configuration (2/3)

```
sla monitor 1
 type echo protocol ipIcmpEcho 194.1.1.2 interface primary
 timeout 1000
 threshold 500
 frequency 20

sla monitor schedule 1 life forever start-time now

track 1 rtr 1 reachability

route primary 0.0.0.0 0.0.0.0 194.1.1.2 1 track 1
route backup 0.0.0.0 0.0.0.0 195.1.1.2 100
```

ASA1 Configuration (3/3)

```
access-list 100 extended permit ip 192.168.1.0 255.255.255.0 10.1.1.0
255.255.255.0

crypto ipsec ikev1 transform-set SET1 esp-aes-256 esp-sha-hmac

crypto map MAP1 10 match address 100
crypto map MAP1 10 set peer 213.1.1.10
crypto map MAP1 10 set ikev1 transform-set SET1
crypto map MAP1 interface primary
crypto map MAP1 interface backup

crypto ikev1 enable primary
crypto ikev1 enable backup

crypto ikev1 policy 10
 authentication pre-share
 encryption aes-256
 hash sha
 group 5
 lifetime 86400

tunnel-group 213.1.1.10 type ipsec-l2l
tunnel-group 213.1.1.10 ipsec-attributes
 ikev1 pre-shared-key *****
```

ASA2 Configuration (1/2)

```
interface GigabitEthernet0/0.75
 nameif inside
 security-level 100
 ip address 10.1.1.10 255.255.255.0

interface GigabitEthernet0/0.76
 nameif outside
 security-level 0
 ip address 213.1.1.10 255.255.255.0

route outside 0.0.0.0 0.0.0.0 213.1.1.2 1

crypto ikev1 enable outside

crypto ikev1 policy 10
 authentication pre-share
 encryption aes-256
 hash sha
 group 5
 lifetime 86400
```

ASA2 Configuration (2/2)

```
access-list 100 extended permit ip 10.1.1.0 255.255.255.0 192.168.1.0
255.255.255.0

crypto ipsec ikev1 transform-set SET1 esp-aes-256 esp-sha-hmac

crypto map MAP1 10 match address 100
crypto map MAP1 10 set peer 194.1.1.10 195.1.1.10
crypto map MAP1 10 set ikev1 transform-set SET1
crypto map MAP1 interface outside

tunnel-group 194.1.1.10 type ipsec-l2l
tunnel-group 194.1.1.10 ipsec-attributes
 ikev1 pre-shared-key *****

tunnel-group 195.1.1.10 type ipsec-l2l
tunnel-group 195.1.1.10 ipsec-attributes
 ikev1 pre-shared-key *****
```

Possible Problems

- IP SLA and tracking are limited on ASA
  - only ICMP Echo is supported
  - multiple operations with AND/OR logic are not supported
  - single failure causes SLA and tracking object go down

  CSCti67445   ENH: Implement "delay up/down" command in ASA Object Tracking

  - it is possible to configure “num-packets” in IP SLA, but all of ICMP probes are sent at once

- IP SLA is not supported in multiple context mode

  CSCug56848   ENH: SLA Monitoring support in Multi-Context Mode

Basic Troubleshooting

- Verify routing, verify that TCP conn goes over tunnel to primary int.

```
ASA1# show track
Track 1
  Response Time Reporter 1 reachability
  Reachability is Up
  2 changes, last change 00:01:20
  Latest operation return code: OK
  Latest RTT (millisecs) 1
  Tracked by:
    STATIC-IP-ROUTING 0

ASA1# show route
…
C    192.168.1.0 255.255.255.0 is directly connected, inside
C    195.1.1.0 255.255.255.0 is directly connected, backup
C    194.1.1.0 255.255.255.0 is directly connected, primary
S*   0.0.0.0 0.0.0.0 [1/0] via 194.1.1.2, primary


ASA1# show conn long
…
TCP primary: 10.1.1.2/23 (10.1.1.2/23) inside: 192.168.1.2/51887
(192.168.1.2/51887), flags UIO , idle 6s, uptime 6s, timeout 1h0m, bytes 71
```

Switchover (ISP1 failure)

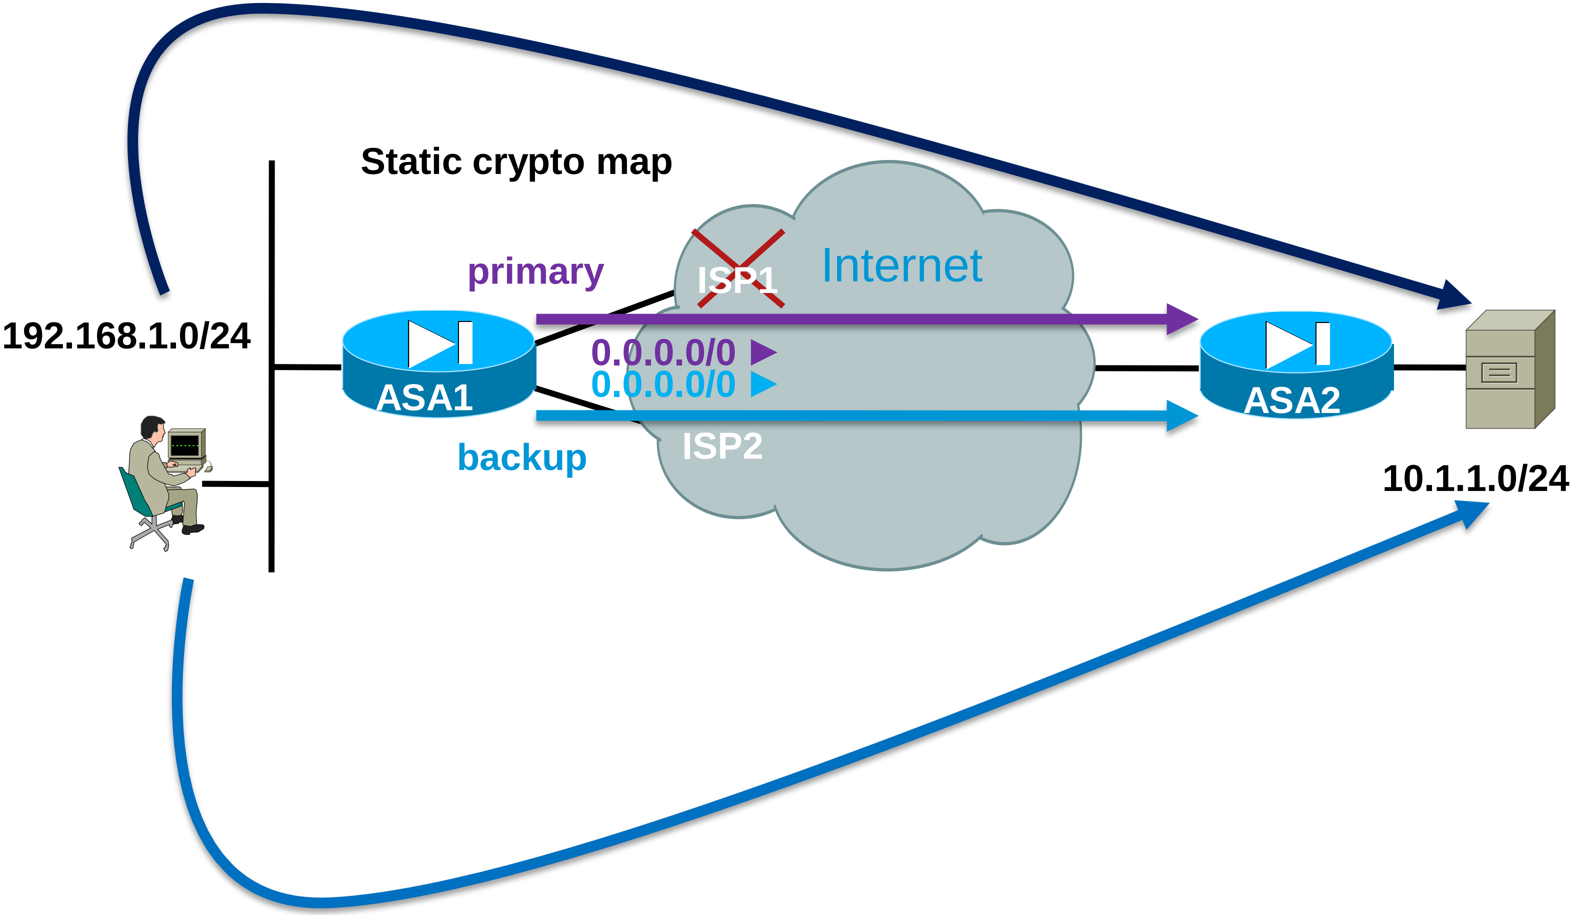

Switchover (ISP1 failure)

- DPD on ASA1 and remote ASA2 detects that ISP1 has failed
- (def. DPD timers on ASA: isakmp keepalive threshold 10 retry 2)

```
%ASA-3-713123: Group = 213.1.1.10, IP = 213.1.1.10, IKE lost contact with
remote peer, deleting connection (keepalive type: DPD)

%ASA-5-713259: Group = 213.1.1.10, IP = 213.1.1.10, Session is being torn
down. Reason: Lost Service

%ASA-4-113019: Group = 213.1.1.10, Username = 213.1.1.10, IP =
252.122.158.200, Session disconnected. Session Type: LAN-to-LAN, Duration:
0h:04m:01s, Bytes xmt: 914, Bytes rcv: 845, Reason: Lost Service
```

Switchover (ISP1 failure)

- IP SLA detects that ISP1 has failed
- Tracking installs new default route via backup interface

```
%ASA-6-622001: Removing tracked route 0.0.0.0 0.0.0.0 194.1.1.2, distance 1,
table Default-IP-Routing-Table, on interface primary

ASA1# show route
…
S*   0.0.0.0 0.0.0.0 [100/0] via 195.1.1.2, backup
```

- Note that TCP connection remains in the conn table via primary interface…

```
ASA1# show conn long
…
TCP primary: 10.1.1.2/23 (10.1.1.2/23) inside: 192.168.1.2/51887
(192.168.1.2/51887), flags UIO , idle 1m42s, uptime 2m30s, timeout 1h0m, bytes
71
```

Switchover (ISP1 failure)

- New TCP connection is initiated
- Which causes new tunnel to be established (via backup interface)

```
%ASA-5-713041: IP = 213.1.1.10, IKE Initiator: New Phase 1, Intf inside, IKE
Peer 213.1.1.10  local Proxy Address 192.168.1.0, remote Proxy Address
10.1.1.0,  Crypto map (MAP1)

ASA1# show conn long
…
TCP backup: 10.1.1.2/23 (10.1.1.2/23) inside: 192.168.1.2/19371
(192.168.1.2/19371), flags UIO , idle 1m53s, uptime 1m53s, timeout 1h0m, bytes
71

TCP primary: 10.1.1.2/23 (10.1.1.2/23) inside: 192.168.1.2/51887
(192.168.1.2/51887), flags UIO , idle 5m11s, uptime 5m59s, timeout 1h0m, bytes
71
```

- The same sequence of events takes place when ISP1 recovers (in this specific topology)

What will happen with old TCP connection?

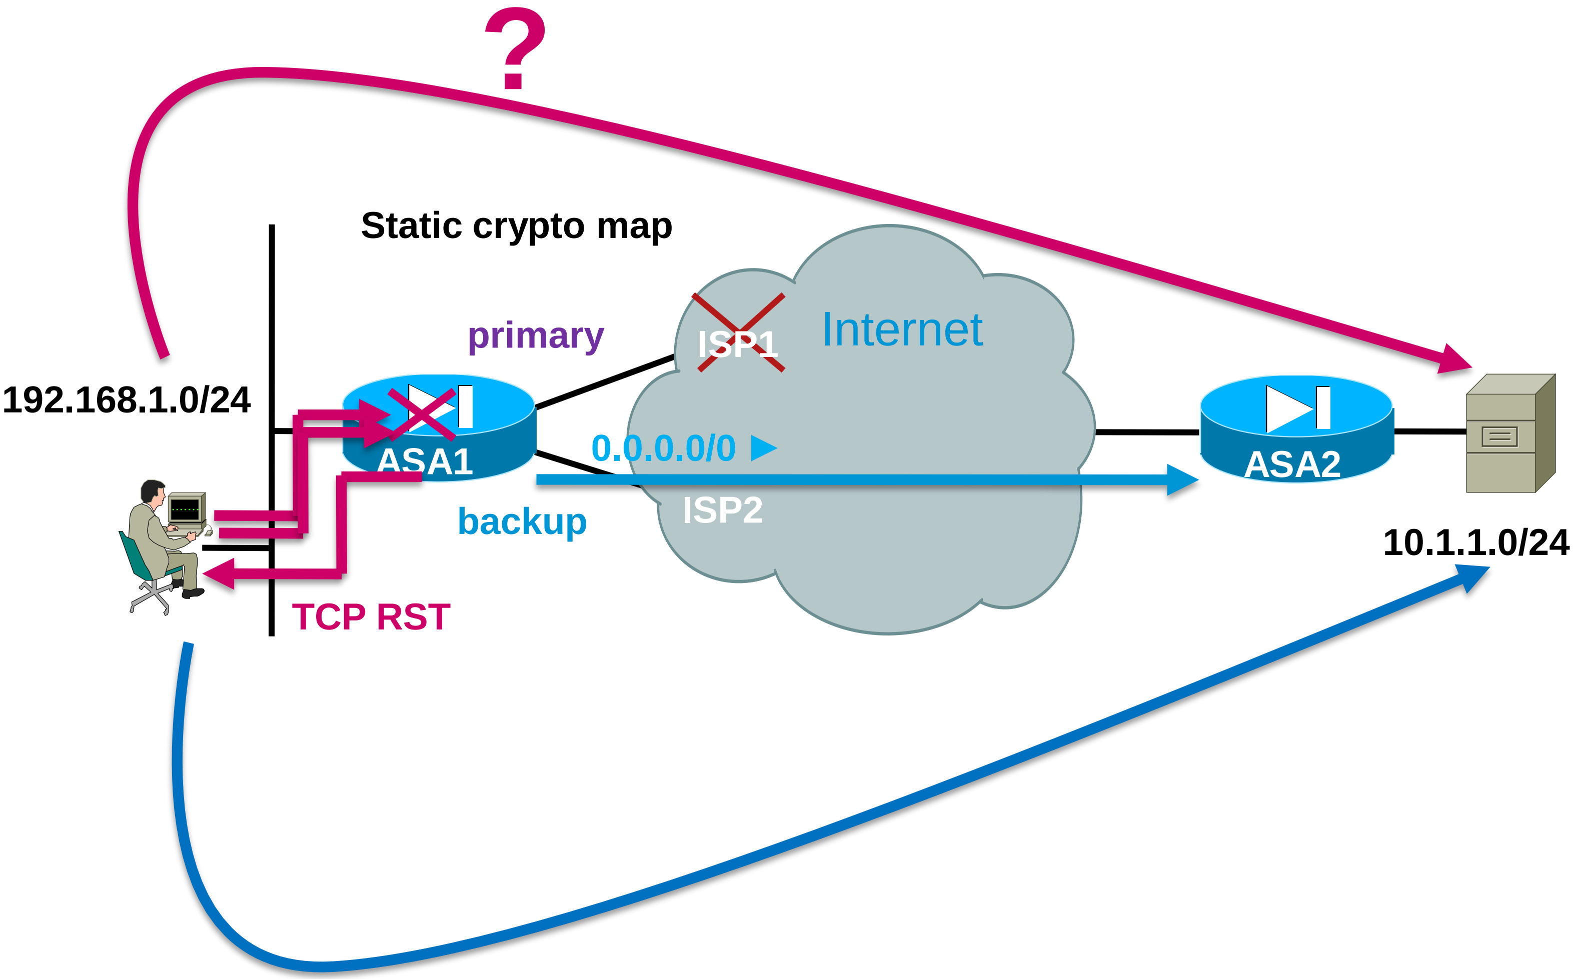

What happens with TCP connections

- But what will happen with old TCP connection that goes over non-existent tunnel?
- It will be torn down on the ASA as soon as next packet is received (the connection remains active on the client at this point)
- ASA responds with TCP RST to 1st client retransmit which tears down this TCP connection on the client (this happens because "service resetoutbound" is enabled on ASA by default)
- The TCP connection remains active on the server until it times out due to inactivity
- UDP connections are transparently re-created on ASA by traffic

What happens with TCP connections

- The following syslog messages are produced:

```
%ASA-6-302014: Teardown TCP connection 144 for primary:10.1.1.2/23 to
inside:192.168.1.2/51887 duration 0:08:00 bytes 71 Tunnel being brought up or
torn down

%ASA-6-106015: Deny TCP (no connection) from 192.168.1.2/51887 to 10.1.1.2/23
flags PSH ACK  on interface inside
```

- Also, the following ASP drop counters are incremented:

```
ASA1# show asp drop

Frame drop:
  IPSEC tunnel is down (ipsec-tun-down)             1
  First TCP packet not SYN (tcp-not-syn)            1
```

- Standard conn teardown message is produced for UDP:

```
%ASA-6-302016: Teardown UDP connection 4539 for primary:10.1.1.2/12345 to
inside:192.168.1.2/54321 duration 0:01:30 bytes 64
```

Possible Problems

- This process looks tricky
- Why don’t we tear down ALL connections at once when their underlying tunnel goes down? Performance / CPU load reasons!
- Can we guarantee that all connections are torn down? This is very important to avoid tunnel flapping! Read on!
- What will happen with connections if their underlying tunnel doesn’t go down, for example, when DPD is disabled? Read on!

Connections Creation and Teardown

Side Note: Clear-text Connections

- For clear-text connections we have two mechanisms to facilitate convergence when outgoing interface changes due to routing change

- The first one was implemented in 8.0(5) software release

  CSCso42904 When routes change, connections should be updated automatically

- Typical diagnostics is:

```
%ASA-6-110003: Routing failed to locate next hop for TCP from
inside:192.168.1.4/51394 to primary:4.4.4.4/23
```

- Connection is torn down and “no-adjacency” ASP drop counter is incremented

- This mechanism is enabled by default and cannot be disabled

Side Note: Clear-text Connections

- Second mechanism “enhances” the previous one for cases when floating-static routes (routes that have bigger admin distance) are used, for example:

```
route primary 0.0.0.0 0.0.0.0 194.1.1.2 1 track 1
route backup 0.0.0.0 0.0.0.0 195.1.1.2 100
```

- With such configuration switchback from the backup interface to the primary interface doesn’t tear down connections running over the backup interface. This may sometimes cause problems

- Switchover from the primary interface to the backup interface works fine and connections are torn down

Side Note: Clear-text Connections

- The fix was integrated into 8.2(4.2), 8.3(2.12), 8.4(1.1) releases

  CSCsy19222 Conns should update when using dynamic protocol and floating statics

- In order to tear down connections during switchback, explicit configuration is required:

```
timeout floating-conn 0:00:30
```

- We lookup egress interface when the first next packet arrives on a flow. If there is a change in egress interface, we schedule the flow for termination after user configured time-out (floating-conn).
- All connection traffic is sent to the backup interface in the meantime, even though default route points to the primary interface

Side Note: Clear-text Connections

- The "floating-conn" timeout is off by default
- The minimal configurable timeout is 30 seconds
- It is only needed when it is necessary to torn down existing connections through the backup interface when primary path becomes available
- Longer timeout allows existing short-lived flows to complete. If timeout is off, sort of load-balancing can be achieved in some scenarios
- Note that both fixes do not work in case of connections running over tunnels!

Connections Over Tunnels

- In certain scenarios IPsec tunnel may not go down even if the peer is unreachable
- Example: DPD is disabled on the ASA
  - Tunnel stays up even if the peer is unreachable, TCP connections running over the tunnel are not torn down and hang
- Example: switchback from a backup ISP to a primary ISP in certain PPPoE scenarios with DPD enabled
  - Outside of the scope of this presentation

Connections Over Tunnels

- So, we may end up in a situation when tunnel is up, but cannot pass any traffic and all connections over the tunnel hang
- And at the same time routing can re-converge (for example, IP SLA or dynamic routing) and default route can point to a new interface
- Fortunately, IPsec re-convergence is usually triggered by new TCP/UDP connections which are routed to new interface after routing change. This triggers new P2 tunnel over new interface. Old tunnel over old interface is torn down
- The following message can be produced:

```
%ASA-3-713258: IP = 213.1.1.1, Attempting to establish a phase2 tunnel on
backup interface but phase1 tunnel is on primary interface. Tearing down old
phase1 tunnel due to a potential routing change.
```

Solution for Connections Over Tunnels

- But what will happen if we don’t have “new” connections, for example, when ASA is used as an encryption device to encrypt and encapsulate GRE traffic exchanged by two routers?
- The answer was given by the CSCsz04730 in 8.2(5.20), 8.3(2.29), 8.4(3.1), 8.4(4), 9.0(1) versions

  CSCsz04730 PIX/ASA: When route changes connections over IPSEC tunnel not torn down

- How it was fixed:

  Modify the IPsec packet injection code to monitor for changes in the egress interface and, if a change is detected, tear down all tunnels to the affected peer

Solution for Connections Over Tunnels

- The following can be seen in fixed versions when DPD is disabled, primary ISP experiences failure, routing installs new default route via backup interface, next TCP packet arrives to the ASA and should be sent to the tunnel running over primary interface:

```
%ASA-5-713259: Group = 213.1.1.1, IP = 213.1.1.1, Session is being torn down.
Reason: Administrator Reset

%ASA-4-113019: Group = 213.1.1.1, Username = 213.1.1.1, IP = 213.1.1.1,
Session disconnected. Session Type: LAN-to-LAN, Duration: 0h:04m:52s, Bytes
xmt: 3644, Bytes rcv: 3765, Reason: Administrator Reset
```

- TCP connection remains in the ASA conn table at this point. It will be torn down as before when next packet is received by ASA:

```
%ASA-6-302014: Teardown TCP connection 261 for primary:10.1.1.2/23 to
inside:192.168.1.2/26198 duration 0:04:51 bytes 1443 Tunnel has been torn down
```

Example of the Problem

- Still, there are scenarios, when this fix doesn't help… Example:

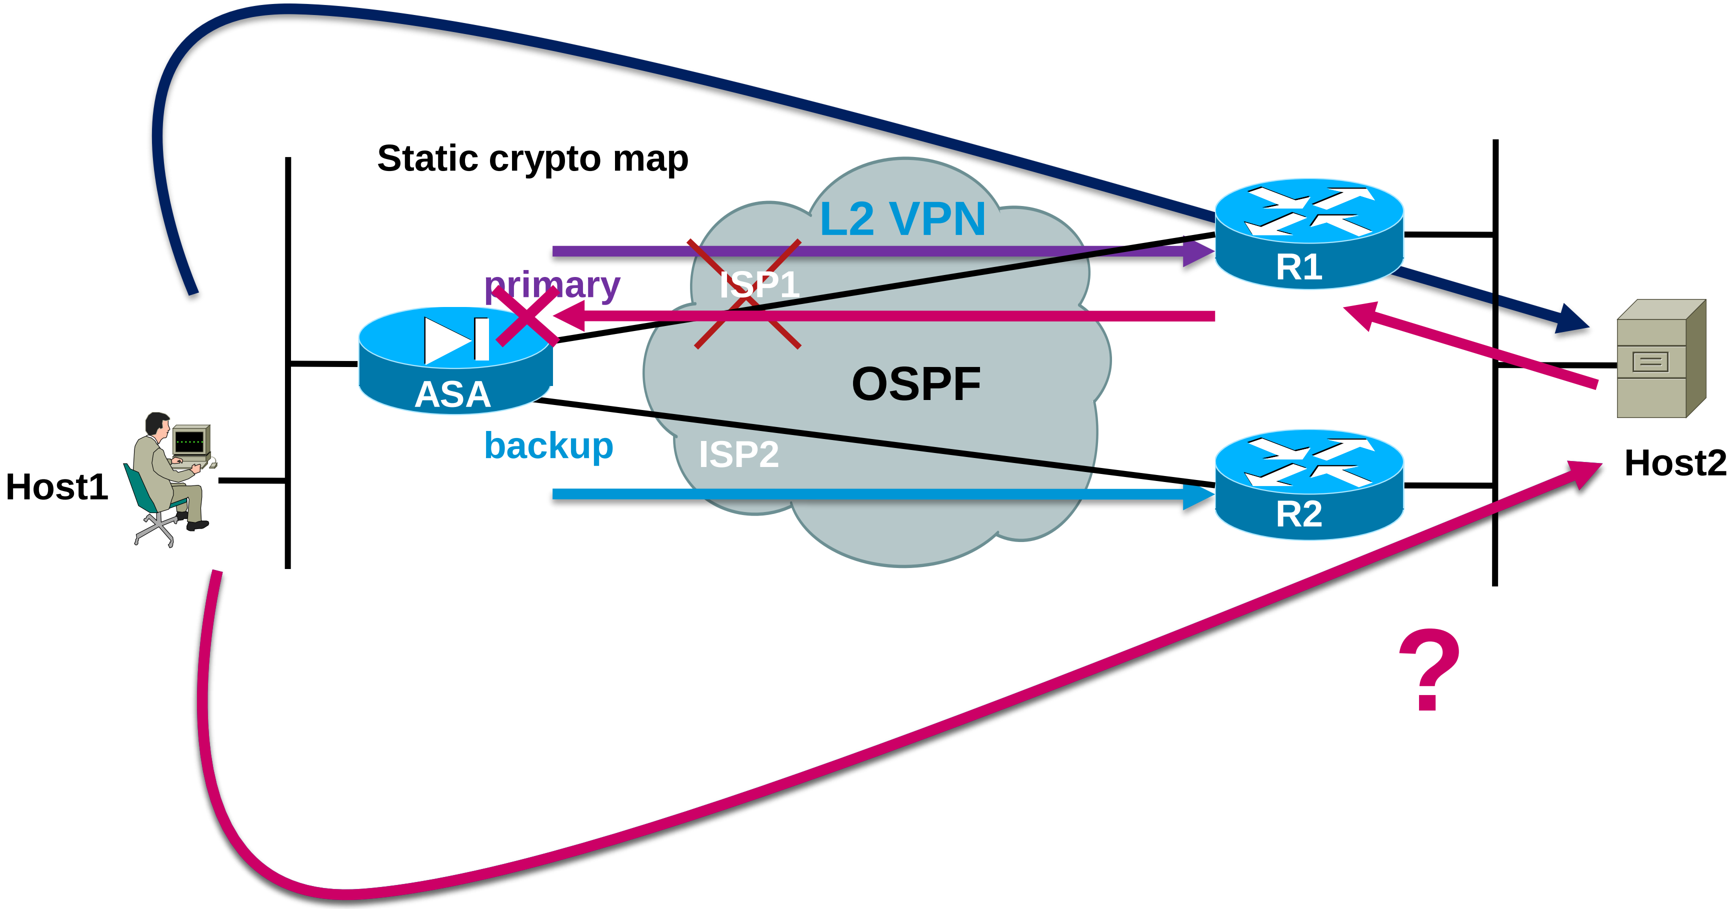

Example of the Problem

- Consider the following sequence of events:
  - ISP1 experienced failure and OSPF removed routes through primary interface
  - TCP connection was opened from Host1 to Host2 over ASA-R2 tunnel
  - ISP1 recovered and OSPF installed routes via primary int. with better metric
  - Completely new TCP connection is now opened from Host1 to Host2
  - And is routed to primary interface
  - This triggers new P2 tunnel from ASA to R1
  - Old tunnel (ASA-R2) is not torn down in this case
  - If traffic goes over old TCP connection it is routed to old tunnel by the ASA
  - The traffic arrives to Host2 and reply is sent back to R1 (preferred path)
  - R1 has tunnel to the ASA and sends the reply to the tunnel
  - ASA receives the reply from R1 (over the new tunnel, over primary interface)
  - And drops it, because conn entry for this connection points to the backup interface (asymmetric routing!)

Example of the Problem

- The root cause of the problem is the fact that ASA allows two identical IPsec (P2) tunnels (tunnels with identical “proxy identities”) to two different IPsec peers (routers in this case)

- The following bug was recently opened to investigate this behavior:

  CSCui57181   ASA: Do not allow two IPsec tunnels with identical proxy IDs

Notes

- It should be mentioned that ASA doesn't allow two identical tunnels to the same IPsec peer
- If ASA is initiator, one tunnel is torn down with the:

```
%ASA-3-713258: IP = 213.1.1.1, Attempting to establish a phase2 tunnel on
backup interface but phase1 tunnel is on primary interface. Tearing down old
phase1 tunnel due to a potential routing change.

%ASA-4-113019: Group = 213.1.1.1, Username = 213.1.1.1, IP = 213.1.1.1,
Session disconnected. Session Type: LAN-to-LAN, Duration: 0h:18m:04s, Bytes
xmt: 2612, Bytes rcv: 2924, Reason: User Requested
```

- If it is a responder and the same peer attempts to establish identical P2 tunnel to the ASA, we can see:

```
%ASA-5-713259: Group = 213.1.1.1, IP = 213.1.1.1, Session is being torn down.
Reason: Peer Reconnected

%ASA-4-113019: Group = 213.1.1.1, Username = 213.1.1.1, IP = 213.1.1.1,
Session disconnected. Session Type: LAN-to-LAN, Duration: 0h:04m:14s, Bytes
xmt: 554, Bytes rcv: 445, Reason: Peer Reconnected
```

Notes

- Finally, the bug CSCui57181 is only valid if the ASA is P2 initiator
- If it is a responder and some other peer attempts to establish identical tunnel to the ASA, we can see:

```
%ASA-7-715068: Group = 194.1.1.1, IP = 194.1.1.1, Duplicate remote proxy
(10.1.1.0/255.255.255.0) detected. Replacing old tunnel. Old peer:
195.1.1.2:500; New peer: 194.1.1.1:500

%ASA-5-713259: Group = 195.1.1.2, IP = 195.1.1.2, Session is being torn down.
Reason: Peer Address Changed

%ASA-4-113019: Group = 195.1.1.2, Username = 195.1.1.2, IP = 195.1.1.2,
Session disconnected. Session Type: LAN-to-LAN, Duration: 0h:01m:52s, Bytes
xmt: 474, Bytes rcv: 445, Reason: Peer Address Changed
```

Another Example

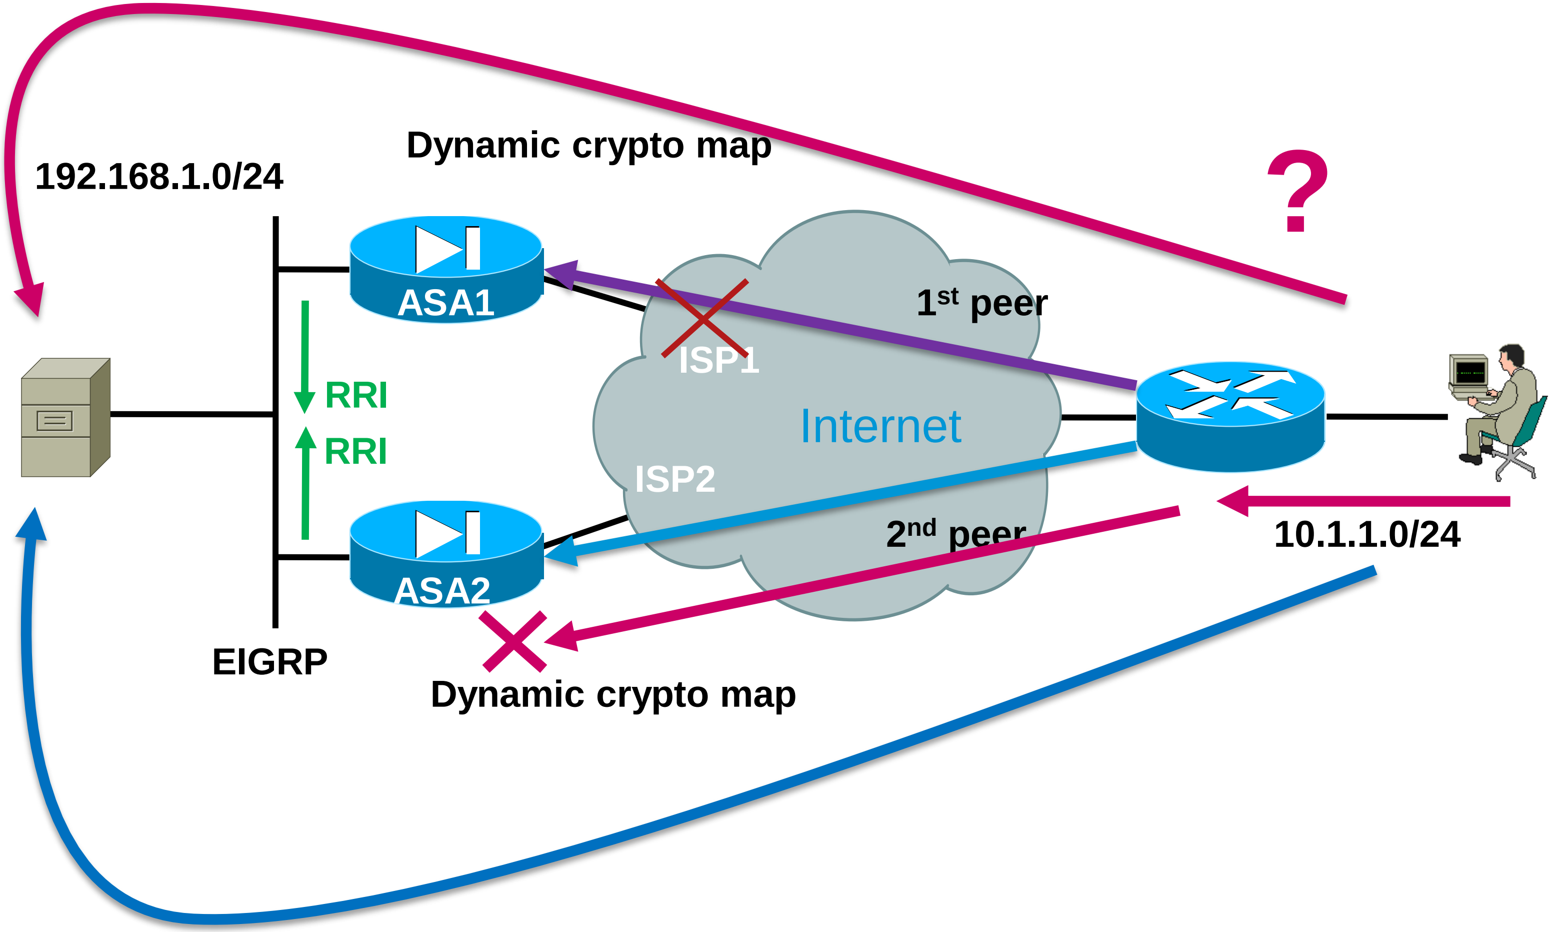

Another Example

- Consider the following sequence of events:
  - TCP connection 1 was opened over Router-ASA1 tunnel
  - After that ISP1 experienced failure
  - DPD tore down the failed tunnel
  - Completely new TCP connection 2 was opened
  - And new tunnel was created via ISP2 between Router and ASA2
  - The problem is that router doesn't have notion of "TCP connections"
  - If the client sends traffic over connection 1, it is routed to Router-ASA2 tunnel
  - The traffic arrives to ASA2 and is dropped there, because ASA2 doesn't have such a connection in its conn table

Solution for Both Cases

- Obviously, we need something to tear down hung connections on the client and/or server
- The following command does the trick

```
service resetinbound
```

- ASA will send TCP RST to the sender if connection entry doesn’t exist for the traffic
- This command is disabled by default

Summary of Solutions for Connections

- Remember that connection egress interface has priority over routing table
- The following features help rebuild connections running over tunnels in dual-ISP topologies:
  - Tear down connections if their tunnel has been torn down
  - Tear down tunnel if connection egress interface has changed (CSCsz04730)
  - Do not allow tunnels with identical proxy identities
- The “service resetinbound”, “service resetoutbound” help tear down TCP connections on endpoints
- Dead Peer Detection plays crucial role in many scenarios
- The following feature is not compatible and cannot be used in Single ASA – Dual ISP topologies:

  sysopt connection preserve-vpn-flows

Final Notes About Connections

- Even though many protective technologies exist, you can still face with bugs
- Example:

  CSCue97782 ASA: Old connections tear down IPsec VPN tunnel on switchover

  CSCue46275 Connections not timing out when the route changes on the ASA

- This bug can cause VPN flapping due to stale connection entries
- Fixed in: 8.4(6.6), 9.0(3.1), 9.1(2.5)
- It is highly recommended to upgrade before configuring redundant Site-to-Site VPNs on ASA

OSPF Over Tunnels

OSPF Over Tunnels

- OSPF over L2L tunnels is supported on ASA since 7.0
- In other to send OSPF over tunnel it should run on the same interface where crypto map is applied (remember that ASA doesn’t support logical tunnel interfaces)
- ASA negotiates additional tunnel with its peers for OSPF:

  [ASA-IP/32 <-> Remote-IP/32] for IP protocol 89 (OSPF)

- This should be explicitly included into crypto ACL
- ASA cannot send multicast traffic to the tunnel

  “ospf network point-to-point non-broadcast” should be configured

  “neighbor 213.1.1.10 interface primary” should be configured

OSPF Over Tunnels

- Currently this works between ASA devices only

  Routers can be configured with “ip ospf network point-to-multipoint non-broadcast”, but this doesn’t help much: OSPF adjacency cannot be established if ASA outside IP and router IP are in different IP subnets:

  CSCty10786   ENH: Add support for OSPF p2p neighbor on different network

```
OSPF-1 ADJ   Fa0/0: Rcv pkt from 195.1.1.10, area 0.0.0.0 : src not on the
same network
```

- ASA doesn’t support “point-to-multipoint non-broadcast” OSPF network type

  This means that ASA can have only one OSPF neighbor on each interface, i.e. you need as many interfaces as you have tunnels!

  CSCsh17456   OSPF: Implement point-to-multipoint for supporting multiple neighbors

```
ERROR: Only one neighbor allowed on point-to-point interfaces
```

OSPF Over Tunnels

- Explicit host route to tunnel destination needs to be configured on each ASA

  route primary 213.1.1.10 255.255.255.255 194.1.1.2

- ASA runs OSPF on its outside interface and announces corresponding network to its peer. Without static host route this would cause recursive routing failures and tunnel flapping

```
ASA1# show route
…
O    213.1.1.0 255.255.255.0 [110/2] via 213.1.1.10, 0:34:10, primary
S    213.1.1.10 255.255.255.255 [1/0] via 194.1.1.2, primary
O    10.1.1.0 255.255.255.0 [110/11] via 213.1.1.10, 0:34:10, primary
C    192.168.1.0 255.255.255.0 is directly connected, inside
C    195.1.1.0 255.255.255.0 is directly connected, backup
C    194.1.1.0 255.255.255.0 is directly connected, primary
S*   0.0.0.0 0.0.0.0 [1/0] via 194.1.1.2, primary
```

OSPF Over Tunnels

- Overall, this feature can be used as a “replacement” for IP SLA, DPD and RRI in a typical “triangle” topology:
  - OSPF Hello packets will check reachability, instead of DPD (but DPD is still needed to tear down failed tunnels)
  - OSPF will rebuild routing table in case of ISP failure (but IP SLA is still needed to provide redundancy for Internet access)
  - RRI and redistribution are not needed anymore
- To achieve ISP redundancy you need two interfaces
  - Typical topology is head-office with two ASAs and remote office with single ASA connected to two different ISPs via two different physical or sub-interfaces
  - The problem is that you need new outside interface on each ASA to connect new remote office (CSCsh17456)! This is not scalable!
- The best solution would be to implement logical tunnel interfaces!
  - CSCtu06739 ENH: Introduce support for virtual interfaces for VPN on ASA

OSPF Over Tunnels

- Sample configuration (see also http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_configuration_example09186a00804acfea.shtml)

```
interface GigabitEthernet0/0.98
 vlan 98
 nameif primary
 security-level 0
 ip address 194.1.1.10 255.255.255.0
 ospf cost 1
 ospf network point-to-point non-broadcast

interface GigabitEthernet0/0.103
 vlan 103
 nameif inside
 security-level 100
 ip address 192.168.1.10 255.255.255.0
```

OSPF Over Tunnels

- Sample configuration (see also http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_configuration_example09186a00804acfea.shtml)

```
access-list 100 extended permit ip 192.168.1.0 255.255.255.0 10.1.1.0
255.255.255.0
access-list 100 extended permit ospf host 194.1.1.10 host 213.1.1.10

crypto ipsec ikev1 transform-set SET1 esp-aes-256 esp-sha-hmac

crypto map MAP1 10 match address 100
crypto map MAP1 10 set peer 213.1.1.10
crypto map MAP1 10 set ikev1 transform-set SET1
crypto map MAP1 interface primary

router ospf 1
 network 192.168.1.0 255.255.255.0 area 0
 network 194.1.1.0 255.255.255.0 area 0
 neighbor 213.1.1.10 interface primary

route primary 213.1.1.10 255.255.255.255 194.1.1.2 1
```

Conclusion

Conclusion

- Fault-tolerant Site-to-Site IPsec VPNs are possible with ASA
- Each topology requires its own set of features, such as RRI, etc.
- Connections should be torn down and reestablished, unless stateful failover is used
- The “service resetoutbound” and “service resetinbound” commands may help
- Dynamic routing over tunnels with OSPF is possible in certain topologies, but this solution doesn’t scale well
- Logical GRE/VTI tunnels are not implemented on ASA
  - CSCtu06739 ENH: Introduce support for virtual interfaces for VPN on ASA
- New routing code is being ported to ASA and will change certain things in the future

Опрос #3

Опрос #3: Какие темы семинаров по безопасности вам интересны?

1. Использование AnyConnect с ASA и маршрутизаторами
2. Использование ASR1k для организации L2L VPNs
3. Clustering в ASA 9.x
4. IPv6 на ASA
5. Smart Call-Home на ASA и другие возможности Smart-поддержки Cisco
6. Продукты IronPort

Q & A

Эксперт ответит на некоторые Ваши вопросы. Используйте Q&A панель, чтобы задать еще вопросы

Сессия «Спросить Эксперта»

Получить дополнительную информацию, а также задать вопросы экспертам в рамках данной темы вы можете в течение двух недель, на странице, доступной по ссылке

https://supportforums.cisco.com/community/russian/expert-corner

Вы можете получить видеозапись данного семинара и текст сессии Q&A в течении ближайших 5 дней по следующей ссылке

https://supportforums.cisco.com/community/russian/expert-corner/webcast

Приглашаем Вас активно участвовать в Cisco Support Community и социальных сетях

https://supportforms.cisco.com/community/russian

http://www.facebook.com/CiscoRu

http://twitter.com/CiscoRussia

http://www.youtube.com/user/CiscoRussiaMedia

http://itunes.apple.com/us/app/cisco-technical-support/id398104252?mt=8

http://www.linkedin.com/groups/CSC-Cisco-Support-Community-3210019

Newsletter Subscription:
https://tools.cisco.com/gdrp/coiga/showsurvey.do?surveyCode=589&keyCode=146298_2&PHYSICAL%20FULFILLMENT%20Y/N=NO&SUBSCRIPTION%20CENTER=YES

Спасибо за
Ваше время

Пожалуйста, участвуйте в опросе

Thank you.

cisco