Configuration exaple

Light-weight method of redirection and re-insertion of traffic after centralized security clearing

Rafal Jan Szarecki

Table of Content

Topology

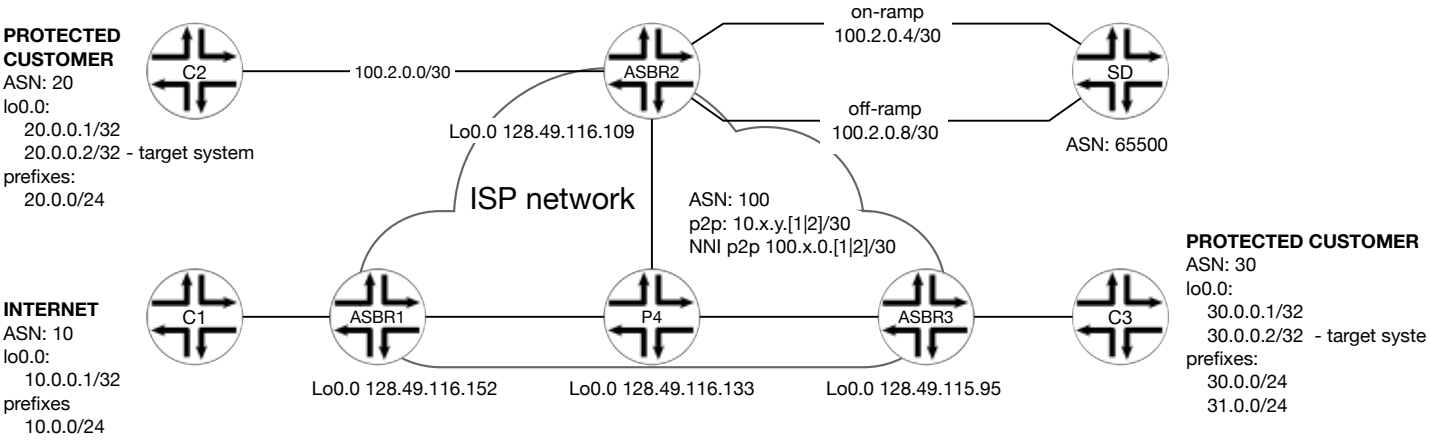

Figure 1 Reference topology

Configuration

Please note we've chosen to highlight only the relevant parts of the configuration.

ASBR3 configuration template (internet edge, where customer is connected)

Rib-group	routing-options {
	rib-groups {
	SecuredCustomer {
	import-rib [inet.0 inet.3];
	}
	}
	autonomous-system 100;
	}

eBGP session → both rib	protocols {
	bgp {
	group iBGP {
	[…]
	}
	group eBGP {
	type external;
	peer-as 30;
	neighbor 100.3.0.2 {
	family inet {
	unicast {
	rib-group SecuredCustomer;
	}
	}

	} } } }

iBGP sessions	protocols {
	bgp {
	group iBGP {
	type internal;
	local-address 128.49.115.95;
	import iBGP-import;
	export NHS;
	neighbor 128.49.116.109 {
	family inet {
	unicast;
	labeled-unicast {
	rib {
	inet.3;
	}
	}
	}
	}
	neighbor 128.49.115.95;
	neighbor 128.49.116.152;
	}
	}
	}
	policy-options {
	policy-statement iBGP-import {
	term no-LU {
	to rib inet.3;
	then reject;
	}
	}
	}

ASBR1 configuration template (peering – no customers connected)

Rib-group	routing-options {
	autonomous-system 100;
	}

eBGP session → both rib	protocols {
	bgp {
	group eBGP {
	type external;
	neighbor 100.1.0.2 {
	peer-as 10;
	}
	}
	}
	}

iBGP sessions	protocols {
	bgp {
	group iBGP {
	type internal;

```
            local-address 128.49.116.152;
            export NHS;
            neighbor 128.49.115.95;
            neighbor 128.49.116.109;
        }
    }
}
```

ASBR2 configuration template (screening cluster and customer are connected)

Configuration in *gray-italic* needed only if same router connects also protected customer networks

Rib-group

```
routing-options {
    rib-groups {
        off-ramp {
            import-rib [ inet.3 off-ramp.inet.0 ];
        }
        SecuredCustomer {
            import-rib [ inet.0 inet.3 off-ramp.inet.0 ];
        }
    }
    autonomous-system 100;
}
```

eBGP session → both rib

```
protocols {
    BGP
        group eBGP {
            type external;
            peer-as 20;
            neighbor 100.2.0.2 {
                family inet {
                    unicast {
                        rib-group SecuredCustomer;
                    }
                }
            }
        }
        group screen-cluster {
            type external;
            peer-as 65500;
            neighbor 100.2.0.6;
        }
    }
}
```

iBGP sessions

```
protocols {
    bgp {
        group iBGP {
            type internal;
            local-address 128.49.116.109;
            export NHS;
            neighbor 128.49.115.95 {
                family inet {
                    unicast;
                    labeled-unicast {
                        rib-group off-ramp;
                        rib {
```

```
                                inet.3;
                            }
                        }
                    }
                }
                neighbor 128.49.116.152;
            }
        }
}
```

Verification

ASBR1 (peering)

The 30.0.0.2/32 and 20.0.0.2/32 are redirected for screening.

Routes in RIB - summary	`regress@ASBR1-re> show route summary`
	`Autonomous system number: 100`
	`Router ID: 128.49.115.41`
	
	`inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden)`
	`Direct: 4 routes, 4 active`
	`Local: 3 routes, 3 active`
	`BGP: 6 routes, 6 active`
	`Static: 4 routes, 4 active`
	`IS-IS: 8 routes, 8 active`
	`LDP: 1 routes, 1 active`
	
	`inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)`
	`LDP: 3 routes, 3 active`
	
	`iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Direct: 1 routes, 1 active`
	
	`mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)`
	`MPLS: 4 routes, 4 active`
	`LDP: 4 routes, 4 active`
	
	`inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)`
	`Direct: 2 routes, 2 active`
	`INET6: 1 routes, 1 active`
		• Only default routing instances
		• Static routes are lab OOB
		• No BGP routes in inet.3. Only LDP (looopbacks of ASBR)
		• Only LDP entries in mpls.o RIB (plus 4 well-known labbes)

BGP routes	`regress@ASBR1-re> show route protocol bgp`
	
	`inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden)`
	`+ = Active Route, - = Last Active, * = Both`
	
	`10.0.0.0/24 *[BGP/170] 20:50:15, localpref 100`
		• inspected destination are redirected to ASBR2 (`128.49.114.132`)
		•

	AS path: 10 I, validation-state: unverified > to 100.1.0.2 via ge-0/0/0.0 20.0.0.0/24 *[BGP/170] 20:39:08, localpref 100, from 128.49.114.132 AS path: 20 I, validation-state: unverified > to 10.1.4.2 via ge-0/0/1.0, Push 299776 20.0.0.2/32 *[BGP/170] 20:39:08, localpref 100, from 128.49.114.132 AS path: 65500 I, validation-state: unverified > to 10.1.4.2 via ge-0/0/1.0, Push 299776 30.0.0.0/24 *[BGP/170] 15:50:39, localpref 100, from **128.49.114.112** AS path: 30 I, validation-state: unverified > to 10.1.4.2 via ge-0/0/1.0, Push 299824 30.0.0.2/32 *[BGP/170] 20:39:08, localpref 100, from 128.49.114.132 AS path: 65500 I, validation-state: unverified > to 10.1.4.2 via ge-0/0/1.0, Push 299776 31.0.0.0/24 *[BGP/170] 15:50:39, localpref 100, from 128.49.114.112 AS path: 30 I, validation-state: unverified > to 10.1.4.2 via ge-0/0/1.0, Push 299824	
Prefixes in IP FIB's	% cli show route forwarding-table family inet | grep -A 1 "[1-3][0-1]\.0\.0|Routing table|Destination" | grep -v "\-\-" Routing table: default.inet Internet: Destination Type RtRef Next hop Type Index NhRef Netif 10.0.0.0/8 user 0 0:0:5e:0:1:50 ucst 341 6 fxp0.0 10.0.0.0/24 user 0 100.1.0.2 ucst 587 4 ge-0/0/0.0 10.1.4.0/30 intf 0 rslv 585 1 ge-0/0/1.0 20.0.0.0/24 user 0 indr 1048574 4 10.1.4.2 Push 299776 593 2 ge-0/0/1.0 20.0.0.2/32 user 0 indr 1048574 4 10.1.4.2 Push 299776 593 2 ge-0/0/1.0 30.0.0.0/24 user 0 indr 1048575 3 10.1.4.2 Push 299824 594 2 ge-0/0/1.0 30.0.0.2/32 user 0 indr 1048574 4 10.1.4.2 Push 299776 593 2 ge-0/0/1.0 31.0.0.0/24 user 0 indr 1048575 3 10.1.4.2 Push 299824 594 2 ge-0/0/1.0	

Path advertised to iBGP peers	regress@ASBR1-re> show route advertising-protocol bgp 128.49.114.112
	
	inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden)
	  Prefix Nexthop MED Lclpref AS path
	* 10.0.0.0/24 Self 100 10 I
	
	regress@ASBR1-re> show route advertising-protocol bgp 128.49.114.132
	
	inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden)
	  Prefix Nexthop MED Lclpref AS path
	* 10.0.0.0/24 Self 100 10 I
	
	regress@ASBR1-re> show route advertising-protocol bgp 100.1.0.2
	
	inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden)
	  Prefix Nexthop MED Lclpref AS path
	* 20.0.0.0/24 Self 20 I
	* 30.0.0.0/24 Self 30 I
	* 31.0.0.0/24 Self 30 I	

Table 1

ASBR3 (edge device)

Routes in RIB - summary	regress@ASBR3-re> show route summary Autonomous system number: 100 Router ID: 128.49.114.112 inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden) Direct: 4 routes, 4 active Local: 3 routes, 3 active BGP: 6 routes, 6 active Static: 4 routes, 4 active IS-IS: 8 routes, 8 active LDP: 1 routes, 1 active inet.3: 6 destinations, 6 routes (5 active, 0 holddown, 1 hidden) BGP: 3 routes, 2 active LDP: 3 routes, 3 active iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden) Direct: 1 routes, 1 active mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden) MPLS: 4 routes, 4 active LDP: 4 routes, 4 active VPN: 2 routes, 2 active inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden) Direct: 2 routes, 2 active INET6: 1 routes, 1 active	• Only default routing instances • Static routes are lab OOB • Some (2) BGP routes in inet.3. (networks pfx of customer eligible for protection/inspection) • Some (2) BGP (VPN) entries in mpls.0 RIB (plus 4 well-known labbes)
BGP routes	regress@ASBR3-re> show route protocol bgp inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both 10.0.0.0/24 *[BGP/170] 17:52:38, localpref 100, from 128.49.115.41 AS path: 10 I, validation-state: unverified > to 10.3.4.2 via ge-0/0/0.0, Push 299808 20.0.0.0/24 *[BGP/170] 17:52:46, localpref 100, from **128.49.114.132** AS path: 20 I, validation-state: unverified > to 10.3.4.2 via ge-0/0/0.0, Push **299776** 20.0.0.2/32 *[BGP/170] 17:52:46, localpref 100, from **128.49.114.132** AS path: 65500 I, validation-state: unverified > to 10.3.4.2 via ge-0/0/0.0, Push **299776** 30.0.0.0/24 *[BGP/170] 22:49:25, localpref 100 AS path: 30 I, validation-state: unverified > to **100.3.0.2 via ge-0/0/1.0** 30.0.0.2/32 *[BGP/170] 17:52:46, localpref 100, from **128.49.114.132**	• In inet.3 there are only prefixes of locally connected customer that are eligible for inspection/ protection. •

	AS path: 65500 I, validation-state: unverified > to 10.3.4.2 via ge-0/0/0.0, Push **299776** 31.0.0.0/24 *[BGP/170] 18:03:58, localpref 100 AS path: 30 I, validation-state: unverified > to 100.3.0.2 via ge-0/0/1.0 inet.3: 6 destinations, 6 routes (5 active, 0 holddown, 1 hidden) + = Active Route, - = Last Active, * = Both 30.0.0.0/24 *[BGP/170] 22:49:25, localpref 100 AS path: 30 I, validation-state: unverified > to **100.3.0.2 via ge-0/0/1.0** 31.0.0.0/24 *[BGP/170] 18:03:58, localpref 100 AS path: 30 I, validation-state: unverified > to **100.3.0.2 via ge-0/0/1.0** […] regress@ASBR3-re> show route protocol vpn […] mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both 299888 *[VPN/170] 17:53:02 > to 100.3.0.2 via **ge-0/0/1.0, Pop** 299888(S=0) *[VPN/170] 17:53:02 > to 100.3.0.2 via **ge-0/0/1.0, Pop**	
Prefixes in FIB's	% cli show route forwarding-table family inet | grep -A 1 "^[1-3][0-1]\.0\.0|Routing table|Destination" | grep -v "\-\-" Routing table: default.inet Internet: Destination Type RtRef Next hop Type Index NhRef Netif default perm 0 rjct 36 1 10.0.0.0/8 user 0 0:0:5e:0:1:50 ucst 341 6 fxp0.0 10.0.0.0/24 user 0 indr 1048575 2 10.3.4.2 Push 299808 596 2 ge-0/0/0.0 20.0.0.0/24 user 0 indr 1048574 4 10.3.4.2 Push 299776 595 2 ge-0/0/0.0 20.0.0.2/32 user 0 indr 1048574 4 10.3.4.2 Push 299776 595 2 ge-0/0/0.0 30.0.0.0/24 user 0 100.3.0.2 ucst 587 7 ge-0/0/1.0 30.0.0.2/32 user 0 indr 1048574 4 10.3.4.2 Push 299776 595 2 ge-0/0/0.0 31.0.0.0/24 user 0 100.3.0.2 ucst 587 7 ge-0/0/1.0	• There is no table derived form inet.3 • There are only default routing-instances

	regress@ASBR3-re> show route forwarding-table table default family mpls Routing table: default.mpls MPLS: Destination Type RtRef Next hop Type Index NhRef Netif default perm 0 dscd 50 1 0 user 0 recv 49 4 1 user 0 recv 49 4 2 user 0 recv 49 4 13 user 0 recv 49 4 299840 user 0 10.3.4.2 Pop 589 2 ge-0/0/0.0 299840(S=0) user 0 10.3.4.2 Pop 590 2 ge-0/0/0.0 299856 user 0 10.3.4.2 Swap 299776 591 2 ge-0/0/0.0 299872 user 0 10.3.4.2 Swap 299808 592 2 ge-0/0/0.0 **299888** user 0 100.3.0.2 **Pop** 593 2 ge-0/0/1.0 **299888**(S=0) user 0 100.3.0.2 **Pop** 594 2 ge-0/0/1.0	
Path advertised to iBGP peers	regress@ASBR3-re> show route advertising-protocol bgp 128.49.114.132 inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden) Prefix Nexthop MED Lclpref AS path * 30.0.0.0/24 Self 100 30 I * 31.0.0.0/24 Self 100 30 I [..] regress@ASBR3-re> show route advertising-protocol bgp 128.49.114.132 detail table inet.3 inet.3: 6 destinations, 6 routes (5 active, 0 holddown, 1 hidden) * 30.0.0.0/24 (1 entry, 1 announced) BGP group iBGP type Internal Route Label: **299888** Nexthop: Self Flags: Nexthop Change Localpref: 100 AS path: [100] 30 I Entropy label capable * 31.0.0.0/24 (1 entry, 1 announced) BGP group iBGP type Internal Route Label: **299888** Nexthop: Self Flags: Nexthop Change Localpref: 100 AS path: [100] 30 I Entropy label capable	• Prefixes of local customers are advertised w/o label and for protection/inspection eleigible also with label. • Note advertised label value and compare with MPLS FIB contenet above.

Table 2

ASBR2 (SC_PE)

Routes in RIB - summary	regress@ASBR2-re> show route summary Autonomous system number: 100 Router ID: 128.49.114.132 inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden) Direct: 5 routes, 5 active Local: 4 routes, 4 active BGP: 6 routes, 6 active Static: 4 routes, 4 active IS-IS: 7 routes, 7 active LDP: 1 routes, 1 active inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden) BGP: 3 routes, 3 active LDP: 3 routes, 3 active **off-ramp.inet.0:** 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden) Direct: 1 routes, 1 active Local: 1 routes, 1 active BGP: 3 routes, 3 active iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden) Direct: 1 routes, 1 active mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden) MPLS: 4 routes, 4 active LDP: 4 routes, 4 active VPN: 2 routes, 2 active inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden) Direct: 2 routes, 2 active INET6: 1 routes, 1 active	• Single non-default instance exist • Static routes are lab OOB • Some (3) BGP routes in inet.3. (networks pfx of customer eligible for protection/inspection) • Some (2) BGP (VPN) entries in mpls.0 RIB (plus 4 well-known labbes)
BGP routes	regress@ASBR2-re> show route protocol bgp inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both 10.0.0.0/24 *[BGP/170] 23:06:36, localpref 100, from 128.49.115.41 AS path: 10 I, validation-state: unverified > to 10.2.4.2 via ge-0/0/0.0, Push 299808 20.0.0.0/24 *[BGP/170] 23:16:29, localpref 100 AS path: 20 I, validation-state: unverified > to **100.2.0.2 via ge-0/0/1.0** **20.0.0.2/32** *[BGP/170] 23:16:26, localpref 100 AS path: 65500 I, validation-state: unverified	• Traffic toward protected targets send to screening device (100.2.0.6) • Routing to non-directly connected customer networks send with single label (LDP; 299824) • Routing to non-directly connected customer networks that are eligible for protection/inspection: send from off-ramp VR with dual labels (BGP-LU 299888, LDP 299824)

```
                    > to 100.2.0.6 via ge-0/0/2.0
30.0.0.0/24          *[BGP/170] 18:18:15, localpref
       100, from 128.49.114.112
                      AS path: 30 I, validation-
       state: unverified
                    > to 10.2.4.2 via ge-0/0/0.0,
       Push 299824
30.0.0.2/32          *[BGP/170] 23:16:26, localpref
       100
                      AS path: 65500 I, validation-
       state: unverified
                    > to 100.2.0.6 via ge-0/0/2.0
31.0.0.0/24          *[BGP/170] 18:18:15, localpref
       100, from 128.49.114.112
                      AS path: 30 I, validation-
       state: unverified
                    > to 10.2.4.2 via ge-0/0/0.0,
       Push 299824

inet.3: 6 destinations, 6 routes (6 active, 0
       holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.0.0.0/24          *[BGP/170] 23:16:29, localpref
       100
                      AS path: 20 I, validation-
       state: unverified
                    > to 100.2.0.2 via ge-0/0/1.0
30.0.0.0/24          *[BGP/170] 18:18:14, localpref
       100, from 128.49.114.112
                      AS path: 30 I, validation-
       state: unverified
                    > to 10.2.4.2 via ge-0/0/0.0,
       Push 299888, Push 299824(top)
31.0.0.0/24          *[BGP/170] 18:18:15, localpref
       100, from 128.49.114.112
                      AS path: 30 I, validation-
       state: unverified
                    > to 10.2.4.2 via ge-0/0/0.0,
       Push 299888, Push 299824(top)

off-ramp.inet.0: 5 destinations, 5 routes (5 active,
       0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.0.0.0/24          *[BGP/170] 23:16:29, localpref
       100
                      AS path: 20 I, validation-
       state: unverified
                    > to 100.2.0.2 via ge-0/0/1.0
30.0.0.0/24          *[BGP/170] 18:18:14, localpref
       100, from 128.49.114.112
                      AS path: 30 I, validation-
       state: unverified
                    > to 10.2.4.2 via ge-0/0/0.0,
       Push 299888, Push 299824(top)
31.0.0.0/24          *[BGP/170] 18:18:15, localpref
       100, from 128.49.114.112
                      AS path: 30 I, validation-
       state: unverified
                    > to 10.2.4.2 via ge-0/0/0.0,
       Push 299888, Push 299824(top)

iso.0: 1 destinations, 1 routes (1 active, 0
       holddown, 0 hidden)
```

	mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden) inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden) regress@ASBR2-re> show route protocol vpn [..] mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both 299856 *[VPN/170] 18:18:45 > to 100.2.0.2 via ge-0/0/1.0, Pop 299856(S=0) *[VPN/170] 18:18:45 > to 100.2.0.2 via ge-0/0/1.0, Pop inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)	
Prefixes in FIB's		•

Path advertised to iBGP peers	regress@ASBR2-re> show route advertising-protocol bgp 128.49.114.112 inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden) Prefix Nexthop MED Lclpref AS path * 20.0.0.0/24 Self 100 20 I * 20.0.0.2/32 Self 100 65500 I * 20.0.0.2/32 Self 100 65500 I inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden) Prefix Nexthop MED Lclpref AS path * 20.0.0.0/24 Self 100 20 I regress@ASBR2-re> show route advertising-protocol bgp 128.49.114.112 detail table inet.3 inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden) * 20.0.0.0/24 (1 entry, 1 announced) BGP group iBGP type Internal Route Label: 299856 Nexthop: Self Flags: Nexthop Change Localpref: 100 AS path: [100] 20 I Entropy label capable regress@ASBR2-re> show route advertising-protocol bgp 128.49.115.41 inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden) Prefix Nexthop MED Lclpref AS path * 20.0.0.0/24 Self 100 20 I * 20.0.0.2/32 Self 100 65500 I * 30.0.0.2/32 Self 100 65500 I	• Locally connected customer (protection eligible) pfx advertised with label – if screening cluster exist ion other location, this routes (lables) would be used for traffic re-insertion. • More specific prefixes (20.0.0.2/32, 30.0.0.2/3) of protected inspected targets are advertised as plain IPv4. When installed on other ASBR will attract traffic to screening device. • Only plain BGP, customer aggregated prefixes advertised to peering router (128.49.115.41).
Prefixes advertised to Screening device (off-ramp)	regress@ASBR2-re> show route advertising-protocol bgp 100.2.0.10 off-ramp.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden) Prefix Nexthop MED Lclpref AS path * 20.0.0.0/24 Self 20 I * 30.0.0.0/24 Self 30 I * 31.0.0.0/24 Self 30 I	•

Prefixes recived from Screening device (on-ramp)	regress@ASBR2-re> show route receive-protocol bgp 100.2.0.6 inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden) Prefix Nexthop MED Lclpref AS path * 20.0.0.2/32 100.2.0.6 65500 I * 30.0.0.2/32 100.2.0.6 65500 I	•

Table 3

Service Cluster configuration

To simulate the Services Cluster, an MX router was used.

The 1:1 sampling of entire transit traffic is configured in order to show that traffic is passing simulated screening device,

Please note that this configuration shows only routing and forwarding part, not actual packet processing that are subject of inspection.

eBGP session → on_ramp	protocols {
	bgp {
	group on-ramp {
	type external;
	import discard-all;
	export screen;
	peer-as 100;
	neighbor 100.2.0.5;
	}
	}
	}
	policy-options {
	policy-statement discard-all {
	then reject;
	}
	policy-statement screen {
	from community screen;
	then {
	community add no-export;
	accept;
	}
	}
	community no-export members no-export;
	community screen members 65500:1;
	}

eBGP session → off_ramp	protocols {
	bgp {
	group off-ramp {
	type external;

	export discard-all;
	peer-as 100;
	neighbor 100.2.0.9;
	}
	}
	}
	policy-options {
	policy-statement discard-all {
	then reject;
	}
	}

Static routes → destination system enabled for screening	routing-options {
	static {
	route 30.0.0.2/32 {
	discard;
	no-install;
	community 65500:1;
	}
	route 20.0.0.2/32 {
	discard;
	no-install;
	community 65500:1;
	}
	}
	autonomous-system 65500;
	}
	policy-options {
	community screen members 65500:1;
	}