




NOCTION
NETWORK INTELLIGENCE



Intelligent Routing Platform

Installation and Configuration Guide

[Version 4.2]



1901 S Bascom Ave, Campbell, CA 95008

Tel: 1-650-618-9823

Email: info@noction.com

Noction logos, and trademarks or registered trademarks of Noction Holdings Ltd or its subsidiaries in the United States and other countries.

Other names and brands may be claimed as the property of others. Information regarding third party products is provided solely for educational purposes. Noction Holdings Ltd is not responsible for the performance or support of third party products and does not make any representations or warranties whatsoever regarding quality, reliability, functionality, or compatibility of these devices or products.

Copyright ©2013 - 2024 Noction Inc., All Rights Reserved.

Contents

1	Introduction	17
1.1	How to contact support	17
1.2	What is IRP	17
1.2.1	IRP Features	17
1.2.2	IRP Components	18
1.2.3	IRP Technical Requirements	19
1.2.4	IRP Operating modes	23
1.2.5	BGP Monitoring	26
1.2.6	Outage detection	29
1.2.7	VIP Improvements	29
1.2.8	Retry Probing	30
1.2.9	Routing Policies	31
1.2.10	Support for Centralized Route Reflectors	34
1.2.11	Support for Internet Exchanges	34
1.2.12	Optimization for multiple Routing Domains	35
1.2.13	Improvements weight	41
1.2.14	Notifications and events	41
1.2.15	IRP API	44
1.2.16	IRP Failover	45
1.2.17	Inbound bandwidth optimization	49
1.2.18	Inbound performance optimization	51
1.2.19	Flowspec policies	52
1.2.20	Throttling excessive bandwidth use with Flowspec	53
1.2.21	Maintenance windows	53
1.2.22	Optimization of transiting traffic	54
1.2.23	Circuit issues detection	55
1.2.24	Threat Mitigation	56
1.3	IRP Optimization modes	56
1.3.1	Performance optimization	56
1.3.2	Cost optimization	57
1.3.3	Commit Control	58
1.4	What is IRP Global Management Interface	62
1.4.1	GMI Operating Details	62
1.4.2	GMI Technical Requirements	62
2	Configuration	64
2.1	IRP software management	64
2.2	Configuration files	65
2.3	IRP Global Management Interface Configuration	66
2.3.1	GMI Initial Configuration	66
2.3.2	GMI software management	66
2.3.3	Starting, stopping and getting status of GMI backend	66
2.4	Starting, stopping and getting status of IRP components	67
2.5	IRP Initial Configuration	68
2.6	Global and Core Configuration	69
2.7	Collector Configuration	70
2.7.1	Irpflowd Configuration	70
2.7.2	Irpspand Configuration	74

2.8	Explorer Configuration	76
2.8.1	Specific PBR configuration scenarios	76
2.8.2	Flowspec PBR	89
2.9	Bgpd Configuration	90
2.9.1	AS-Path behavior in IRP Bgpd	94
2.9.2	Bgpd online reconfiguration	94
2.9.3	BGP Additional paths	95
2.10	BMP Configuration	95
2.11	Threat Mitigation Configuration	96
2.11.1	BGP Blackholing	96
2.11.2	BGP Redirect	96
2.11.3	FlowSpec Drop	96
2.11.4	FlowSpec Redirect	96
2.11.5	Mitigation prefix size (IPv4 / IPv6)	96
2.12	Administrative Components	97
2.13	Failover Configuration	97
2.13.1	Initial failover configuration	97
2.13.2	Re-purpose operational IRP node into an IRP failover slave	106
2.13.3	Re-purpose operational IRP failover slave into a new master	106
2.13.4	Recover prolonged node downtime or corrupted replication	106
3	Using IRP	108
3.1	Getting started	108
3.1.1	Accessing the system	108
3.1.2	System dashboard	108
3.2	Frontend structure	110
3.2.1	Sidebar	110
3.2.2	Failover status	113
3.2.3	Global search	114
3.2.4	Warning bars	115
3.3	Customizing the dashboard	115
3.3.1	Custom dashboard gadgets	116
3.4	Reports	118
3.4.1	Platform overview	118
3.4.2	Platform performance	119
3.4.3	Loss Improvements	119
3.4.4	Latency Improvements	121
3.4.5	Probe conversion rates	123
3.4.6	Current improvements	123
3.4.7	Improvements to Exchanges	125
3.4.8	Historical records	125
3.4.9	Providers overall	126
3.4.10	Provider performance	126
3.4.11	Cost overview	128
3.4.12	Congestion and outages	128
3.4.13	Top volume prefixes	129
3.4.14	Top volume AS	129
3.4.15	Top problem AS	130
3.4.16	Prefix statistics	130
3.4.17	AS statistics	132
3.4.18	Country statistics	132
3.4.19	Exchanges statistics	133
3.4.20	Current inbound improvements	134
3.4.21	Inbound Improvements history	134
3.4.22	Probes today	135
3.4.23	Monitors status	135
3.5	Graphs	137
3.5.1	Improvements by type	137
3.5.2	Performance improvements	137

3.5.3	Probed prefixes and Improvements	138
3.5.4	New and retry improvements	138
3.5.5	Improvements by probing source	139
3.5.6	Prefixes rerouted from provider	139
3.5.7	Prefixes rerouted to provider	140
3.5.8	Improvements by provider	140
3.5.9	Performance improvements by provider	141
3.5.10	Commit improvements by provider	142
3.5.11	Total and Improved traffic	142
3.5.12	Bandwidth usage by provider	143
3.5.13	Providers bandwidth usage	143
3.5.14	Bandwidth usage and improvements	144
3.5.15	Total bandwidth usage	145
3.5.16	Inbound traffic distribution	145
3.5.17	Provider performance history	145
3.6	Routing Policies	146
3.7	Flowspec policies	149
3.8	Throttling excessive bandwidth use with Flowspec	151
3.9	Maintenance windows	151
3.10	Events	153
3.11	Troubleshooting	155
3.11.1	Traceroute	155
3.11.2	Looking glass	156
3.11.3	Whois	157
3.11.4	Prefix probing	157
3.12	Configuration editor	159
3.12.1	Global configuration	160
3.12.2	BGP and Routers configuration	163
3.12.3	Collector configuration	166
3.12.4	Core configuration	168
3.12.5	Commit Control configuration	172
3.12.6	Explorer configuration	173
3.12.7	Providers and Peers configuration	174
3.12.8	SNMP hosts configuration	182
3.12.9	Notifications and Events	183
3.12.10	Security Configuration	188
3.12.11	Frontend configuration	194
3.12.12	Inbound configuration	195
3.12.13	Wizards	202
4	Using IRP Global Management Interface	216
4.1	Getting started	216
4.1.1	Accessing the system	216
4.2	Frontend structure	217
4.2.1	Right Sidebar	218
4.2.2	IRP Instance Performance Stats	219
4.2.3	Global Search	220
4.2.4	Warning bars	221
4.3	GMI Dashboards	221
4.3.1	Creating a new dashboard	222
4.3.2	Managing Dashboards	224
4.3.3	Deleting Dashboards	225
4.4	Reports	226
4.4.1	ASN Path Problems	227
4.4.2	ASN Statistics	228
4.4.3	ASN Traffic Percentage	228
4.4.4	Country Statistics	229
4.4.5	Current Improvements	230
4.4.6	Exchanges Statistics	230

4.4.7	Historical Records	231
4.4.8	Improved Prefixes Within Probed	231
4.4.9	Improved Traffic Volume	232
4.4.10	Improvements to Exchanges	233
4.4.11	Instances Status (available as a widget only)	233
4.4.12	Monitor History	233
4.4.13	Monitor Status	234
4.4.14	Platform Overview	234
4.4.15	Prefix Improvement Rate by Cause	235
4.4.16	Prefix Statistics	235
4.4.17	Probes Today	236
4.4.18	Providers Efficiency	237
4.4.19	Providers Performance and Probes (last 24 hours)	237
4.4.20	Top Problem ASN	237
4.4.21	Top Volume ASN	238
4.4.22	Top Volume Prefixes	239
4.5	Graphs	239
4.5.1	ASN Traffic Heatmap	240
4.5.2	Bandwidth Usage and Improvements	240
4.5.3	Bandwidth Usage per Provider	241
4.5.4	Commit Control Improvements by Provider	241
4.5.5	Improvements by Cause	241
4.5.6	Improvements by IP Version	242
4.5.7	Improvements by Probing Source	242
4.5.8	Latency Destination	243
4.5.9	Latency Improvements by Provider	243
4.5.10	Latency Values (ms)	244
4.5.11	Loss Improvements by Provider	244
4.5.12	Loss Rates	245
4.5.13	New and Retry Improvements	246
4.5.14	Overall Improvements Type by Provider	246
4.5.15	Overall Rerouted Prefixes by Provider	246
4.5.16	Performance Improvements	247
4.5.17	Prefixes Rerouted from Provider	248
4.5.18	Prefixes Rerouted to Provider	248
4.5.19	Probed and Improved Volumes	248
4.5.20	Probed Prefixes and Improvements	249
4.5.21	Probes per Provider	249
4.5.22	Probes Today	250
4.5.23	Problem Destinations	250
4.5.24	Provider Performance	251
4.5.25	Providers Bandwidth Usage	251
4.5.26	Providers Efficiency	252
4.5.27	Total and Improved Traffic	252
4.5.28	Total Bandwidth Usage	253
4.5.29	Unique Probed/Improved Prefixes	253
4.6	Inbound	254
4.6.1	Inbound Commit Control	254
4.6.2	Inbound Performance Optimization	257
4.7	Routing Policies	261
4.8	Flowspec Policies	263
4.8.1	Policies by Country	265
4.9	Troubleshooting	266
4.9.1	Looking glass	266
4.9.2	Prefix probing	267
4.9.3	Traceroute	268
4.9.4	Whois	268

4.10	Threat Mitigation	269
4.10.1	Monitor	269
4.10.2	Mitigation Rules	270
4.10.3	Mitigation Feed	272
4.10.4	History	273
4.11	Global Commit	273
4.11.1	Global Commit Configuration	273
4.11.2	Global Commit Summary	274
4.11.3	Global Commit Graph	275
4.12	Configuration editor	275
4.12.1	Global configuration	276
4.12.2	BGP and Routers configuration	277
4.12.3	Collector configuration	280
4.12.4	Core configuration	281
4.12.5	Commit Control configuration	284
4.12.6	Inbound performance optimization configuration	287
4.12.7	Explorer configuration	287
4.12.8	Providers and Peers configuration	288
4.12.9	SNMP hosts configuration	296
4.13	Management	299
4.13.1	Instances	299
4.13.2	User Management	300
4.13.3	Access Restriction	304
4.13.4	Frontend	305
4.13.5	Senders	305
4.14	Emailing subscription	305
4.15	Events	306
4.16	Notifications and Events	307
4.16.1	Events	307
4.16.2	SNMP Traps	309
4.16.3	Notifications	310
4.17	User Profile	311
4.18	Maintenance windows	312
5	Configuration parameters reference	315
5.1	Database credentials	315
5.1.1	db.clickhouse.dbname	315
5.1.2	db.clickhouse.host	315
5.1.3	db.clickhouse.http_port	315
5.1.4	db.clickhouse.password	315
5.1.5	db.clickhouse.port	315
5.1.6	db.clickhouse.username	315
5.1.7	db.dbname	316
5.1.8	db.host	316
5.1.9	db.ourhost	316
5.1.10	db.password	316
5.1.11	db.port	316
5.1.12	db.username	316
5.2	Global parameters	317
5.2.1	global.agg_ipv4_max	317
5.2.2	global.agg_ipv6_max	317
5.2.3	global.aggregate	317
5.2.4	global.bw_overusage	317
5.2.5	global.exchanges	317
5.2.6	global.exchanges.auto_config_interval	318
5.2.7	global.failover	318
5.2.8	global.failover.log	318
5.2.9	global.failover.use_communities	318
5.2.10	global.failover_identity_file	318

5.2.11	global.failover_role	318
5.2.12	global.failover_slave.ip	319
5.2.13	global.failover_slave.ipv6	319
5.2.14	global.failover_slave.port	319
5.2.15	global.failover_timer_fail	319
5.2.16	global.failover_timer_failback	319
5.2.17	global.flowspec	319
5.2.18	global.flowspec.pbr	320
5.2.19	global.frontend_acl	320
5.2.20	global.frontend_acl_ips	320
5.2.21	global.ignored.asn	320
5.2.22	global.ignored.unannounced	320
5.2.23	global.ignored_communities	321
5.2.24	global.ignorednets	321
5.2.25	global.improve_mode	321
5.2.26	global.inbound.injection	322
5.2.27	global.inbound_conf	322
5.2.28	global.inbound_transit	322
5.2.29	global.ipv6_enabled	322
5.2.30	global.master_management_interface	323
5.2.31	global.master_probing_interface	323
5.2.32	global.master_rd	323
5.2.33	global.nonintrusive_bgp	323
5.2.34	global.offpeak_hour	323
5.2.35	global.outbound	324
5.2.36	global.png.datadir	324
5.2.37	global.policies	324
5.2.38	global.policy.unpack_max	324
5.2.39	global.private_key	324
5.2.40	global.public_key	324
5.2.41	global.rd_rtt	325
5.2.42	global.rrd.age_max	325
5.2.43	global.rrd.datadir	325
5.2.44	global.slave_management_interface	325
5.2.45	global.slave_probing_interface	326
5.2.46	global.slave_rd	326
5.2.47	global.user_directories_conf	326
5.3	API daemon settings	326
5.3.1	apid.allowed_ips	326
5.3.2	apid.grpc.port	326
5.3.3	apid.listen.master_ip	326
5.3.4	apid.listen.port	327
5.3.5	apid.listen.slave_ip	327
5.3.6	apid.log	327
5.3.7	apid.log.level	327
5.3.8	apid.maxthreads	327
5.3.9	apid.path.mtr	327
5.3.10	apid.path.ping	328
5.3.11	apid.path.ping6	328
5.3.12	apid.path.traceroute	328
5.4	Bgpd settings	328
5.4.1	bgpd.as_path	328
5.4.2	bgpd.as_path_borrowing	329
5.4.3	bgpd.db.timeout.withdraw	329
5.4.4	bgpd.full_control	330
5.4.5	bgpd.improvements.remove.bmp_exact_aggregate	330
5.4.6	bgpd.improvements.remove.hold_time	330
5.4.7	bgpd.improvements.remove.next_hop_eq	330

5.4.8	bgpd.improvements.remove.withdrawn	331
5.4.9	bgpd.improvements.strip_non_irp_communities	331
5.4.10	bgpd.log	331
5.4.11	bgpd.log.level	331
5.4.12	bgpd.mon.guardtime	331
5.4.13	bgpd.mon.holdtime	332
5.4.14	bgpd.mon.internal.flap_guardtime	332
5.4.15	bgpd.mon.keepalive	332
5.4.16	bgpd.mon.longholdtime	332
5.4.17	bgpd.monitor.type	332
5.4.18	bgpd.no_export	333
5.4.19	bgpd.policy.cascade.amount	333
5.4.20	bgpd.prefix.monitor.interval	333
5.4.21	bgpd.prefix.monitor.search_interval	333
5.4.22	bgpd.prefixlist.asn	334
5.4.23	bgpd.prefixlist.prefixes	334
5.4.24	bgpd.rd_local_mark	334
5.4.25	bgpd.retry_probing.new.bmp_path_change	334
5.4.26	bgpd.retry_probing.old.bmp_path_change	335
5.4.27	bgpd.scaninterval	335
5.4.28	bgpd.snmp.concurrent_requests	335
5.4.29	bgpd.snmp.packets_interval	335
5.4.30	bgpd.snmp.simultaneous	336
5.4.31	bgpd.transit.monitor.election_interval	336
5.4.32	bgpd.transit.monitor.fast_reconfirm_interval	336
5.4.33	bgpd.transit.monitor.retries	336
5.4.34	bgpd.transit.monitor.timeout	336
5.4.35	bgpd.updates.split	336
5.5	BGP sessions settings	337
5.5.1	bgpd.peer.X.as	337
5.5.2	bgpd.peer.X.blackholing.ipv4.next_hop	337
5.5.3	bgpd.peer.X.blackholing.ipv6.next_hop	337
5.5.4	bgpd.peer.X.blackholing.localpref	337
5.5.5	bgpd.peer.X.cap_4byte_as	337
5.5.6	bgpd.peer.X.flowspec	338
5.5.7	bgpd.peer.X.flowspec.redirect_type	338
5.5.8	bgpd.peer.X.inbound.ipv4.next_hop	338
5.5.9	bgpd.peer.X.inbound.ipv6.next_hop	339
5.5.10	bgpd.peer.X.inbound.master_localpref	339
5.5.11	bgpd.peer.X.inbound.slave_localpref	339
5.5.12	bgpd.peer.X.keepalive	339
5.5.13	bgpd.peer.X.listen	339
5.5.14	bgpd.peer.X.master_communities	340
5.5.15	bgpd.peer.X.master_localpref	340
5.5.16	bgpd.peer.X.master_our_ip	340
5.5.17	bgpd.peer.X.master_our_ipv6	340
5.5.18	bgpd.peer.X.master_password	341
5.5.19	bgpd.peer.X.master_peer_ip	341
5.5.20	bgpd.peer.X.master_peer_ipv6	341
5.5.21	bgpd.peer.X.master_router_id	341
5.5.22	bgpd.peer.X.med	342
5.5.23	bgpd.peer.X.origin	342
5.5.24	bgpd.peer.X.shutdown	342
5.5.25	bgpd.peer.X.slave_communities	342
5.5.26	bgpd.peer.X.slave_localpref	342
5.5.27	bgpd.peer.X.slave_our_ip	342
5.5.28	bgpd.peer.X.slave_our_ipv6	343
5.5.29	bgpd.peer.X.slave_password	343

5.5.30	bgpd.peer.X.slave_peer_ip	344
5.5.31	bgpd.peer.X.slave_peer_ipv6	344
5.5.32	bgpd.peer.X.slave_router_id	344
5.5.33	bgpd.peer.X.transit.mib	344
5.5.34	bgpd.peer.X.transit.snmp	345
5.5.35	bgpd.peer.X.transit.status	345
5.5.36	bgpd.peer.X.updates.limit.max	345
5.5.37	bgpd.peer.X.updates.limit.ps	345
5.6	BMP monitoring station settings	345
5.6.1	irpbmpd.grpc.port	345
5.6.2	irpbmpd.log	346
5.6.3	irpbmpd.log.level	346
5.6.4	irpbmpd.port	346
5.6.5	irpbmpd.sources	346
5.7	Collector settings	346
5.7.1	collector.detect.explorer_ips	346
5.7.2	collector.export.archive_inbound	347
5.7.3	collector.export.archive_transit	347
5.7.4	collector.export.top_volume_ips	347
5.7.5	collector.export.ttl	347
5.7.6	collector.export.volume.high.top_n	347
5.7.7	collector.export.volume.min	347
5.7.8	collector.export.volume.min_pct	348
5.7.9	collector.flow.all_outbound	348
5.7.10	collector.flow.buffer.size	348
5.7.11	collector.flow.enabled	348
5.7.12	collector.flow.export.inbound_transit.topn	348
5.7.13	collector.flow.listen.nf	349
5.7.14	collector.flow.listen.sf	349
5.7.15	collector.flow.log	349
5.7.16	collector.flow.log.level	349
5.7.17	collector.flow.process_transit_in_outbound	349
5.7.18	collector.flow.sources	349
5.7.19	collector.ournets	350
5.7.20	collector.sessions.max	350
5.7.21	collector.snmp.enhanced.sec	350
5.7.22	collector.span.buffer.size	350
5.7.23	collector.span.enabled	351
5.7.24	collector.span.interfaces	351
5.7.25	collector.span.log	351
5.7.26	collector.span.log.level	351
5.7.27	collector.span.min_delay	351
5.7.28	collector.span.min_delay.probing_queue_size	352
5.7.29	collector.span.size_from_ip_header	352
5.7.30	collector.span.threshold.blackout	352
5.7.31	collector.span.threshold.congestion	352
5.7.32	collector.span.threshold.delay	352
5.7.33	collector.span.threshold.excessive	353
5.7.34	collector.speaking_ips	353
5.8	Core settings	354
5.8.1	core.circuit.high_loss_diff	354
5.8.2	core.circuit.hist_interval	354
5.8.3	core.circuit.inbound	354
5.8.4	core.circuit.recover_hold_time	354
5.8.5	core.circuit.recover_loss_diff	355
5.8.6	core.circuit.recover_monitored_intervals	355
5.8.7	core.circuit.transit	355
5.8.8	core.circuit.warn_loss_diff	355

5.8.9	core.circuit.withdraw_on_warn	356
5.8.10	core.commit_control	356
5.8.11	core.commit_control.agg_bw_min	356
5.8.12	core.commit_control.del_irrelevant	357
5.8.13	core.commit_control.group_balance.check_monitor	357
5.8.14	core.commit_control.inbound.enabled	357
5.8.15	core.commit_control.inbound.improvement.delay	357
5.8.16	core.commit_control.inbound.moderated	358
5.8.17	core.commit_control.inbound.rate.high	358
5.8.18	core.commit_control.inbound.rate.low	358
5.8.19	core.commit_control.inbound.volume_estimation	358
5.8.20	core.commit_control.loss_override	358
5.8.21	core.commit_control.probe_ttl	359
5.8.22	core.commit_control.probing_queue_size	359
5.8.23	core.commit_control.rate.group	359
5.8.24	core.commit_control.rate.high	359
5.8.25	core.commit_control.rate.low	360
5.8.26	core.commit_control.react_on_collector	360
5.8.27	core.commit_control.worst_loss	360
5.8.28	core.cost.worst_ms	361
5.8.29	core.eventqueueulimit	361
5.8.30	core.eventqueueulimit.retry_probe_pct	361
5.8.31	core.eventqueueulimit_ipv6	361
5.8.32	core.flowspec.max	361
5.8.33	core.flowspec.max_ipv6	362
5.8.34	core.global.allow_commit	362
5.8.35	core.global.allow_latency_cost	362
5.8.36	core.global.worst_ms	362
5.8.37	core.improvements.clearslots	362
5.8.38	core.improvements.clearslots.days_max	363
5.8.39	core.improvements.inbound_transit.max	363
5.8.40	core.improvements.inbound_transit.ttl.clean	363
5.8.41	core.improvements.inbound_transit.ttl.max	363
5.8.42	core.improvements.inbound_transit.ttl.min	364
5.8.43	core.improvements.max	364
5.8.44	core.improvements.max_ipv6	364
5.8.45	core.improvements.safe_removal	364
5.8.46	core.improvements.ttl.retry_probe	364
5.8.47	core.log	365
5.8.48	core.log.level	365
5.8.49	core.outage_detection	365
5.8.50	core.outage_detection.limit_pct	365
5.8.51	core.outage_detection.volume_top_n	365
5.8.52	core.overusage.check_interval	365
5.8.53	core.overusage.hold_timer	366
5.8.54	core.overusage.out.average.period	366
5.8.55	core.overusage.out.average.relevant_min	366
5.8.56	core.overusage.out.threshold.throttle	366
5.8.57	core.overusage.out.threshold.trigger	366
5.8.58	core.performance.loss_pct	367
5.8.59	core.performance.rtt.diff_ms	367
5.8.60	core.performance.rtt.diff_pct	367
5.8.61	core.performance.rtt.ix_diff_ms	367
5.8.62	core.performance.rtt.ix_diff_pct	368
5.8.63	core.probes.ttl.failed	368
5.8.64	core.probes.ttl.max	368
5.8.65	core.probes.ttl.min	368
5.8.66	core.problem.outage_timeout	369

5.8.67	core.problem.rtt.diff_pct	369
5.8.68	core.vip.interval.probe	369
5.9	Explorer settings	369
5.9.1	explorer.aipi	369
5.9.2	explorer.algorithms	370
5.9.3	explorer.high_vol_precedence	370
5.9.4	explorer.infra_ips	370
5.9.5	explorer.interval.infra	370
5.9.6	explorer.interval.other	370
5.9.7	explorer.interval.other.trace	371
5.9.8	explorer.ipv4_test	371
5.9.9	explorer.ipv6_test	371
5.9.10	explorer.log	371
5.9.11	explorer.log.level	371
5.9.12	explorer.max_collector_ips	371
5.9.13	explorer.maxthreads	371
5.9.14	explorer.probe.algorithm	372
5.9.15	explorer.probe.indirect_priority	372
5.9.16	explorer.probing.sendpkts.adaptive_max	372
5.9.17	explorer.probing.sendpkts.min	372
5.9.18	explorer.probing.simultaneous	373
5.9.19	explorer.scanning.confirm_ips	373
5.9.20	explorer.scanning.replypkts.min	373
5.9.21	explorer.scanning.rtt.dispersion_ms	373
5.9.22	explorer.scanning.sendpkts.factor	373
5.9.23	explorer.timeout	373
5.9.24	explorer.timeout.infra	374
5.9.25	explorer.trace.algorithms	374
5.9.26	explorer.trace.all	374
5.9.27	explorer.trace.diag_hop_unique	374
5.9.28	explorer.traceroute.retrypkts	375
5.9.29	explorer.traceroute.sendpkts	375
5.9.30	explorer.traceroute.simultaneous	375
5.9.31	explorer.traceroute.simultaneous.infra	375
5.9.32	explorer.traceroute.ttl.max	375
5.9.33	explorer.traceroute.ttl.min	375
5.10	Inbound Performance	376
5.10.1	irpinperfd.enabled	376
5.10.2	irpinperfd.log	376
5.10.3	irpinperfd.log.level	376
5.10.4	irpinperfd.model.debug	376
5.10.5	irpinperfd.model.topn_per_rule	376
5.10.6	irpinperfd.moderated	376
5.10.7	irpinperfd.probing.confirmation_interval	376
5.10.8	irpinperfd.probing.failure_margin	377
5.10.9	irpinperfd.probing.interval	377
5.10.10	irpinperfd.probing.shelf_life	377
5.10.11	irpinperfd.probing.timeout	377
5.11	Threat mitigation	377
5.11.1	irpdetectd.bgp.reaction	377
5.11.2	irpdetectd.bgp.redirect.bgp_peers	377
5.11.3	irpdetectd.bgp.redirect.communities	377
5.11.4	irpdetectd.blackhole.threshold.kpps	378
5.11.5	irpdetectd.blackhole.threshold.mbps	378
5.11.6	irpdetectd.flowspec.ipv4.redirect	378
5.11.7	irpdetectd.flowspec.ipv6.redirect	378
5.11.8	irpdetectd.flowspec.reaction	378
5.11.9	irpdetectd.flowspec.threshold.kpps	379

5.11.10	irpdetectd.flowspec.threshold.mbps	379
5.11.11	irpdetectd.ip_batch_size	379
5.11.12	irpdetectd.ipv4.prefix_size	379
5.11.13	irpdetectd.ipv6.prefix_size	379
5.11.14	irpdetectd.mode	379
5.11.15	irpdetectd.protected_addresses	380
5.11.16	irpdetectd.time.keep	380
5.11.17	irpdetectd.time.monitor	380
5.11.18	irpdetectd.whitelist	380
5.12	Notification and events	380
5.12.1	pushd.email.from	380
5.12.2	pushd.email.host	380
5.12.3	pushd.email.port	380
5.12.4	pushd.listen.port	381
5.12.5	pushd.log	381
5.12.6	pushd.log.level	381
5.12.7	pushd.sms.account_sid	381
5.12.8	pushd.sms.auth_token	381
5.12.9	pushd.sms.gateway	381
5.12.10	pushd.sms.message_size	381
5.12.11	pushd.sms.phone_number	382
5.12.12	pushd.sms.uri.plivo	382
5.12.13	pushd.sms.uri.twilio	382
5.12.14	pushd.templates.datadir	382
5.12.15	pushd.webhook.avatar_emoji	382
5.12.16	pushd.webhook.avatar_url	383
5.12.17	pushd.webhook.botname	383
5.12.18	pushd.webhook.url	383
5.12.19	trap.bgpd_announced_rate_low.limit_pct	383
5.12.20	trap.core_cc_improvements_spike.diff_pct	383
5.12.21	trap.core_cc_improvements_spike.period_sec	384
5.12.22	trap.core_cc_overload.inbound_limit_mbps	384
5.12.23	trap.core_cc_overload.inbound_limit_pct	384
5.12.24	trap.core_cc_overload.limit_mbps	384
5.12.25	trap.core_cc_overload.limit_pct	384
5.12.26	trap.core_cc_provider_overload.inbound_limit_mbps	385
5.12.27	trap.core_cc_provider_overload.inbound_limit_pct	385
5.12.28	trap.core_cc_provider_overload.limit_mbps	385
5.12.29	trap.core_cc_provider_overload.limit_pct	385
5.12.30	trap.core_excessive.prefixes	385
5.12.31	trap.core_excessive_latency.limit	386
5.12.32	trap.core_excessive_loss.limit	386
5.12.33	trap.destination.auth_password	386
5.12.34	trap.destination.auth_protocol	386
5.12.35	trap.destination.auth_username	386
5.12.36	trap.destination.community	386
5.12.37	trap.destination.port	386
5.12.38	trap.destination.priv_password	387
5.12.39	trap.destination.priv_protocol	387
5.12.40	trap.destination.seclvl	387
5.12.41	trap.destination.version	387
5.13	Administrative settings	387
5.13.1	dbcron.api.log	387
5.13.2	dbcron.api.socket	388
5.13.3	dbcron.log	388
5.13.4	dbcron.log.level	388
5.13.5	globalcc.key_service.port	388
5.13.6	globalcc.log	388

5.13.7	globalcc.log.level	388
5.13.8	globalcc.service.port	388
5.13.9	irpdetected.log	388
5.13.10	irpdetected.log.level	389
5.13.11	irpmng.grpc.port	389
5.14	Provider	389
5.14.1	peer.X.95th	389
5.14.2	peer.X.95th.bill_day	389
5.14.3	peer.X.95th.centile	390
5.14.4	peer.X.95th.in	390
5.14.5	peer.X.95th.mode	390
5.14.6	peer.X.aspath_for_ix	390
5.14.7	peer.X.auto_config	390
5.14.8	peer.X.bgp_peer	391
5.14.9	peer.X.blackholing.bgp_peer	391
5.14.10	peer.X.blackholing.community	391
5.14.11	peer.X.bmp	391
5.14.12	peer.X.bmp.check_routes	391
5.14.13	peer.X.cc_disable	391
5.14.14	peer.X.circuit.control	392
5.14.15	peer.X.cost	392
5.14.16	peer.X.description	392
5.14.17	peer.X.diag_hop.interval_max	392
5.14.18	peer.X.diag_hop.interval_min	393
5.14.19	peer.X.flow_agents	393
5.14.20	peer.X.flowspec.ipv4.redirect_community	393
5.14.21	peer.X.flowspec.ipv6.redirect_community	393
5.14.22	peer.X.flowspec.pbr.enabled	393
5.14.23	peer.X.flowspec.pbr.use_bgp_peer	393
5.14.24	peer.X.global_group	394
5.14.25	peer.X.group_loadbalance	394
5.14.26	peer.X.improve_in_group	394
5.14.27	peer.X.inbound.community_base	394
5.14.28	peer.X.ipv4.diag_hop	394
5.14.29	peer.X.ipv4.master_probing	395
5.14.30	peer.X.ipv4.mon	395
5.14.31	peer.X.ipv4.next_hop	395
5.14.32	peer.X.ipv4.next_hop_as	395
5.14.33	peer.X.ipv4.route_server	396
5.14.34	peer.X.ipv4.slave_probing	396
5.14.35	peer.X.ipv4_pbr_check	396
5.14.36	peer.X.ipv6.diag_hop	396
5.14.37	peer.X.ipv6.master_probing	396
5.14.38	peer.X.ipv6.mon	396
5.14.39	peer.X.ipv6.next_hop	397
5.14.40	peer.X.ipv6.next_hop_as	397
5.14.41	peer.X.ipv6.route_server	397
5.14.42	peer.X.ipv6.slave_probing	397
5.14.43	peer.X.ipv6_pbr_check	397
5.14.44	peer.X.limit_load	397
5.14.45	peer.X.mon.ipv4.bgp_peer	398
5.14.46	peer.X.mon.ipv4.external.state	398
5.14.47	peer.X.mon.ipv4.internal.mode	398
5.14.48	peer.X.mon.ipv4.internal.state	398
5.14.49	peer.X.mon.ipv6.bgp_peer	398
5.14.50	peer.X.mon.ipv6.external.state	398
5.14.51	peer.X.mon.ipv6.internal.mode	399
5.14.52	peer.X.mon.ipv6.internal.state	399

5.14.53	peer.X.mon.snmp	399
5.14.54	peer.X.mon.vpn_instance_id	399
5.14.55	peer.X.pbr_check	399
5.14.56	peer.X.precedence	400
5.14.57	peer.X.rd	401
5.14.58	peer.X.routes_config	401
5.14.59	peer.X.shortname	401
5.14.60	peer.X.shutdown	401
5.14.61	peer.X.snmp.enhanced	401
5.14.62	peer.X.snmp.interfaces	402
5.14.63	peer.X.type	402
5.15	Inbound rule	402
5.15.1	inbound.rule.X.bgp_peer	402
5.15.2	inbound.rule.X.communities	402
5.15.3	inbound.rule.X.enabled	402
5.15.4	inbound.rule.X.full_control	402
5.15.5	inbound.rule.X.next_hop	403
5.15.6	inbound.rule.X.prefix	403
5.15.7	inbound.rule.X.providers	403
5.16	Routing Policy	403
5.16.1	policy.X.asn	404
5.16.2	policy.X.cascade	404
5.16.3	policy.X.community	404
5.16.4	policy.X.country	404
5.16.5	policy.X.enabled	404
5.16.6	policy.X.forcelocal	404
5.16.7	policy.X.notes	404
5.16.8	policy.X.policy	405
5.16.9	policy.X.prefix	405
5.16.10	policy.X.priority	405
5.16.11	policy.X.providers	405
5.16.12	policy.X.vip	405
5.17	Exchange	406
5.17.1	provider.X.rule.Y.bgp_peer	406
5.17.2	provider.X.rule.Y.enabled	406
5.17.3	provider.X.rule.Y.next_hop	406
5.17.4	provider.X.rule.Y.next_hop_as	406
5.17.5	provider.X.rule.Y.pbr_check	406
5.17.6	provider.X.rule.Y.probing_dscp	406
5.17.7	provider.X.rule.Y.probing_ip	407
5.17.8	provider.X.rule.Y.shortname	407
5.17.9	provider.X.rule.Y.transit	407
5.18	Global Group	407
5.18.1	globalgroup.X.members	407
5.18.2	globalgroup.X.outbound.95th	407
5.18.3	globalgroup.X.outbound.95th.reserve	407
5.18.4	globalgroup.X.outbound.rate_high	407
5.18.5	globalgroup.X.outbound.rate_low	408
5.18.6	globalgroup.X.shortname	408
5.19	User Directory	408
5.19.1	directory.X.admin_role_groups	408
5.19.2	directory.X.base_dn	408
5.19.3	directory.X.email	408
5.19.4	directory.X.fullname	408
5.19.5	directory.X.hostname	408
5.19.6	directory.X.initial_bind_dn	409
5.19.7	directory.X.initial_bind_password	409
5.19.8	directory.X.name	409

5.19.9	directory.X.order	409
5.19.10	directory.X.port	409
5.19.11	directory.X.secret	409
5.19.12	directory.X.state	409
5.19.13	directory.X.timeout	409
5.19.14	directory.X.tls	410
5.19.15	directory.X.tls_cacertfile	410
5.19.16	directory.X.type	410
5.19.17	directory.X.user_dn	410
5.19.18	directory.X.user_filter	410
5.19.19	directory.X.user_role_attr	410
5.19.20	directory.X.username	410
5.19.21	directory.X.verify_cert	410
5.20	SNMP Host	411
5.20.1	snmp.X.auth_password	411
5.20.2	snmp.X.auth_protocol	411
5.20.3	snmp.X.auth_username	411
5.20.4	snmp.X.community	411
5.20.5	snmp.X.context	411
5.20.6	snmp.X.ip	411
5.20.7	snmp.X.name	411
5.20.8	snmp.X.port	412
5.20.9	snmp.X.priv_password	412
5.20.10	snmp.X.priv_protocol	412
5.20.11	snmp.X.seclevel	412
5.20.12	snmp.X.timeout	412
5.20.13	snmp.X.version	412
5.21	Routing domain	413
5.21.1	rd.X.community.local	413
5.21.2	rd.X.community_worsening	413
5.21.3	rd.X.shortname	413
6	Appendixes	414
6.1	Frontend labels for configuration parameters	414
6.2	Configuration parameter index	423

Chapter 1

Introduction

1.1 How to contact support

If you encounter any problems while using or setting up the Intelligent Routing Platform, please contact us.

You may contact us in the following ways:

- Email us at support@noction.com
- Open a support request at <https://helpdesk.noction.com>

1.2 What is IRP

BGP is a fundamental technology for the fault-tolerance of the Internet, which chooses network paths based on the number of hops traffic must traverse before it reaches its destination. However, BGP does not take into account the important factors of network performance. Even if multi-homing does provide some redundancy, multiple network outages have shown that multi-homing alone is not the solution for risk diversity and business continuity. When major blackouts or even congestions happen, multi-homing gives a fallback link for the “first-mile” connection, rather than providing a way to route around the Internet “middle-mile” issues.

Noction Intelligent Routing Platform (IRP) is a product developed by Noction to help businesses to optimize their multi-homed network infrastructure. The platform sits in the network and gets a copy of the traffic from the edge routers. The system passively analyzes it for specific TCP anomalies and actively probes remote destination networks for such metrics as latency, packet loss, throughput, and historical reliability. It computes a performance- or a cost-improvement network traffic engineering policy and applies the new improved route by announcing it to the network’s edge routers via traditional BGP session.

Noction IRP is a complete network monitoring and troubleshooting solution, which facilitates the detection, diagnosis, and automatic resolution performance issues. It delivers real-time views and dashboards that allow to visually track network performance and generate triggers, alerts and notifications when specific problems occur.

1.2.1 IRP Features

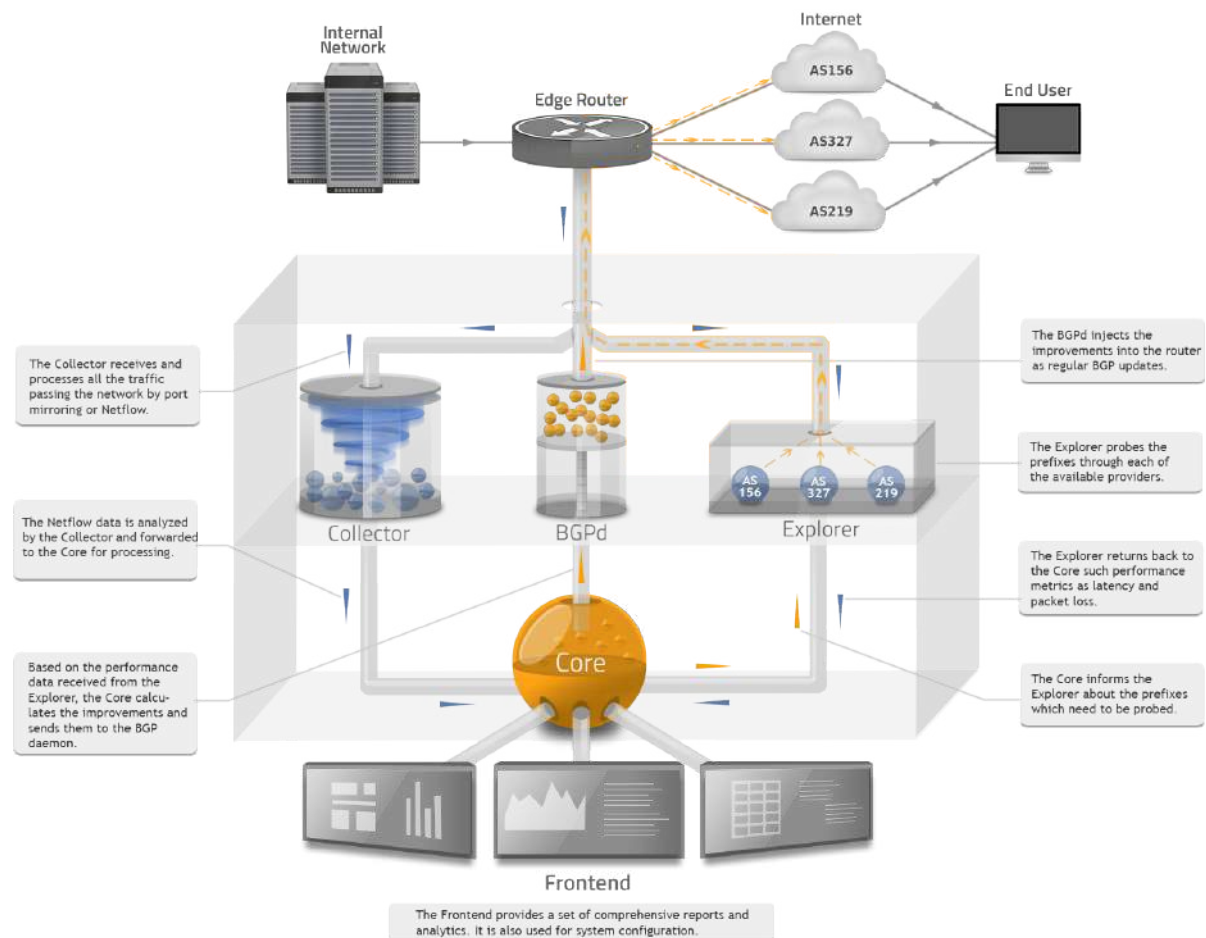
The Intelligent Routing Platform is designed to help Service Providers to improve the performance and to reduce the costs of running a multi-homed BGP network. The system makes intelligent routing decisions by analyzing various network performance metrics and selecting the best performing route for the traffic to pass through. As a result, Noction IRP allows you to:

- Improve overall network performance
- Reroute congestion and outages

- Decrease network downtime
- Reduce latency and packet loss
- Get comprehensive network performance analytics
- Facilitate network troubleshooting
- Decrease network operational costs
- Monitor platform performance
- Reduce the risk of human errors during BGP configuration

1.2.2 IRP Components

The IRP platform has a few interconnected components (see figure 1.2.1) , which are performing together to improve the routing best path selection and to eliminate various common network issues.



A short description of each of the components is given below. The detailed information is available in the next chapters.

To inject the Improvements into the network, the platform needs to ‘tell’ the routers what exactly needs to be changed in the routing table. IRP **BGP daemon** announces the improved prefix with an updated next-hop for the traffic to start flowing through the new path.

Bmpserver receives BMP monitoring sessions to provide route data to IRP components.

Core is the most important part of the system. It runs all the logical operations and interconnects all the components. It handles the performance and cost improvements. It processes, stores and exports data.

Collector **Irpflowd** & **Irpspand** are receiving, analyzing and processing all the traffic passing the network. They have two ways to gather data about the network: by mirroring network traffic or by using NetFlow/sFlow. Collector also gathers interface statistics from the edge routers via SNMP protocol.

Explorer runs all the probes and checks the metrics specified by the platform policies, such as packet loss and latency. This information is sent back to the Core.

Frontend represents a web interface with a comprehensive set of reports, graphs and diagnostic information which can reflect the current and historical network state, as well as the benefits of a single or multiple Intelligent Routing Platforms in terms of network optimization.

Globalcc performs communication between different IRP instances to maintain commit levels globally.

Irppapid is used to serve requests from Frontend (the Global Management Interface) and provides API to external tools.

Irpdetected collects various statistics for DDOS detection and mitigation feature.

Irpiperfd performs statistical analysis for Inbound Performance.


Irpmng is a command-line interface designed to perform various management tasks.

Irppushd is responsible for forwarding events to configured notification channels.

1.2.3 IRP Technical Requirements

In order to plan the IRP deployment in your network, a series of requirements need to be met and specific information has to be determined to configure IRP.

1.2.3.1 Hardware requirements

 In production, a dedicated server for each IRP instance is strongly recommended. The system can also be deployed on a Virtual Machine with matching specifications, provided that this is hardware- or paravirtualization (Xen, KVM, VMware). Os-level virtualization (OpenVZ/Virtuozzo or similar) is not supported.

1. CPU

- Recommended Intel® Xeon® Processor E3/E5 family, for example:
 - 1x Intel® Xeon® Processor E3 family for up to 20 Gbps traffic;
 - 1x Intel® Xeon® Processor E5 family for 40 Gbps or more traffic.

2. RAM

- if providing sFlow/NetFlow data at least 16 GB, recommended - 32 GB;
- if providing raw traffic data by port mirroring:
 - minimum 16 GB for up to 10 Gbps traffic;
 - minimum 32 GB for 40 Gbps traffic
- Additional RAM would be required to maintain large amount of BGP & BMP sessions (for example, Bgpd occupies about 10G of RAM for 16 full view BGP sessions; the estimation may change due to growth of world's BGP table and new IRP features).


3. HDD

- At least 160GB of storage;
- SAS disks are recommended (SSDs are required only for 40Gbps+ networks);
- HDD partitioning:
 - LVM is recommended;
 - At least 100GB disk space usable for /var or separate partition;


- At least 10GB disk space usable for /tmp or separate partition. This is required for big mysql tables manipulation. More disk space might be required under heavy workload.

4. NIC

- if providing sFlow/NetFlow data - at least 1 x 1000Mbps NIC while two NICs are recommended (one will be dedicated to management purposes).
- if providing raw traffic data by port mirroring - additional 10G interfaces are required for each of the configured SPAN ports (Myricom 10G network cards with Sniffer10G license are recommended to be used for high pps networks). When configuring multiple SPAN ports the same number of additional CPU cores are needed to analyze traffic.


 In very large networks carrying hundreds or more of Gbps of traffic and in IRP configurations with very aggressive optimization settings configurations with 2x or 4x CPUs are recommended. The IRP servers in these cases should also allocate double the recommended RAM and use SSD storage.

Noction can size, setup and mail to you an appliance conforming to your needs. The appliance is delivered with OS installed (latest CentOS 7) and IRP software deployed.

 A different supported OS can be installed on customer request.

1.2.3.2 Software requirements

Clean Linux system, with the latest CentOS 7, Red Hat Enterprise Linux 8 or Ubuntu Server LTS for x86_64 architecture installed on the server.

 IRP has a dependency on MySQL/MariaDB server and it expects the latest version from official OS repositories. In case the DBMS has been installed from a different repository it is strongly advised that the database instance and its configuration is purged before proceeding with IRP installation.

IRP requires root access to local database instance during first installation. In case the root access can't be given, use the statements below to grant all necessary privileges to the 'irp' user and database.:

```
GRANT ALL ON $dbdbname.* TO '$dbusername'@$dbourhost' IDENTIFIED BY '$dbpassword' WITH GRANT OPTION
```

```
GRANT SELECT ON $dbdbname_fe.* TO '$dbusername_fe'@$dbourhost' IDENTIFIED BY '$dbpassword_fe'
```

where \$dbdbname (5.1.7), \$dbusername (5.1.12), \$dbpassword (5.1.10), \$dbourhost (5.1.9) are the corresponding parameters from `/etc/noction/db.global.conf` \$dbdbname_fe and \$dbusername_fe from `/etc/noction/db.frontend.conf`

1.2.3.3 Network-related information and configuration

IRP is designed to help Service Providers (AS) optimize a multi-homed BGP network. This implies the basic prerequisites for using IRP:

- Ownership of the AS for the network where IRP is deployed,
- BGP protocol is used for routing and,
- Network is multi-homed.

Eventually the following needs to be performed in order to deploy and configure IRP:

1. Prepare a network diagram with all the horizontal (own) as well as upstream (providers) and downstream (customers) routers included. Compare if your network topology is logically similar to one or more of the samples listed in section [Collector Configuration](#) for example [Flow export configuration](#).
2. Identify the list of prefixes announced by your AS that must be analyzed and optimized by IRP.
3. Review the output of commands below (or similar) from all Edge Routers:

- `'sh ip bgp summary'`
- `'sh ip bgp neighbor [neighbor-address] received-routes'`
- `'sh ru'` (or similar)

The settings relating to BGP configuration, prefixes announced by your ASN, the route maps, routing policies, access control list, sFlow/NetFlow and related interfaces configurations are used to setup similar IRP settings or to determine what settings do not conflict with existing network policies.

4. Provide traffic data by:
 - (a) **sFlow**, **NetFlow** (v1, 5, 9) or **jFlow** and send it to the main server IP. Make sure the IRP server gets both inbound and outbound traffic info.
Egress flow accounting should be enabled on the provider links, or, if this is not technically possible, ingress flow accounting should be enabled on all the interfaces facing the internal network.
NetFlow is most suitable for high traffic volumes, or in the case of a sophisticated network infrastructure, where port mirroring is not technically possible.
Recommended sampling rates:
 - i. For traffic up to 1Gbps: 1024
 - ii. For traffic up to 10Gbps: 2048
 - (b) Or: configure port mirroring (a partial traffic copy will suffice). In this case, additional network interfaces on the server will be required - one for each mirrored port.

See also: [Collector Configuration](#)


5. Setup Policy Based Routing (**PBR**) for IRP active probing.
 - Apart from the main server IP, please add an additional alias IP for each provider and configure PBR for traffic originating from each of these IPs to be routed over different providers.
 - No route maps should be enforced for the main server IP, traffic originating from it should pass the routers using the default routes.
 - Define Provider↔PBR IP routing map

In specific complex scenarios, traffic from the IRP server should pass multiple routers before getting to the provider. If a separate probing Vlan cannot be configured across all routers, GRE tunnels from IRP to the Edge routers should be configured. The tunnels are mainly used to prevent additional overhead from route maps configured on the whole IRP↔Edge routers path.

 If network has Flowspec capabilities then alternatively Flowspec policies can be used instead of PBR. Refer for example [Flowspec policies](#), [global.flowspec.pbr](#).

1. Configure and provide **SNMP** for each provider link, and provide the following information:
 - SNMP interface name (or ifIndex)
 - SNMP IP (usually the router IP)
 - SNMP community

This information is required for the report generation, Commit Control decision-making and prevention of overloading a specific provider with an excessive number of improvements.

 The above is applicable in case of SNMP v2c. If SNMP v3 is used further details will be required depending on security services used.

1. To setup cost related settings as well as the Commit Control mechanism, provide the maximum allowed interface throughput for each provider link as well as the cost per Mbps for each provider.

1.2.4 IRP Operating modes

The IRP platform can operate in two modes, which can be used at different stages of the deployment process. During the initial installation and configuration, it is recommended for the system not to inject any improvements into the network until the configuration is completed.

After running several route propagation tests, the system can be switched to the full Intrusive mode.

1.2.4.1 Non-intrusive mode

While running in this mode, the system will not actually advertise any improvement to the network, and will only reflect the network improvements and events in the platform reports and graphs.

1.2.4.2 Intrusive mode

After the system configuration is completed, and manual route propagation tests were performed in order to ensure that the edge routers behavior is correct, the system can be switched to Intrusive mode. While running in this mode, the system injects all the computed improvements into the edge router(s) routing tables, allowing the traffic to flow through the best performing route.

1.2.4.3 Going Intrusive

While IRP operates in non-intrusive mode it highlights the potential improvements within client's environment. Going Intrusive will realize IRP potential.

The difference between Intrusive and Non-Intrusive operating modes is that IRP advertises improvements to edge routers. In order to switch to Intrusive we follow a controlled process. The highlights of the process are as follows:

1. The optimizing component of IRP (Core) is taken offline and existing improvements are purged. The Core being offline guarantees IRP will not automatically insert new improvements into its Current Improvement table and hinder the Go Intrusive process.

Listing 1.1: Stop IRP Core and purge existing improvements

```
root@server ~ $ service core stop
root@server ~ $ mysql irp -e 'delete_from_improvements;'
```

2. Enable Intrusive Mode and adjust IRP Core and Bgpd parameters as follows:

Listing 1.2: Switch to Intrusive Mode and adjust IRP Core and Bgpd parameters

```
root@server ~ $ nano /etc/noction/irp.conf
global.nonintrusive_bgp = 0
core.improvements.max = 100
bgpd.improvements.remove.next_hop_eq = 0
bgpd.improvements.remove.withdraw = 0
```

3. Improvements towards test networks are introduced manually so that client's traffic is not affected. The improvements are chosen so that they cover all client's providers. Any public networks could be used for test purposes. Just keep in mind that preferably your network shouldn't have traffic with chosen test network in order to do not re-route the real traffic. Use the template below in order to insert the test improvements:

Listing 1.3: Inserting test improvements

```
mysql> insert into improvements
(ni_bgp, prefix, peer_new, ipv6, asn)
```

```

values
(0, '10.10.10.0/24', 1, 0, 48232),
(0, '10.10.11.0/24', 2, 0, 48232),
(0, '10.10.12.0/24', 3, 0, 48232);

```

4. Make sure that 'route-reflector-client' is set for IRP BGP session.
5. Make sure that 'next-hop-self' is not configured for IRP BGP session.
6. On iBGP sessions (between edge routers, route-reflectors; except session with IRP) where 'next-hop-self' is configured, the following route-map should be applied:


Listing 1.4: Remove next-hop-self route-map (RM-NHS) example

```

route-map RM-NHS
set ip next-hop peer-address
neighbor X.X.X.X route-map out RM-NHS

```

where X.X.X.X is the iBGP neighbor

 route-map contents should be integrated into existing route-map in case other route-map already configured on the iBGP session.

7. Use the commands below to restart IRP Bgpd to use actual IRP configuration and establish BGP session(s) and verify if BGP updates are being announced:

Listing 1.5: Restart IRP Bgpd

```

root@server ~ $ service bgpd restart
root@server ~ $ tail -f /var/log/irp/bgpd.conf

```

Wait **for** the following lines **for** each BGP session:

```

NOTICE: Adding peer X
NOTICE: BGP session established X
INFO: N update(s) were sent to peer X

```

where X is the router name and N is the number of the updates sent towards the X router.

8. Verify if IRP BGP announcements are properly propagated across all the network. Run the following commands on each router (the commands vary depends on the router brand):

Listing 1.6: Show BGP information for specified IP address or prefix

```

show ip bgp 10.10.10.1
show ip bgp 10.10.11.1
show ip bgp 10.10.12.1

```

Analyze the output from all the routers. If the IRP BGP announcements are properly propagated, you should see /25 (refer to 5.4.35) announcements and the next-hop for each announcement should be the improved provider's next-hop:

```

10.10.10.1 - provider 1 next-hop
10.10.11.1 - provider 2 next-hop
10.10.12.1 - provider 3 next-hop
(refer to 5.14.31, 5.14.39, 5.17.3).

```

Run the following commands in order to check if IRP improvements are announced and applied:

Listing 1.7: Traceroute destination networks

```
root@server ~ $ traceroute -nn 10.10.10.1
root@server ~ $ traceroute -nn 10.10.11.1
root@server ~ $ traceroute -nn 10.10.12.1
```

Again, you should see corresponding providers' next-hops in the traces.

9. If the tests are successful perform the steps below:

(a) Delete test improvements

Listing 1.8: Delete test improvements

```
root@server ~ $ mysql -e "delete_from_improvements_where_prefix_
like_'10.10.1%';"
```

(b) Configure at most 100 improvements and revert Bgpd configuration

Listing 1.9: Configure the maximum improvements limit and revert Bgpd configuration

```
root@server ~ $ nano /etc/noction/irp.conf
core.improvements.max = 100
bgpd.improvements.remove.next_hop_eq = 1
bgpd.improvements.remove.withdraw = 1
```

(c) Restart IRP Core and Bgpd

Listing 1.10: Restart IRP Core and Bgpd

```
root@server ~ $ service core restart
root@server ~ $ service bgpd restart
```

10. If everything goes well, after 1-2 hours the maximum number of improvements announced is increased to 1000 and after 24 hour to 10000.

As a rollback plan we have to revert the changes and switch the system to non-intrusive mode:

1. Delete test improvements

Listing 1.11: Delete test improvements

```
root@server ~ $ mysql -e "delete_from_improvements_where_prefix_like_
'10.10.1%';"
```

2. Switch the system to non-intrusive mode

Listing 1.12: Switch the system to non-intrusive mode

```
root@server ~ $ nano /etc/noction/irp.conf
global.nonintrusive_bgp = 1
```

3. Restart IRP Core and Bgpd

Listing 1.13: Restart IRP Core and Bgpd

```
root@server ~ $ service core restart
root@server ~ $ service bgpd restart
```

1.2.5 BGP Monitoring

IRP uses two types of BGP monitors and a BMP monitoring station to collect data, diagnose and report mainly the state of the BGP session between the edge routers and the providers, as well as the network reachability through a specific provider. The information provided by monitors enables IRP to avoid announcing routing updates that would result in traffic misrouting for example by sending improvements to a failed provider but also to better inform IRP probing and improvement decisions.

1.2.5.1 Internal monitor

Internal BGP Monitor is checking the state of the Edge Router→Provider BGP session by regularly polling the router via SNMP. When queried, the SNMP protocol returns variables describing the session status to be used by the IRP's Internal BGP Monitor. If the session between the edge router and the provider is down, SNMP will return a value, representing session failure and IRP will react as follows:

- the provider will be marked as FAILED,
- all the improvements towards this provider will be withdrawn from the routing tables to avoid creating black holes,
- new improvements towards this providers will not be made.

In some cases (e.g. DDoS attack or various factors causing router CPU over-usage) there may be no response to the SNMP queries at all. In this case a timeout status will be reported to the Internal Monitor and a 30 minutes timer (called longhold timer) ([bgpd.mon.longholdtime](#)) will be started. During this time the monitor will be sending ICMP/UDP ping requests toward the configured provider's next-hop IP address ([peer.X.ipv4.next_hop](#) or [peer.X.ipv6.next_hop](#)). The requests will be sent once in keepalive period (a parameter adjustable in the BGP daemon configuration interface) ([bgpd.mon.keepalive](#)). If the next-hop stops responding to these requests, another 30 seconds timer (called hold timer) ([bgpd.mon.holdtime](#)) will be started. If according to the ping response the session is reestablished during this time, the hold timer will be discarded while the longhold timer continues. In case one of the timers expires, the provider is switched to a FAIL state and all the improvements towards this provider will be withdrawn from the routing table. However, if the BGP session with the provider is re-established, the system will start rerouting traffic to this provider.

When the Bgpd is started, the monitors are initialized and one SNMP query is sent towards each router, in order to check the status of the BGP sessions with providers. If there is no reply, the Internal Monitor will send up to two more SNMP requests, separated by a keepalive interval.

i Internal monitors for Internet Exchange peering partners are not initialized until there is an improvement made towards it. When running in non-intrusive mode internal monitors for IX peers are not initialized at all.

If none of the SNMP queries returned a status confirming that the sessions with providers are up, the provider will be assigned a FAIL status and the Internal Monitor will continue the periodical SNMP polling (each 60 seconds), to recheck providers sessions' status.

Then, the BGP session with the edge routers is initialized and Bgpd starts retrieving the routing table from the edge routers. While IRP retrieves the routing table, SNMP request may timeout due to the high CPU usage on the edge routers.

For details refer:

[bgpd.mon.guardtime](#)

[bgpd.mon.keepalive](#)

[bgpd.mon.holdtime](#)

[bgpd.mon.longholdtime](#)

1.2.5.2 External monitor

External BGP Monitor analyzes the network reachability through a specific provider. It performs ICMP/UDP ping requests towards the configured remote IP address(es) ([peer.X.ipv4.mon](#) or [peer.X.ipv6.mon](#)) through the monitored provider. If any of the configured IP addresses are accessible, the monitor is marked as OK. If the monitored remote IP addresses do not reply through the examined provider IRP will react as follows:

- the provider will be marked as FAILED,
- all the improvements towards this provider will be withdrawn from the routing table,
- new improvements towards this providers will not be made.

If for some reason (e.g. when the provider's interface goes down state), the Next-Hop of the Policy Based Routing rule does not exist in the routing table, then the packets forwarding may return to the default route. In that case, the External BGP Monitor will return a false-positive state. To avoid that by properly configuring PBR, please consult "Specific PBR configuration scenarios" ([Specific PBR configuration scenarios](#)).

The External BGP Monitor status does not depend on the state of the BGP session(s) between the edge router and the provider (which is monitored by the Internal BGP Monitor). Therefore, in the case that the BGP session with one of the providers goes down, the External Monitor still shows an OK state which will remain unchanged as long as the packets are successfully routed towards the monitored destination. We do recommend adding at least two remote IP addresses, in order to prevent false-positive alerts.

When both BGP monitors are enabled ([peer.X.mon.ipv4.internal.state](#), [peer.X.mon.ipv4.external.state](#)), they function in conjunction with each other. If any of them fails, the provider will be declared as FAILED and IRP will react as described above. The BGP monitors' statuses are displayed on the system dashboard as shown in the screenshot below.

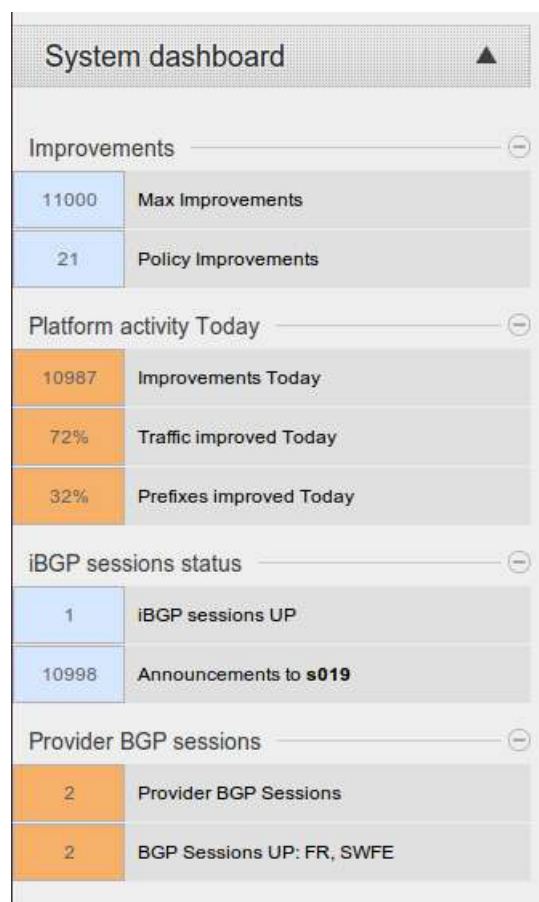



Figure 1.2.2: System Dashboard

 Starting with version 1.8.5, IRP requires at least the Internal Monitor to be configured. Otherwise, the system Frontend will return an error as shown below.

Attention! The Internal BGP Monitor is not configured. This may cause traffic misrouting in case of provider failures!



Figure 1.2.3: Error: Internal BGP Monitor not configured

For details refer:

peer.X.mon.snmp
 peer.X.ipv4.mon
 peer.X.ipv6.mon
 peer.X.mon.ipv4.bgp_peer
 peer.X.mon.ipv6.bgp_peer

1.2.5.3 BMP monitoring station

A BMP monitoring station is included in IRP starting with version 3.9. It implements the monitoring station specified in RFC 7854 BGP Monitoring Protocol (BMP). The BMP monitoring station requires a monitored router to communicate over BMP the detailed routing information received from neighbors. The BMP monitoring station exposes detailed routing data to other IRP components so that better and timelier decisions are made, for example:

- BMP lists both active and inactive routes advertised by peers on an Internet Exchange. The additional information is used by IRP to evaluate and identify the best candidate peers at all times. Without BMP data IRP has knowledge about active routes only which only point to a single peer on the IX while all the alternatives are hidden.
- route changes even for inactive routes are visible via BMP. This allows IRP the opportunity to revisit previously made probes and improvements not only at predefined re-probing intervals but also when route changes are detected for both active and inactive routes.
- prefix monitors for IX improvements consume significant router CPU resources in order to service the SNMP requests traversing the router's relevant OIDs. More so this information is at times inaccurate and vendor dependent. When BMP data is available IRP uses this routing data to determine if IX peers still advertise the routes and no longer makes the SNMP requests for those prefixes thus significantly reducing the CPU overhead especially on routers servicing very large IX.
- IRP reconstructs the AS Path for candidate providers in order to make accurate iBGP announcements of improvements. Unfortunately network configuration practices might cause some errors during reconstruction of AS Paths using traceroute. BMP data makes the reconstruction of AS Path redundant and more accurate as this BGP attribute can be retrieved from actual (inactive) routes received from neighbors.
- improvements can be re-visited on AS Path changes. Both new and old provider AS Path attributes are monitored via BMP for changes. When changes are detected IRP re-probes the prefix to ensure the network uses the best available route. Note that re-probing can be triggered on any AS Path changes or only on major ones - when AS Path traverses a different set of autonomous systems.
- Accurate as_path reconstruction helps Inbound performance feature detecting path changes and correlate it with performance changes.

The possible benefits of passing BMP data to IRP are many. To benefit from them the monitored router must support BMP too. Configuration is fully performed on monitored router by pointing it to the IRP BMP monitoring station IP address and port. The monitored router establishes the TCP connection and communicates the data while the IRP BMP monitoring station continuously listens and accepts fresh routing data as it comes.

i As per BMP RFC requirements IRP BMP monitoring station never attempts to establish BMP or any other connections with the monitored router leaving the full scope of decisions regarding when and if BMP data is communicated in network's responsibility.

w Any route filtering applied to a BGP session with Partial Routing provider or IX peering partner wouldn't not be taken into consideration by a BMP sender. If filtering is deemed as important, then IRP shouldn't use BMP data to find routes (`peer.X.bmp.check_routes`).

1.2.6 Outage detection

A complete traffic path from source to destination typically passes through multiple networks, with different AS numbers. This is reflected in the traceroute results. If the Outage detection is enabled in the IRP configuration, the system gathers network performance information for all traceroute hops over which the traffic passes to the remote networks. Next, each hop is translated into AS numbers. In case any network anomalies are detected on a specific ASN, then this ASN and the immediate neighbor ASN are declared a problematic AS-pattern. The system then re-probes the prefixes that pass through this AS-pattern. In case the issue is confirmed, all related prefixes are rerouted to the best performing alternate provider.

The Outage detection uses a statistical algorithm for selecting the best routing path. Rerouting will occur for all the prefixes that are routed through the affected as-pattern, despite their current route .

i Several improvements-related reports need to indicate the original route for a specific prefix. This value is taken from the last probing results for this prefix, even if the results are outdated (but not older than 24h). Since the outage-affected prefixes are rerouted in bulk by as-pattern, in some cases the reports can show the same provider for both the old and the new route.

1.2.7 VIP Improvements

The VIP Improvements is a feature that allows manual specification of a list of prefixes or AS numbers that will be periodically probed by IRP and optimized in compliance with the probing results. This allows the system to monitor specific networks or Autonomous Systems, without reference to the data provided by the IRP collector.

Possible usage scenarios include, but are not limited to:

- monitoring and optimizing traffic to commercial partners that should have access to your networks via the best performing routes
- monitoring and optimizing traffic to your remote locations, operating as separate networks
- monitoring and optimizing traffic to AS, which are known for the frequent accessibility issues due to geographical or technical reasons

i If a prefix is being announced from multiple Autonomous Systems, you can see different ASNs in [Current improvements](#) report in the prefixes translated from ASN

IRP performs the proactive (more frequent than regular probing) monitoring of the VIP prefixes/ASNs that allows VIPs to be constantly improved.

For future reference, see:


[core.vip.interval.probe](#)

1.2.8 Retry Probing

Retry Probing is a feature that allows reconfirmation of the initial and already reconfirmed improvements validity. The feature is applicable to all types of improvements made by the system (Performance, Cost and Commit Control improvements). The improvements that were made more than a retry probing period ago (`core.improvements.ttl.retry_probe`) are being sent to retry probing. If the probing results confirm that the current improvement is still valid, it stays in the system and its description is updated. Otherwise, it will be removed with the log message further described in this section.

During Retry Probing reconfirmation the improvement details will be updated in the following cases:

- Performance and Cost improvements
 - An old provider has been removed from the system configuration.
Example: “Old provider and performance metrics not known. New packet loss 55%, avg rtt 105 ms.”
- Commit Control improvements
 - An old provider’s has been removed from the system configuration.
Example: “Previous provider not known. Rerouted 1 Mbps to Peer5[5] (250 Mbps, 50%)”
 - An old provider’s bandwidth statistics are not available.
Example: “Rerouted 6 Mbps from Peer1[1] to Peer5[5] (250 Mbps, 50%)”
 - A new provider’s bandwidth statistics are not available.
Example: “Rerouted 6 Mbps from Peer1[1] (250 Mbps, 50%) to Peer5[5]”
 - The old and new providers’ bandwidth statistics are not available.
Example: “Rerouted 6 Mbps from Peer1[1] to Peer5[5]”

 Commit control improvements are reconfirmed based on their average bandwidth usage (and not on current bandwidth usage). This way if performance characteristics allow it, even when current bandwidth usage is low but the average is still relevant, the improvement is preserved thus anticipating network usage cycles and reducing number of route changes.

During Retry Probing reconfirmation the improvements will be removed from the system and the details will be logged into the log file (`core.log`) in the following cases:

- The Commit Control feature has been disabled.
Example: “Prefix 1.0.2.0/24 withdrawn from Improvements (Commit Control is disabled)”
- The low prefix traffic volume is less than the configured bandwidth limits (`core.commit_control.agg_bw_min`).
Example: “Prefix 1.0.2.0/24 withdrawn from Improvements (low traffic volume, irrelevant for the Commit Control algorithm)”
- The system has been switched from the Cost mode to the Performance mode (applied for cost improvements only).
Example: “Prefix 1.0.2.0/24 withdrawn from Improvements (Performance Improvements mode)”
- A prefix has been added to the ignored networks/ASN list.
Example: “Prefix 1.0.2.0/24 withdrawn from Improvements (added to ignored networks/ASN)”
- The improvement’s performance metrics are not the best ones anymore.
Example: “Prefix 1.0.2.0/24 withdrawn from Improvements (Performance Improvement not actual anymore)”
- The maximum number of improvements limits (`core.improvements.max`, `core.improvements.max_ipv6`) are exceeded
Example: “Prefix 1.0.2.0/24 withdrawn from Improvements (no more available Improvement slots)”


For future references, see:


`core.eventqueue.limit.retry_probe_pct`
`core.improvements.ttl.retry_probe`

1.2.9 Routing Policies


Routing Policies brings the capability of defining specific routing policies according to business objectives. This feature permits denying or allowing providers to be used for probing and reaching a specific prefix or ASN. It also provides the possibility to set a static or static exact route through a particular provider. Static route policy will only apply improvements for prefixes learned from the BGP table or external sources. Static exact route policy will announce improvements according to the exact policy details provided by the user, regardless of prefixes being present in the BGP table or external data sources (BGP split does not apply to static exact policies).

Within an Allow or Deny policy, you can choose between VIP or non-VIP (regular) probing mechanisms to be used.

 Policies can be configured for networks (ASN) or aggregate prefixes. IRP works with prefixes from the global BGP routing table. Each prefix is checked against a set of policies, matching a single policy in accordance with the policy priority value, those with the highest value being checked first.

 In version 3.9 IRP added support for policies by Country. Note that IRP maps individual prefixes to a country and does not use a transitive inference based on AS records. These prefix mappings allow IRP to more accurately work with large transcontinental AS.

Starting with version 3.9 IRP introduces a priority attribute that defines what policy to choose in cases when different policies include the same unpacked specific prefix. The policy with the highest priority will apply for such a prefix.


 Note that after upgrade all existing policies are assigned (implicitly) the lowest default priority of 0.


Below you can find some typical scenarios of Routing Policies usage. For instance, there may be a specific prefix that consumes a high volume of traffic and IRP redirects it through the most expensive provider due to performance considerations. At the same time you have another two less costly available providers. In this case, you could deny the expensive provider for the specific prefix and IRP will choose the best-performing provider between the remaining two. This can be achieved by applying a Deny policy to the expensive provider or an Allow policy to the less costly providers.

The table below shows which cases a specific policy can be applied in, depending on the number of providers.

No. of providers \ Policy	Allow	Deny	Static	Static Exact
1 provider	No	Yes	Yes	Yes
2 providers or more (maximum: total number of providers - 1)	Yes	Yes	No	No
All providers (VIP probing disabled)	No	No	No	No
All providers (VIP probing enabled)	Yes	No	No	No

If a Static or Static Exact Route policy is applied to a prefix, VIP probing through each of the providers is unnecessary. Regular probing will suffice for detection of a provider failure that would trigger IRP to reroute the traffic to a different provider. Therefore IRP does not allow using VIP probing within a Static or Static Exact Route policy.

 Routing policies are designed to control outgoing paths for destination networks that are outside your infrastructure. This feature should not be used to manipulate your infrastructure network behavior.

 Avoid setting up of policies that point to same prefix. When improvement by aggregate is enabled, multiple prefixes can point to the same aggregate and this can cause unpredictable behaviour.

If a provider is added or removed, suspended or shutdown, the routing policies are adjusted by IRP in the following way:

A new provider is added.

Policy	Result
Allow all providers (VIP probing enabled)	The new provider is automatically included into the configured policy and probed by the VIP probing mechanism.
Allow selected providers only	The new provider is automatically ignored by the configured policy and not probed by the probing mechanism.
Deny	The new provider is automatically included into the configured policy and probed by the selected probing mechanism.
Static Route	The new provider is automatically ignored by the configured policy.
Static Exact	The new provider is automatically ignored by the configured policy.

The provider under the policy is removed.

Policy	Result
Allow all providers (VIP probing enabled)	The new provider is automatically removed from the configured policy and not probed by the VIP probing mechanism. If there is only one provider left, the policy is automatically deactivated.
Allow selected providers only	The provider is automatically removed from the configured policy and not probed by the probing mechanism. If there is only one provider left, the policy is automatically deactivated.
Deny	The provider is automatically removed from the configured policy and not probed by the probing mechanism. If there is no any other provider under this policy, it is automatically deactivated.
Static Route	The policy is automatically deactivated.
Static Exact Route	The policy is automatically deactivated.


The provider under the policy is suspended.

Policy	Result
Allow all providers (VIP probing enabled)	The provider is temporarily removed from the configured policy and not probed by the VIP probing mechanism. If there is only one provider left, the policy is temporarily deactivated.
Allow selected providers only	The provider is temporarily removed from the configured policy and not probed by the probing mechanism. If there is only one provider left, the policy is temporarily deactivated.
Deny	The provider is temporarily removed from the configured policy and not probed by the probing mechanism. If there is no any other provider under this policy, it is temporarily deactivated.
Static Route	The policy is temporarily deactivated.
Static Exact Route	The policy is temporarily deactivated.

The provider under the policy is shutdown.

Policy	Result
Allow all providers (VIP probing enabled)	The provider is temporarily removed from the configured policy and not probed by the VIP probing mechanism. If there is only one provider left, the policy is temporarily deactivated.
Allow selected providers only	The provider is temporarily removed from the configured policy and not probed by the probing mechanism. If there is only one provider left, the policy is temporarily deactivated.
Deny	The provider is temporarily removed from the configured policy and not probed by the probing mechanism. If there is no any other provider under this policy, it is temporarily deactivated.
Static Route	The policy is temporarily deactivated.
Static Exact Route	The policy is temporarily deactivated.


Policies that target an AS can be cascaded. Cascading applies the same policy to AS that are downstream to target AS, i.e. to AS that are transited by target AS.


 The use case of cascading is to apply a policy to a remote AS that transits a few other AS. Still, a cascading policy can cover a huge number of down-streams. This number is parameterized and can be set to values that best fit customer's needs. Refer for example [5.4.19](#).

When multiple routing domains are configured a policy can be configured to prevent global improvements. Refer to [Optimization for multiple Routing Domains](#) for further details about routing domains.

Policies can be assigned a specific community and all policy based improvements will be marked with the designated value to allow their further manipulation on edge routers.

All Routing Policies are stored in `/etc/noction/policies.conf` file, which is automatically generated by the system.

 Do not alter the `/etc/noction/policies.conf` file manually because any modifications will be overwritten by the system. Any changes can only be performed from the Configuration -> Routing Policies section in the system Frontend.

 The rule will be ignored by the platform if the Routing Policy contains a syntax/logical error.

Please check ([Routing Policy](#)) for a detailed description of Routing Policies parameters.

1.2.10 Support for Centralized Route Reflectors




Figure 1.2.4: Support for Centralized Route Reflectors

IRP gives the possibility to advertise routes into one or more route reflectors which subsequently advertise improvements into upper and/or lower routing layers (such as edge, core or distribution).

In case the iBGP sessions can't be established between the IRP appliance and edge routers, a route reflector is used. The following restrictions apply for such a solution:

- The Next-Hop-Self option should not be used. A direct iBGP session is required between IRP and each of the divergence points (Reflector1, Edge2) in case it is enabled. This restriction is not applicable between reflectors and core/distribution layers.
- Next-Hop addresses should be accessible (exist in the routing table) where next-hop-self is not applied (Either static routes or IGP is used).
- An Internal Monitor should be configured to retrieve the eBGP session state from the device where the corresponding eBGP session is terminated. For example, ISP1 should be monitored on Reflector1, ISP2 on Edge1, and ISP3 and ISP4 on Edge2.
- Injecting routes to reflector(s) can cause temporary routing loops.

In order to announce improvements into route reflector, it should be configured as a BGP router in the “Configuration”→“BGP and routers” section and should be assigned to all the related providers in the “Configuration”→“Providers and Peers” section.

1.2.11 Support for Internet Exchanges

A transit provider can deliver traffic to any destination on the Internet. However, within an Internet Exchange, a peering partner gives access only to the set of prefixes originated or transiting its network. Therefore, when IRP evaluates the Exchange as a best path, it has to know the prefixes announced by each peer, to avoid inefficient probing of paths that cannot lead to the desired destination.

With this purpose, IRP gets the routing table from the edge router containing the list of IPs and the corresponding next-hop; this represents the next router's IP address to which a packet is sent as it traverses a network on its journey to the final destination. IRP matches the prefix with the corresponding next-hop among the configured peers, allowing it to select for probing only those peers that have access to a specific prefix. This process is also performed in the case of a transit provider that gives access only to a limited set of prefixes, rather than the entire Internet.




Figure 1.2.5: IRP configuration in a multi-homed network connected to transit providers as well as and Internet Exchange

In the case of multiple transit providers, there is an additional IP alias added on the IRP platform for each provider. The edge router is configured in such a way that traffic originating from each of these IPs is routed over different providers. This is done with the help of Policy Based Routing (PBR) or Flowspec policies.

With PBR, a network engineer has the ability to dictate the routing behavior based on a number of different criteria other than the destination network. These PBR rules are applied to make sure that IRP probes are following the desired paths. However, when it comes to Internet Exchanges, configuring hundreds of IP aliases on the platform would result in inefficient IP address usage and an unmanageable setup.

To avoid this, a set of PBR rules are applied making sure that the probes to be sent through a specific provider are originating from one of the configured IPs with a specific DSCP code assigned. DSCP - Differentiated Services Code Point - is a field in an IP packet that enables different levels of service to be assigned to network traffic. Since DSCP can take up to 64 different values, one configured IP can be associated with up to 64 peers. Although, due to this mechanism, the number of required IP addresses for aliases to be configured has decreased considerably, hard work would still be needed to configure the PBR on the edge router as described above.

To solve this, IRP implemented a built-in PBR config-generator which provides the configuration code to be used for a specific router model. By running this generated set of commands, network administrators can easily configure the required PBR rules on the router.

1.2.12 Optimization for multiple Routing Domains

Overview

Some networks have multiple Points of Presence interconnected both internally via inter-datacenter links and externally via multiple transit providers. The diagram below depicts an example diagram with the available routes to one destination on the Internet.

IRP uses the concept of Routing Domains to separate the locations. A Routing Domain's main characteristic is that its routing tables are mainly built on data received from its locally connected providers and the preferred routes are based on locally defined preferences.

The process of optimizing outbound network traffic in such a configuration is to mainly find better alternative routes locally (within a Routing Domain) and only reroute the traffic to other Routing Domains via inter-datacenter links when local routes are completely underperforming.

It must be noted that a multiple Routing Domain configuration works best if the Points of Presence are not too far away (ex. a network with POPs in San Francisco, Palo Alto and Danville is perfectly suitable under this scenario).




Figure 1.2.6: City wide network

POPs situated at larger distances, for example in Las Vegas and Salt Lake City are still supported by a single IRP instance running in San Francisco.




Figure 1.2.7: Regional network

Intercontinental links for POPs in Tokyo and Melbourne are way too far away from the IRP instance in San Francisco and in such a case multiple IRP instances are required.

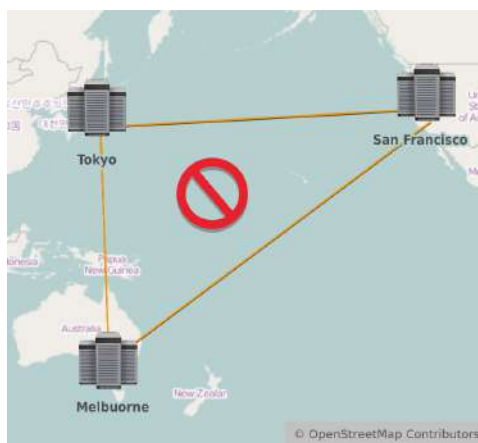


Figure 1.2.8: Intercontinental network

Multiple Routing Domains implementation attributes To further detail the multiple routing

domain attributes the following diagram will be used:

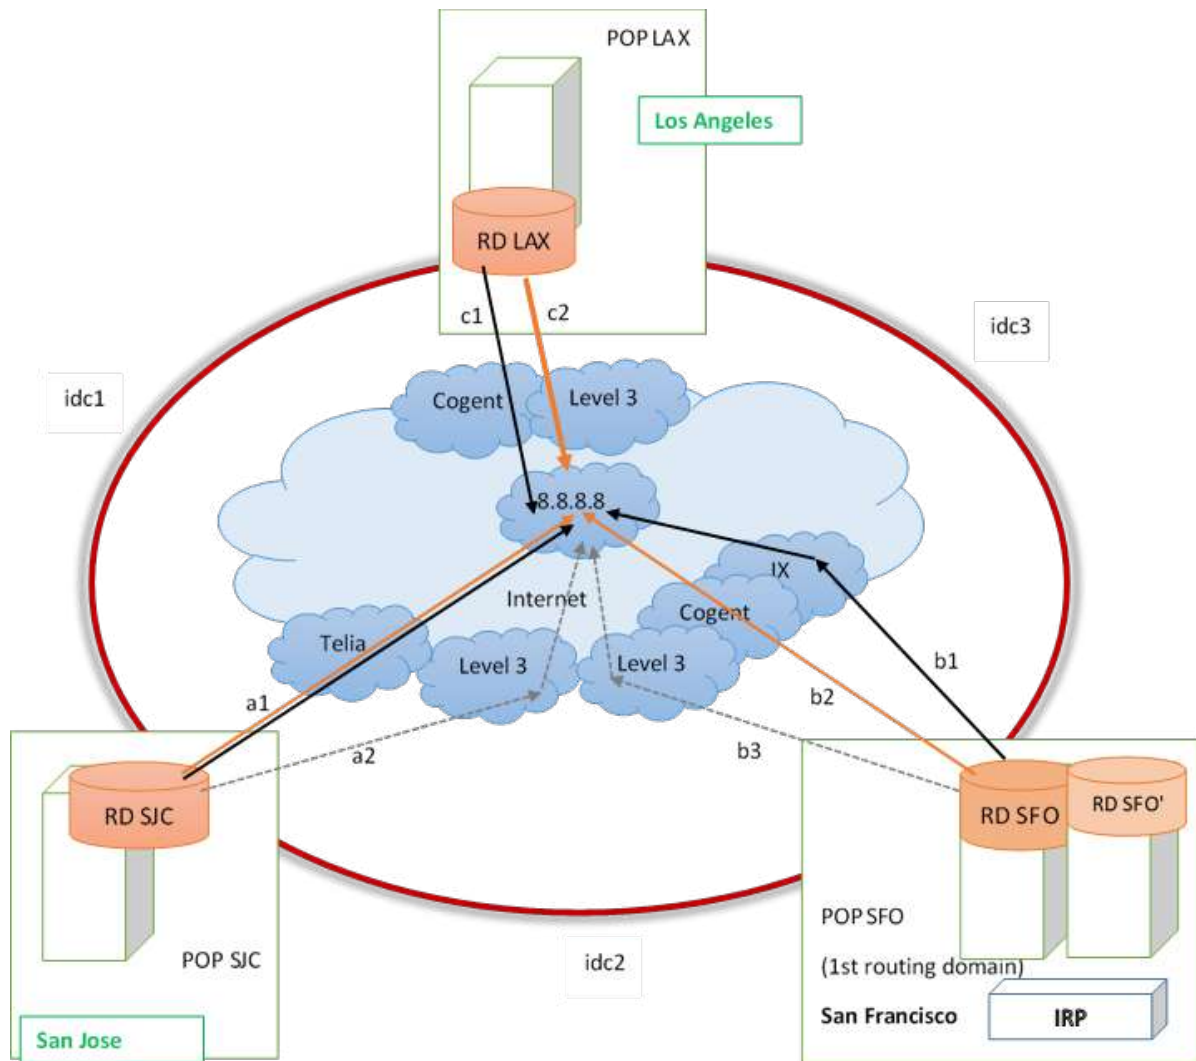


Figure 1.2.9: Multiple routing domains

A multiple Routing Domain configuration has a series of attributes:


- Multiple locations belonging to the same network (AS) represented in the diagram by POP SJC, POP SFO and POP LAX (of course, more than 3 routing domains are supported).
- The locations are distinguished by the different Routing Domains within which they operate (depicted by RD SJC, RD SFO, and RD LAX)
- The Routing Domains are managed by edge routers belonging to different locations
- Nearby locations that process routing data differently should be split into different Routing Domains, even if they have the same upstream providers. In the diagram above RD SFO and RD SFO' are depicted as part of a single Routing Domain. A decision to split or keep in the same routing domain should be made based on exact knowledge on how routing data is processed.
- Inter-datacenter loop interconnects the different locations (depicted by idc1, idc2 and idc3 segments)
- Data flows between locations take only the short path (in the example POP SJC can be reached from POP SFO via idc2 path (short) or idc3 + idc1 path (long))
- Each Routing Domain has different providers and different preferred routes to reach a specific destination (a1, b1, c1)

- A single IRP instance collects statistics about traffic (Irpflowd only), probes available destinations and makes improvements towards specific prefixes/networks on the Internet.
- IRP assumes RTT of zero and unlimited capacity to route traffic within a Routing Domain
- IRP assumes that Sites are not physically too far away. It is ok to have different sites in the same city or region as at this scale inter-datacenter links have predictable characteristics. When taking intercontinental links into consideration this is quite probably not the case.
- Distances between sites (idc1, idc2, idc3 delays) are measured in advance and specified in IRP's configuration.

Inter-datacenter link characteristics Support for Multiple Routing Domains relies on existence of inter-datacenter links. These links should be independent of upstream providers.

Example of inter-datacenter links that multiple routing domains is designed for are:

- private connections,
- L2 links with guaranteed service,
- MPLS links

 VPNs via public Internet could be used with Multi Routing Domain feature but is a suboptimal choice. Under such conditions IRP MUST be prevented from making Global Improvements. This way IRP will do only local optimizations in each Routing Domain and will operate similarly to multiple IRP instances (while probing excessively because it probes destinations via remote locations too).

Constraints At the moment IRP multiple Routing Domains implementation does not cover the following:

- IRP does not take measurements of inter-datacenter link delays (idc1, idc2 and idc3). This values are configurable.
- IRP does not monitor if inter-datacenter links are operating normally. In case such a link is broken it is expected IRP to loose BGP connectivity with routing domain routers and this will cause IRP improvements to be withdrawn till the link is restored.
- IRP does not try to detect if the traffic is following long or short paths on the inter-datacenter links. In the image above traffic from RD SJC can follow path idc1 (short) or idc2+idc3 (long). IRP always assumes the short path is being followed internally.
- IRP does not take measurements of inter-datacenter link capacity and current bandwidth usage. At this stage IRP assumes there is enough inter-datacenter link capacity to also carry the (few) global improvements. Also, IRP tries to minimize usage of inter-datacenter links.

Routing domains

Routing domain is a generic term used to distinguish a logical location that works with different routing tables. The differences are caused by the fact that a router composes its routing table according to routes received from different providers. It is possible to have multiple routing domains in the same datacenter if routing data is received by different routers (even from same or different sources) and data flows are distributed via different routers by different policies. In the image above RD SFO and RD SFO' can represent a single routing domain or multiple routing domains depending on what routing policies are applied.

Different routing domains are assigned identifiers in the range 1-100. Routing Domain identifier is assigned individually to each provider via parameter `peer.X.rd`. It must be noted the Routing domain that hosts the IRP instance is considered as the first routing domain (RD=1).

Parameter `global.rd_rtt` gives the distances between routing domains. The format of the parameter is

```
rda:rdb:rtt
```

for example if RD SJC has Routing Domain id = 42, RD SFO - 1 (since it hosts IRP), RD LAX - 3 then the idc1, idc2 and idc3 rtt is defined as the collection:

```
global.rd_rtt = 3:42:20 42:1:17 1:3:35
```

This parameter will be validated for correctness and besides the format above it requires that RD SJC and RD SFO values are different and already configured (RD1 is always present).

i Round trip time between one routing domain and another is calculated by executing PING towards edge routers and taking average integer value:

```
$ ping X -c 10 -q
PING X (X) 56(84) bytes of data.
--- X ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9085ms
rtt min/avg/max/mdev = 40.881/41.130/41.308/0.172 ms
```

Flow agents

A very natural constraint for Multiple Routing Domain networks is that IRP can rely only on Flow statistics - NetFlow or sFlow.

- SPAN cannot be used because it does not carry attributes to distinguish traffic between different providers

Flow collector needs to know the exact details of such a configuration in order to correctly determine the overall provider volume and active flows. For this each provider in an MRD setup must be assigned Flow agents to enable IRP to match Flow statistics accordingly. Refer [Flow agents](#) for further details.

Global and local improvements

Local improvements Local improvements represent better alternative routes identified within a routing domain. If in the example image above current routes are represented by black lines then local improvements are depicted by orange lines b2 and c2. Keep in mind that a1 just reconfirmed an existing current route and no improvements are made in such a case.

Local improvements are announced in their routing domains and this has the characteristic that local traffic exits customer's network via local providers. This also means that inter-datacenter interconnects are free from such traffic.

IRP prefers local routes and Improvements to Global improvements.

Parameter `bgpd.rd_local_mark` specifies a community marker that distinguishes local improvements from Global Improvements. A BGP speaker should not advertise these improvements outside its Routing Domain. It must be noted that a single marker is used for all the routing domains and each of them shall be configured to advertise local improvements within the domain and filter it out for inter-domain exchanges.

Local improvements should be stopped from propagating across routing domains. A route map is used to address this. Below are listed sample route maps for Cisco IOS and JUNOS 9.

Cisco IOS

- Refer your router capabilities in order to produce the correct route map. The route map **MUST** be integrated into existing route maps. It is not sufficient to simply append them.

```

neighbor <neighbor-from-another-RD> send-community (should be configured
    for all iBGP sessions)


ip community-list standard CL-IRP permit 65535:1
route-map RM-IRP-RD deny 10
    match community CL-IRP
route-map RM-IRP-RD permit 20

router bgp AS
    neighbor <neighbor-from-another-RD> route-map RM-IRP-RD out

```

Refer [Route-Maps for IP Routing Protocol Redistribution Configuration](#)

JUNOS 9

 Refer your router capabilities in order to produce the correct route map. The route map **MUST** be integrated into existing route maps. It is not sufficient to simply append them.

```

policy-options{
    policy-statement IRP-CL {
        term 0 {
            from {
                protocol bgp;
                community IRP-RD;
            }
            then reject;
        }
        term 1 {
            then accept;
        }
    }
    community IRP-RD members 65535:1;
}
protocols {
    bgp {
        group ebgp {
            type external;
            neighbor 10.0.0.1 {
                export IRP-CL;
            }
        }
    }
}

```

Refer [Policy Framework Configuration Guide; Release 9.3](#)

Global improvements Global improvements are made when IRP identifies an alternative route that even after factoring in the latencies incurred by inter-datacenter interconnects are better than all existing alternatives. Such an example can be represented by alternative route c2 in the image above. A global improvement is made when one routing domain alternative is better than the best local alternatives in all other routing domains even considering the latencies incurred by inter-datacenter interconnects. In the image above c2 will become a global improvement if his loss characteristic is best to all alternatives and its latency:

- (c2+idc1 - margin) is better than best local alternative a1 in RD SJC

- (c2+idc3 - margin) is better than best local alternative b2 in RD SFO

where:

- a1, b2 and c2 represent roundtrip times determined by IRP during probing of a destination.
- idc values are configurable and are set as one entry of `global.rd_rtt` parameter.
- margin is given by `core.global.worst_ms`.

i Global improvements can degrade performance in some routing domains. If performance for some routing domains degrades, IRP announcements for the global improvement also carry designated BGP community attributes set by `rd.X.community_worsening`.

Global improvements move traffic via inter-datacenter interconnects and as such are less desirable to local routes. Global improvements make sense when defined as above and even more sense when packet loss is taken in consideration and routing via a different datacenter reduces packet loss significantly.

1.2.13 Improvements weight

IRP assigns each improvement a weight. The weight takes into consideration many network and environment aspects of the change such as policy or VIP destinations, loss and latency differences, cost or commit control type of the improvement. Based on all of the above the improvement gathers more or less weight as appropriate.

Later on, instead of replacing oldest improvements that might still bring significant benefits with new improvements just because they are fresh, IRP relies on the weights to decide whether the benefit of the new improvement is sufficient to replace an existing one. More so, besides preserving the most relevant improvements this feature reduces route flapping by blocking announcement of new improvements and withdrawal of existing ones if the changes are not offering a good enough return.

1.2.14 Notifications and events

IRP produces a huge number of various events and some of them are critical for customer's awareness. Notifications allow customers to subscribe to any of the available events using the following channels:

- SMS
- Email
- Slack (via Webhook)
- SNMP Traps

IRP service `Irrpushd` provides this feature. In order for Notifications to be delivered correctly the corresponding channel configuration shall be provided. By default only email notifications can be delivered since IRP uses the embedded system email service to send them.

More so, users should subscribe for specific events.

- Only events for valid subscriptions using correctly configured channels will be delivered.

Refer section [Notifications and Events](#) for details about configuring, subscribing and contents of notifications.

Refer section [Notification and events](#) for details about individual configuration parameter.

Events

The list of events monitored by IRP that can generate notifications is provided below.

When one of the IRP components detects a transition from normal to abnormal traffic behavior or back it fires these events:

- Abnormal correction: irpflowd
- Abnormal correction: irpspand
- Inbound traffic low: SPAN
- Inbound traffic low: Flow
- Inbound traffic normal: Flow
- Inbound traffic normal: SPAN
- Outbound traffic low: SPAN
- Outbound traffic low: Flow
- Outbound traffic normal: Flow
- Outbound traffic normal: SPAN

When Commit Control limits are exceeded per provider or overall one of the following events fires. Refer section 5.12 for configuring the actual limits of the events.

- Commit Control overload by X Mbps
- Commit Control overload by X%
- Commit Control provider X overloaded by Y Mbps
- Commit Control provider X overloaded by Y%

When an IRP component (re)loads the configuration it validates it and depending on results fires one of the following events:

- Configuration Invalid: Bgpd
- Configuration Invalid: Core
- Configuration Invalid: Explorer
- Configuration Invalid: Irpapid
- Configuration Invalid: Irpflowd
- Configuration Invalid: Irpspand
- Configuration Ok: Bgpd
- Configuration Ok: Core
- Configuration Ok: Explorer
- Configuration Ok: Irpapid
- Configuration Ok: Irpflowd
- Configuration Ok: Irpspand

Outage detection algorithm fires one of the following events when it confirms congestion or outage problems and reroutes traffic around it:

- Congestion or Outage
- Outage: Confirmed and rerouted

Explorer periodically checks the PBRs and its expected probing performance and triggers the following events:

- Failed PBR (IPv6) check for provider
- Failed PBR (IPv4) check for provider
- Successful PBR (IPv4) check for provider
- Successful PBR (IPv6) check for provider
- Explorer performance low
- High number of VIP prefixes degrades IRP performance

IRP BGP Internal and External monitors fire the following events:

- ExternalMonitor (IPv4) Failed status for a provider. All improvements towards the provider will be withdrawn.
- ExternalMonitor (IPv4) OK status for a provider. All improvements towards the provider will be announced.
- ExternalMonitor (IPv6) Failed status for a provider. All improvements towards the provider will be withdrawn.
- ExternalMonitor (IPv6) OK status for a provider. All improvements towards the provider will be announced.
- InternalMonitor (IPv4) Failed status for a provider. All improvements towards the provider will be withdrawn.
- InternalMonitor (IPv4) OK status for a provider. All improvements towards the provider will be announced.
- InternalMonitor (IPv6) Failed status for a provider. All improvements towards the provider will be withdrawn.
- InternalMonitor (IPv6) OK status for a provider. All improvements towards the provider will be announced.

When statistics collection over SNMP is up or down IRP fires the following events:

- Provider SNMP stats down: X
- Provider SNMP stats up: X

Bgpd raises these events when BGP sessions are established/disconnected:

- IRP BGP session disconnected
- IRP BGP session established

When IRP identifies conditions to re-route traffic (make an improvement) and additionally it considers the differences to be excessive it raises these events:

- Excessive packet latency for prefix
- Excessive packet loss for prefix
- Improvements spike

- Low rate of announced IPv4 improvements
- Low rate of announced IPv6 improvements
- New improvement

Once an IRP component is started, stopped or restarted it raises the following events:

- Service started: Bgpd
- Service started: Core
- Service started: Explorer
- Service started: Irpapid
- Service started: Irpflowd
- Service started: Irpspand
- Service stopped: Bgpd
- Service stopped: Core
- Service stopped: Explorer
- Service stopped: Irpapid
- Service stopped: Irpflowd
- Service stopped: Irpspand

SNMP Traps

SNMP traps is a widely used mechanism to alert about and monitor a system's activity.

IRP SNMP traps not only notify about some IRP platform event but also include the list of varbinds which contain detailed information related to the thrown trap. The complete list of traps and varbinds with their descriptions can be found at `/usr/share/doc/irp/NOCTION-IRP.mib`

1.2.15 IRP API

IRP exposes a web API that uses HTTP verbs and a RESTful endpoint structure. Request and response payloads are formatted as JSON.

The API is running on the IRP instance and is reachable by default over SSL at port 10443. If called directly from the IRP instance server the API can be accessed at `https://localhost:10443` Use `https://hostname:10443/` in order to access the API from elsewhere on the network.

An IRP user id is required to access most of the API services. Use IRP's Frontend to manage users or configure external User Directories (refer [Security Configuration](#)). IRP API uses an authenticating mechanism based on authentication tokens. The token is passed as a query parameter for all API requests that require authentication.

The API comes with an API Reference available from IRP's Help menu and the API References also includes a sample PHP application with source code to aid in development.


IRP's API is powered by Irpapid service that can be started, stopped, configured like any other IRP service. Refer to the [5.3](#) section for Irpapid configuration parameter details.

1.2.16 IRP Failover

Overview


IRP offers failover capabilities that ensure Improvements are preserved in case of planned or unplanned downtime of IRP server.


IRP's failover feature uses a master-slave configuration. A second instance of IRP needs to be deployed in order to enable failover features. For details about failover configuration and troubleshooting refer [Failover Configuration](#).

 A failover license is required for the second node. Check with Noction's sales team for details.

IRP's failover solution relies on:

- slave node running same version of IRP as the master node,
- MySQL Multi-Master replication of 'irp' database,
- announcement of the replicated improvements with different LocalPref and/or communities by both nodes,
- monitoring by slave node of BGP announcements originating from master node based on higher precedence of master's announced prefixes,
- activating/deactivating of slave IRP components in case of failure or resumed work by master,
- syncing master configuration to slave node.

 For exact details about IRP failover solution refer to configuration guides (2.13, 3.12.13.5), template files, and (if available) working IRP configurations. For example, some 'irp' database tables are not replicated, 'mysql' system database is replicated too, some IRP components are stopped.

 IRP versions 3.5 and earlier do not offer failover capabilities for Inbound improvements. It is advised that in these versions only one of the IRP instances is configured to perform inbound optimization in order to avoid contradictory decisions. In case of a failure of this instance inbound improvements are withdrawn.

An overview of the solution is presented in the following figure:

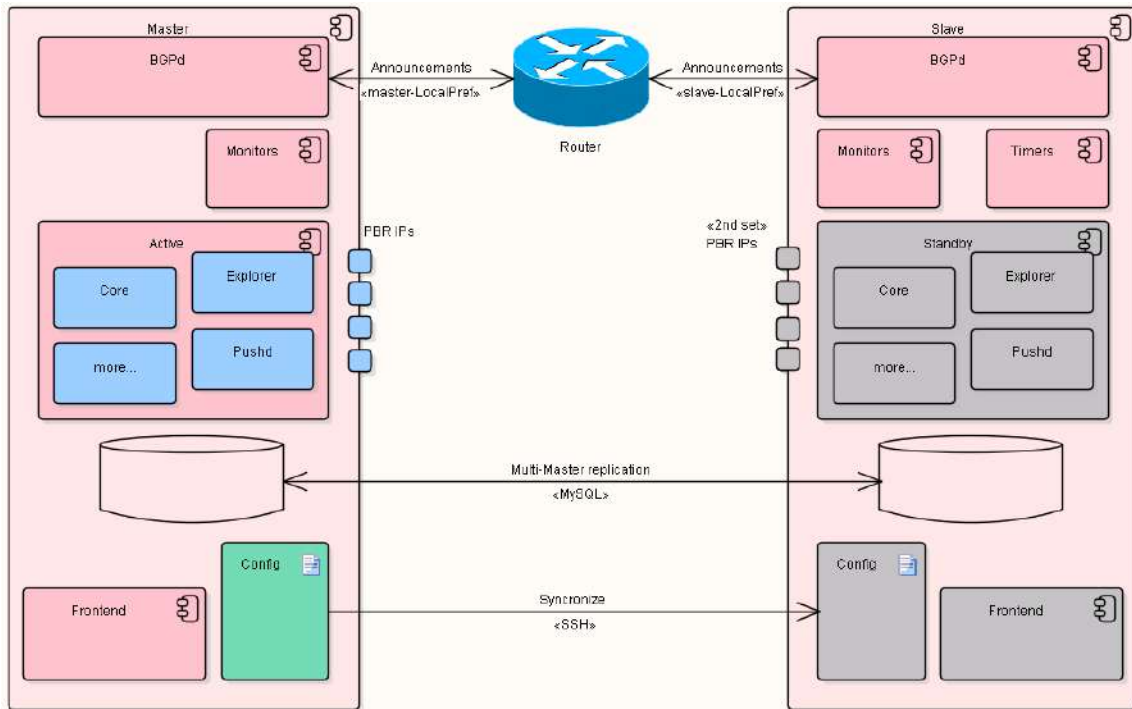


Figure 1.2.10: Failover high level overview

The diagram highlights the:

- two IRP nodes - Master and Slave,
- grayed-out components are in stand-by mode - services are stopped or operating in limited ways. For example, the Frontend detects that it runs on the slave node and prohibits any changes to configuration while still offering access to reports, graphs or dashboards.
- configuration changes are pushed by master to slave during synchronization. SSH is used to connect to the slave.
- MySQL Multi-Master replication is setup for 'irp' database between master and slave nodes. Existing MySQL Multi-Master replication functionality is used.
- master IRP node is fully functional and collects statistics, queues for probing, probes and eventually makes Improvements. All the intermediate and final results are stored in MySQL and due to replication will make it into slave's database as well.
- Bgpd works on both master and slave IRP nodes. They make the same announcements with different LocalPref/communities.
- Bgpd on slave node monitors the number of master announcements from the router (master announcements have higher priority than slave's)
- Timers are used to prevent flapping of failover-failback.

Requirements

The following additional preconditions must be met in order to setup failover:

1. second server to install the slave,
2. MySQL Multi-Master replication for the irp database.

❌ MySQL replication is not configured by default. Configuration of MySQL Multi-Master replication is a mandatory requirement for a failover IRP configuration. Failover setup, and specifically MySQL Multi-Master replication should follow a provided failover script. Only a subset of tables in irp database are replicated. Replication requires extra storage space, depending on the overall traffic and platform activity, for replication logs on both failover nodes.

1. a second set of BGP sessions will be established,
2. a second set of PBR IP addresses are required to assign to the slave node in order to perform probing,
3. a second set of improvements will be announced to the router,
4. a failover license for the slave node,
5. Key-based SSH authentication from master to slave is required. It is used to synchronize IRP configuration from master to slave,
6. MySQL Multi-Master replication of 'irp' database,
7. IRP setup in Intrusive mode on master node.

❌ In case IRP failover is setup in a multiple Routing Domain configuration and IRP instances are hosted by different RDs this must be specified in IRP configuration too. Refer [Optimization for multiple Routing Domains](#), `global.master_rd`, `global.slave_rd`.

Failover

IRP failover relies on the slave node running the same version of IRP to determine if there are issues with the master node and take over if such an incident occurs.

Slave's Bgpd service verifies that announcements are present on a router from master. If announcements from master are withdrawn for some reason the slave node will take over.


⚠️ In order for this mechanism to work IRP needs to operate in Intrusive mode and master's node announcements must have higher priority than the slave's.

During normal operation the slave is kept up to date by master so that it is ready to take over in case of an incident. The following operations are performed:


- master synchronizes its configuration to slave. This uses a SSH channel to sync configuration files from master to slave and process necessary services restart.
- MySQL Multi-Master replication is configured on relevant irp database tables so that the data is available immediately in case of emergency,
- components of IRP such as Core, Explorer, Irppushd are stopped or standing by on slave to prevent split-brain or duplicate probing and notifications,
- slave node runs Bgpd and makes exactly the same announcements with a lower BGP LocalPref and/or other communities thus replicating Improvements too.

❌ It is imperative that master's LocalPref value is greater than slave's value. This ensures that master's announcements are preferred and enables slave to also observe them as part of monitoring.

In case of master failure its BGP session(s) goes down and its announcements are withdrawn.

 Slave node only considers that master is down and takes over only if master's Improvements are withdrawn from all edge routers in case of networks with multiple edge routers.


The same announcements are already in router's local RIB from slave and the router chooses them as best.

 This is true only if LocalPref and/or communities assigned to slave node are preferred. If other most preferable announcements are sent by other network elements, no longer announcements from slave node will be best. This defeats the purpose of using IRP failover.

At the same time, Failover logic runs a set of timers after master routes are withdrawn (refer [global.failover_timer_fail](#)). When the timers expire IRP activates its standby components and resumes optimization.


Failback

IRP includes failback feature too. Failback happens when master comes back online. Once Bgpd on the slave detects announcements from master it starts its failback timer (refer [global.failover_timer_failback](#)). Slave node will continue running all IRP components for the duration of the failback period. Once the failback timer expires redundant slave components are switched to standby mode and the entire setup becomes normal again. This timer is intended to prevent cases when master is unstable after being restored and there is a significant risk it will fail again.

 During failback it is recommended that both IRP nodes are monitored by network administrators to confirm the system is stable.

Recovery of failed node


IRP failover configuration is capable to automatically restore its entire failover environment if downtime of failed node is less than 24 hours.

 Recovery speed is constrained by restoring replication of MySQL databases. On 1Gbps non-congested links replication for a full day of downtime takes approximately 30-45 minutes with 200-250Mbps network bandwidth utilization between the two IRP nodes. During this time the operational node continues running IRP services too.

If downtime was longer than 24 hours MySQL Multi-Master replication is no longer able to synchronize the databases on the two IRP nodes and manual MySQL replication recovery is required.

Upgrades

Failover configurations of IRP require careful upgrade procedures especially for major versions.

 It is imperative that master and slave nodes are not upgraded at the same time. Update one node first, give the system some time to stabilize and only after that update the second node.

1.2.17 Inbound bandwidth optimization

1.2.17.1 Inbound Commit Control

Starting with version 3.4 IRP introduced optimization of Inbound traffic. Inbound bandwidth control reshapes the traffic from different providers targeting your sub-prefixes .

IRP uses well known and proven BGP mechanisms to instruct your routers to adjust their advertisements of your network segments to upstream providers and subsequently to the World. The adjusted advertisements take advantage of existing BGP policies implemented by edge routers worldwide in order to increase or decrease the preference of your internal network segments as advertised by one or another AS to the world. This allows more traffic to lean towards some of your upstream providers and less towards others. In case of an incident that your multihomed configuration is designed to be resilient against, the entire world still knows of the alternative routes towards your network and will be able to adjust accordingly.

Noction's IRP Inbound feature:

- advises your edge routers to advertise different prefixes of your network with different counts of BGP prepends to each of your upstream providers;
- monitors inbound and outbound traffic on your network interfaces using standard protocols (SNMP, NetFlow, sFlow) to determine if there is need for action;
- continuously assesses network health and performance to ensure that when overloads are possible it knows good and reliable alternative routes to use;
- uses a proprietary inferring mechanism that takes into account the inertial nature of the Internet to react to inbound preference changes and thus dampens the urge to “act now” and destabilize the network;
- provides you with configuration, monitoring and visualization tools that allow you to cross-check and validate its actions.

The entire Inbound optimization solution covers:

1. Segmenting your network into prefixes that are controlled by IRP
2. Coordinating communication over BGP between IRP and routers
3. Setting network conditions for making Inbound Improvements
4. Reviewing Inbound optimization results

i Starting with version 3.7 IRP has the capability to also monitor and improve transit routes. Refer [Optimization of transiting traffic](#) for details.

The use cases below highlight typical scenarios when IRP's inbound bandwidth control capabilities might come handy:

Lets start with the case of unbalanced inbound and outbound peering You have a peering agreement with one of your neighbors to exchange traffic. Unfortunately the rest of the neighboring network configuration significantly unbalances the volumes of inbound and outbound traffic.




Figure 1.2.11: Unbalanced inbound/outbound peering

You rely on manipulating prefix announcements towards neighbors in order to shape the traffic. Unfortunately this is a reactive solution and consumes a lot of time while at the same time pushing the balance either one way or another. Often this pushes the network into the second typical scenario.

Fluctuating traffic shape A multihomed configuration overwhelms some links while other links remain barely used. Your network administrators frequently get alerts during peak network use and they manually add or remove preprends or altogether remove some inbound prefixes from being announced to affected neighboring links. These changes are sometime forgotten or other times just push the bulk of traffic towards other links and the problem re-appears.




Figure 1.2.12: Fluctuating inbound traffic

For both above scenarios Inbound commit control in IRP can automate most of the inbound traffic shaping operations. It uses a typical solution for traffic shaping and specifically manipulates prepend counts announced to different providers and this eventually propagates across the Internet thus diverting traffic as desired. IRP automates your routine day to day traffic shaping actions and probably makes significantly more adjustments than you can afford to. At the same time it offers you reliable reviewing and fine-tuning options that will allow you to adjust IRP's work in a changing and evolving network.

Refer [Inbound configuration](#), [Current inbound improvements](#), [Inbound traffic distribution](#), [Inbound rule](#) for details.

1.2.17.2 Inbound Performance Optimization

1.2.18 Inbound performance optimization

Starting with version 4.0 IRP is able to perform loss and latency optimization of inbound traffic.

This feature allows improving the inbound traffic performance by deflecting it from the worst-performing provider. The improvement actions do not guarantee the best path for the inbound traffic but minimize the usage of the worst path. For the traffic deflection consistency, IRP uses provider traffic engineering capabilities.

i Some network providers give customers the ability to influence their traffic based on the announced BGP community, the so-called traffic engineering capability. Customers of such providers can check if the upstreams support this functionality by reviewing the upstreams BGP routing policies, formerly published in Whois/RDAP for provider's Autonomous System, or by contacting providers directly. In other words, traffic engineering capability, therefore, represents a list of BGP communities that a provider configures on their side for the customer's use. When prefixes tagged with a specific BGP community from such lists get announced by clients, the upstream provider applies the corresponding action to those prefixes.

The requirements for the Inbound performance feature are:

- Provider traffic engineering capability using BGP community
- Ability to re-announce providers' BGP table to IRP. Configuring the BGP monitoring Protocol (BMP [BMP monitoring station](#)) is recommended
- Flow export (The feature is not supported in the SPAN-only setups)

w Calculation of an inbound performance improvement takes more time, when compared to the regular outbound improvement. The reason is that the inbound traffic is affected much slower than the outbound traffic and the calculation method is different.


For the Inbound performance feature to work, IRP must know what traffic should be monitored and improved. For that, the Inbound Performance rule(s) must be configured. More about Inbound performance rules: [Inbound Performance Rules](#)

IRP performs the analysis of network performance indicators and applies improvements per individual preconfigured rules in the following order:

- Probes prefixes that generate inbound traffic to build a statistical model of the route performance per each inbound rule
- Reprobes prefixes according to the configured inbound reprobing interval ([irpinperfd.probing.interval](#))
- Compares the reprobed results with the statistical model
- In case the poor inbound traffic metrics are obtained – IRP makes an improvement by announcing the customer prefixes with the corresponding TE community
- Once IRP applies an inbound improvement, its relevance is checked on subsequent probing sessions


There are two modes in which the inbound performance improvements can be enabled:

- Automated - IRP enables/disables all the inbound performance improvements automatically once an improvement is calculated
- Moderated - stands for manual confirmation of the inbound performance improvements. Once IRP detects an improvement, a suggestion to enable it gets displayed in the GMI interface. IRP does not perform any improvements without user confirmation. If an improvement is not enabled in time before a subsequent cycle of checks, IRP reviews the proposal and adjusts it as needed, keeping the recommendation, making a different one, or removing the improvement proposal


 The inbound performance improvements can be enabled manually as per a particular rule. This is done at the discretion of the user. IRP does not make any calculations/proposals in such cases

To view how to configure the Inbound Performance feature, go to: [Inbound performance optimization configuration](#)

1.2.19 Flowspec policies

 Flowspec policies can be used only in conjunction with Flowspec capable routers.

Starting with version 3.5 IRP has support of Flowspec policies. This means that Flowspec capability is recognized and can be used accordingly for BGP sessions established by IRP. In short Flowspec defines matching rules that routers implement and enforce in order to ensure that only desirable traffic reaches a network. Flowspec by relying on BGP to propagate the policies uses a well understood and reliable protocol.

 As specified by BGP, when a session disconnects all the announcements received from that session are discarded. The same is generally true for Flowspec policies. Still, since some vendors recognize Flowspec policies but implement them using capabilities other than BGP a confirmation is needed whether on session disconnect the specific router model indeed removes all the underlying constructs and reverts to a known state.

IRP not being involved in direct packet forwarding expects that Flowspec policies are implemented at least by your edge routers. If upstream providers also offer Flowspec support these policies can be communicated upstream where their implementation is even more effective.

Eventually Flowspec policies help ensure that traffic entering or exiting your network conforms to your plans and expectations. The main use cases that can be accomplished with Flowspec policies in IRP allows:

- controlling bandwidth usage of your low priority traffic towards external services, for example throttling bandwidth usage originating on your backup systems towards off-premises services.
- anticipating inbound traffic towards your services and shaping bandwidth use in advance, for example anticipating low numbers of legitimate customers from Russia, China or India on your e-commerce services and setting high but controllable rate limits on packets originating in those networks.
- reacting on a packet flooding incident by dropping specific packets, for example dropping all packets targeting port 53.
- redirecting some traffic for scrutiny or cleansing, for example forwarding port 80 packets through an intelligent device capable of detecting RUDY, slow read or other low-bandwidth/amplification attacks.

IRP Flowspec policies rely on a minimal set matching rules and actions that offer most of the capabilities while keeping the learning curve low and integration simple:

- Source or destination IP address specified as either CIDR format prefix or direct IP address
- Traffic protocols, for example TCP, UDP or ICMP
- Source or destination TCP/UDP ports
- Throttle, drop and redirect actions.

i It is important to note that IRP does not cross-validate Flowspec policies with improvements. While it is possible that for example a Flowspec redirect action pushes some traffic a different way to what an improvement advises, usually improvements cover many prefixes and while there will be a contradiction for one prefix there will be many other prefixes that IRP improves to compensate for these unitary abnormalities. It is recommended that Flowspec policies take precedence over improvements in order to benefit from this compensating nature of improvements.

Consider that depending on whether source or destination prefix belongs to your network the policy applies to either inbound or outbound traffic while the choice of ports allows targeting different traffic types.

Compare Flowspec policies to the already well known [Routing Policies](#) . For further details regarding Flowspec configuration refer [Flowspec policies](#).

1.2.20 Throttling excessive bandwidth use with Flowspec

IRP can be configured to automatically add throttling Flowspec policies for prefixes that started using abnormal volumes of traffic. This feature can be used only if your network has Flowspec capabilities. Refer [Flowspec policies](#) for details about Flowspec.

- In case thresholds for excessive bandwidth use are set to very aggressive levels IRP can create large numbers of Flowspec policies.

This feature can be described as:

- configure excess threshold and throttling multipliers
- periodically determine current and average prefix bandwidth usage for this hour of the day
- verify if current usage exceeds the average by a larger factor then the threshold multiplier
- rate-limit excessive bandwidth usage prefixes at their average use times throttling multiplier.

Past throttling rules are revised if/when prefix abnormal usage pattern ends.

For example, when a prefix usually consumes 1-2Mbps of traffic and its current bandwidth spikes tenfold and the spike is sustained for a significant period of time a throttling rule limiting usage by this prefix at 5-6Mbps will still offer ample bandwidth for normal service use and will also protect other services from network capacity starvation.

Refer [Throttling excessive bandwidth use with Flowspec, Core configuration](#) for details.

1.2.21 Maintenance windows

Maintenance works are on everybody's agenda in current fast paced and continuously evolving networks. During maintenance network engineers are very busy and will welcome any help their systems can offer in carrying out those works with the least amount of headaches. IRP is clearly not in the top of network engineer's priorities and asking to suspend or shutdown a provider immediately before a maintenance window starts and restart the provider back once the maintenance works end is not very helpful if not even annoying.

Instead IRP offers the facility to plan maintenance windows in advance. Knowing when a maintenance window starts and ends, IRP excludes that specific provider link from either performance optimization or bandwidth control. More so, IRP has the capability to reshape the traffic flowing in and out of a network to anticipate any downtime on a link.

i Setting a maintenance window by router sets each of the providers on the router with a maintenance window of their own.

Properly configured maintenance windows allows IRP time to move most of the outbound traffic and deflect most of inbound traffic away from the provider link that is scheduled for maintenance. Having only a small fraction of traffic or none at all on the maintenance link before the downtime starts avoids any (shall we say, catastrophic) spikes, possible overloads and consequently unpredictable behavior of the remaining live network equipment.

Specifically the following applies:

- a maintenance window is configured in advance and can be removed/revised at any time,
- a maintenance window sets details for single provider. If needed multiple maintenance windows can be setup and even overlapping maintenance windows are OK.
- IRP highlights maintenance windows in IRP's Frontend sidebar so that it is easy to spot current maintenance window status.
- optionally IRP can preserve existing improvements so that once the maintenance window ends improvements are reimplemented. It is advised that this feature is used only when the maintenance window is very short (a few minutes).
- an unloading period can be setup. During unloading IRP actively re-routes outbound prefixes through other available providers. While IRP is able to make most of the unloading improvements fast consideration shall be given to the announcement rate limitations setup in Bgpd in order for all the improvements to reach network routers in time for maintenance window starting time.
- a prepend time can be setup. This is only applicable if Inbound optimization is operational. If this time is setup then IRP will prepend configured inbound prefixes with the maximum allowed number of prepends through the provider link under maintenance in order to deflect inbound traffic towards other providers. Refer to [Inbound bandwidth optimization](#) for more details about Inbound optimization.

Refer [Maintenance windows](#) for details how to configure, review, create, edit, delete maintenance windows.

1.2.22 Optimization of transiting traffic

Starting with IRP 3.7 IRP introduces optimization of transiting traffic capabilities. Optimization of transiting traffic is an enhancement of [Inbound bandwidth optimization](#).

Optimization of transiting traffic relies on the same method of influencing Internet-wide routing - manipulating best path selection by increasing AS Path length for a prefix carrying traffic on an undesirable interface. The most secure and effective way of AS Path manipulation is to apply these changes on a router facing a provider. Edge routers prepend routes according to designated communities.

Besides being an enhancement of Inbound optimization, transit prefixes are governed by a different set of constraints:

- The number of potential routes is very large and only some will/should be targeted for optimization. For this IRP uses filters by ASN/prefix allowing network administrators to configure what segments of the Internet to focus IRP's attention to.
- The improvements are visible on the Internet and excessive route changes can be flagged by external monitoring services as flapping or route instability. IRP protects against this by rate limiting number of route changes through a specific provider.
- Once a route has been improved all its traffic might be diverted away from this network and no new statistics will be available to make further inferences. In such cases IRP reverts old transit improvements during network's off-peak hour by either decreasing the number of advertised prepends or withdrawing the improvement altogether (configurable).
- Transit improvements apply to the same prefixes as do outbound improvements. The outbound improvements are withdrawn in order to avoid the risk of contradictory routing decisions.


Implementing optimization of transiting traffic introduces a series of risks and some inherent drawbacks, for example:

- Potential blackholing of traffic when all alternative routes are withdrawn. IRP implements a protection against this by traversing all RIB-in entries for improved transit prefixes and confirming that the route is still being announced by other providers besides IRP. Still there can be a short time period between confirmations when all alternative routes have been withdrawn and IRP did not yet get a chance to re-confirm this. Attempting to reduce this time period by setting up a higher frequency of confirmations leads to increased load on the router and a tradeoff needs to be made for this. For example when the number of providers is quite large the probability that a route will be withdrawn through all of them is quite small and thus the frequency of confirmations can be reduced too.
- Additional CPU load on edge router(s) for servicing mandatory SNMP requests that check alternative route presence and BGP Best Path selection inside IRP.
- Working with missing BGP attributes that is not available over SNMP.
- Strict upper limits on the number of possible Inbound improvements are imposed by the trade-off required to reduce excessive router's CPU load.

Refer [Optimization of transiting traffic](#) for details.

1.2.23 Circuit issues detection

Starting with IRP 3.8 IRP adds excessive loss circuit issues detection features.


 Circuit issues detection feature is available for transit providers only.

When this feature is enabled for a provider IRP uses past probing data to detect when it suffers from excessive levels of packet loss. To determine excessive loss IRP compares a provider's average loss over an immediate past time horizon, number of probes and average loss difference from other providers. Depending on packet loss thresholds IRP can attempt different actions on the network.

Every time a circuit issue is detected IRP will raise corresponding alerts that can be subscribed to. Network engineers or external network management systems can act on them.

Additionally IRP even though it is constrained on how much can do, ensures that the following will take place:

- provider is marked with an issue badge and excluded from being considered a candidate for performance improvements,
- reprobing is performed for destination prefixes routed through affected provider,
- outbound improvements through affected provider are withdrawn,
- max prepends are announced for inbound and transit prefixes through affected provider,
- FlowSpec rules to induce drop of BGP session towards affected provider are added,


 Note that this is only possible if FlowSpec is enabled and the network is capable of processing FlowSpec rules. Dropping the BGP session with the affected provider causes all outbound and inbound traffic through this provider to be re-routed through known good providers.

- affected provider is monitored to detect if the issue was temporary and if loss averages return to normal restore it to a good state.

Enable this feature for each of the designated provider as detailed in [Configuration editor: Provider name](#) and review circuit issue detection thresholds as detailed in [Core configuration](#).

1.2.24 Threat Mitigation

Starting with IRP 4.1, a new feature called Threat Mitigation is added to the platform. It incorporates the statistics collection as well as the blackholing mechanism, present in the previous product versions, nevertheless offering a fully automated threshold-based threat mitigation instrument that introduces Flowspec in addition to the RTBH.

 RTBH feature allows redirection of traffic to a non-existent resource (a so-called black hole), or the blocking of the unwanted traffic in a provider's network to prevent such traffic from entering the user's network.

Threat Mitigation can operate in the following modes:

Automated IRP performs a particular threat mitigation action automatically when an attack is detected

Moderated stands for manual confirmation of the Threat Mitigation action. Once IRP detects an attack, a suggestion to enable the mitigation rule gets displayed in the GMI interface. IRP does not perform any actions without user confirmation. If the action is not enabled in time before a subsequent cycle of checks, IRP reviews the proposal and adjusts it as needed, keeping the recommendation, making a different one, or removing it altogether

Manual Allows to set up only manual threat mitigation rules

Disabled IRP does not perform any threat mitigation actions.

The Threat Mitigation feature is based on threshold rules set by IRP users. In the context of DoS/DDoS attack detection, threshold values refer to the rate of kilo packets or megabits per second. Specifically, a detection threshold represents the rate at which an IRP instance raises an alert for an attack and takes appropriate action as per the predefined rule. For instance, if the Flowspec rule threshold is set to 80 kilo packets per second, an alert will be generated once incoming packets flow towards a destination exceeds this bound and the consecutive Flowspec action gets triggered. Hence, indicating an appropriate detection threshold value in rules is critical to providing an efficacious response to DoS threats while reducing false alerts.

FlowSpec mitigation method can be chosen to filter out smaller attacks while the Remote Triggered Blackhole should be sent to providers to block large volume attacks.

1.2.24.1 Configuring Blackholing

A Provider in IRP should be configured before it could be used for blackholing.

IRP should know next-hop (`bgpd.peer.X.blackholing.ipv4.next_hop`, `bgpd.peer.X.blackholing.ipv6.next_hop`), localpref (`bgpd.peer.X.blackholing.localpref`) and community (`peer.X.blackholing.community`) values to be able to send a route to a user's network.

A user's router is responsible to distinguish communities sent by an IRP instance and advertise blackholing routes to a provider's router used to receive such routes.

1.2.24.2 Configuring FlowSpec

FlowSpec feature should be enabled globally in `global.flowspec` then for each BGP session in `bgpd.peer.X.flowspec`.

1.3 IRP Optimization modes

1.3.1 Performance optimization

Performance optimization mode makes sure that traffic flows through the best performing routes by reducing packet loss and latency, while ignoring other characteristics such as provider bandwidth usage

or transit cost. The system analyzes each of the connected providers and compares only their performance metrics in order to choose the best candidate and make an improvement.

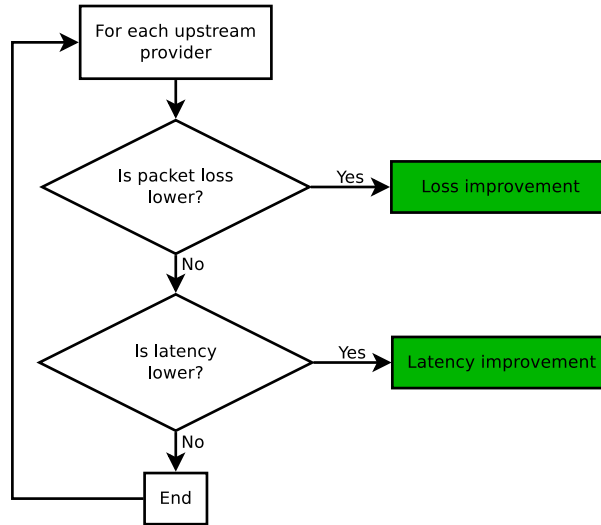


Figure 1.3.1: Performance optimization algorithm

First of all, IRP considers packet loss. If loss is lower and the difference is greater than a predefined value, then the system checks if sending the traffic through this provider will not cause any network congestion. If it confirms that the traffic can flow freely, the system declares the provider as the best one to route through.

However, if loss values are equal or the difference between providers is smaller than predefined, the system continues by comparing latency values. If latency is lower and the difference in latency between providers is greater than predefined, then the system declares a latency-based improvement.

1.3.2 Cost optimization

Cost optimization mode decreases packet loss and improves the cost while not worsening latency more than allowed. If IRP cannot improve costs it tries to reduce latency. The platform runs the same algorithm for loss comparison as in performance optimization mode.

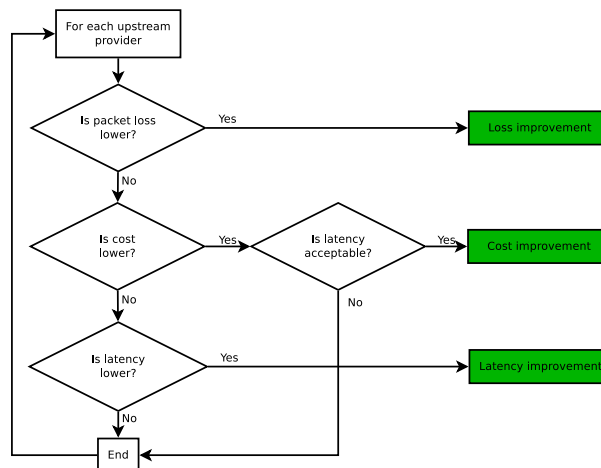


Figure 1.3.2: Cost optimization algorithm

i Note that the diagram does not highlight higher cost cases. IRP operating in cost optimization mode can make improvements towards higher cost providers only when current routes suffer from loss.

Before comparing latency values IRP compares the transit cost for each provider. If the cost of the new provider is better, the system goes further by checking the latency and validates the cost improvement only if the latency is not worsening more than predefined. However, if the cost cannot be improved, IRP tries to make a latency improvement using the same algorithm as in performance mode. Thus, if there is no way to make a cost improvement, system reroutes the traffic to the best performing provider.

1.3.3 Commit Control

Commit Control, allows to keep the commit levels for each provider at a pre-configured level. It includes bandwidth control algorithms for each provider as well as the active traffic rerouting, in case bandwidth for a specific provider exceeds the configured limit. Commit Control also includes passive load adjustments inside each provider group.

A parameter called “precedence” (see [peer.X.precedence](#)) is used to set the traffic unloading priorities, depending on the configured bandwidth cost and providers throughput. The platform will reroute excessive bandwidth to providers, whose current load is less than their 95th percentile. If all providers are overloaded, traffic is rerouted to the provider with the smallest precedence - usually this provider has either the highest available bandwidth throughput, or the lowest cost. The higher is the precedence, the lower is the probability for the traffic to be sent to a provider, when its pre-configured 95th percentile usage is higher.

i IRP usually allows CC improvements when the candidate providers have better or equal loss to current route. This can be configured under [core.commit_control.loss_override](#).

See also: [core.commit_control](#).

1.3.3.1 Flexible aggressiveness of Commit algorithm based on past overloads

Metering bandwidth usage by the 95th presents the following alternative interpretation - the customer is allowed to exceed his limits 5% of times. As such, IRP assumes there's a schedule of overloads based on the current time within a month and an actual number of overloads already made.




Figure 1.3.3: More aggressive when actual exceeds schedule

The sample image above highlights the remaining scheduled and the actual amount of allowed overloads for the month decreasing. Whenever IRP depicts that the actual line goes below schedule (meaning the number of actual overloads exceeds what is planned) it increases its aggressiveness by starting unloading traffic earlier than usual. For example, if the commit level is set at 1Gbps and IRP will start unloading traffic at possibly 90% or 80% depending on past overloads count.

The least aggressive level is set at 99% and the most aggressive level is constrained by configuration parameter `core.commit_control.rate.low` + 1%.

This is a permanent feature of IRP Commit Control algorithm.

1.3.3.2 Trigger commit improvements by collector

Commit control improvements are made after the flows carrying traffic are probed and IRP has fresh and relevant probe results. Subsequent decisions are made based on these results and this makes them more relevant. Still, probing takes some time and in case of fluctuating traffic patterns the improvements made will have reduced impact due to flows ending soon and being replaced by other flows that have not been probed or optimized.

For networks with very short flows (average flow duration under 5 minutes) probing represents a significant delay. In order to reduce the time to react to possible overload events IRP added the feature to trigger commit control improvements on collector events. When Flow Collector detects possible overload events for some providers, IRP will use data about past probed destinations in order to start unloading overloaded providers early. This data is incomplete and a bit outdated but still gives IRP the opportunity to reduce the wait time and prevent possible overloads. Later on, when probes are finished another round of improvements will be made if needed.

Due to the fact that the first round of improvements is based on older data, some of the improvements might become irrelevant very soon. This means that routes fluctuate more than necessary while on average getting a reduced benefit. This is the reason this feature is disabled by default. Enabling this feature represents a tradeoff that should be taken into consideration when weighing the benefits of a faster react time of Commit Control algorithm.

This feature is configurable via parameter: `core.commit_control.react_on_collector`. After enabling/disabling this feature IRP Core service requires restart.

1.3.3.3 Commit Control improvements on disable and re-enable

IRP up to version 2.2 preserved Commit Control improvements when the function was disabled globally or for a specific provider. The intent of this behavior was to reduce route fluctuation. In time these improvements are overwritten by new improvements.

We found out that the behavior described above ran contrary to customer's expectations and needs. Usually, when this feature is disabled it is done in order to address a more urgent and important need. Past Commit Control improvements were getting in the way of addressing this need and was causing confusion. IRP versions starting with 2.2 aligns this behavior with customer expectations:

- when Commit Control is disabled for a provider (`peer.X.cc_disable = 1`), this Provider's Commit Control improvements are deleted;
- when Commit Control is disabled globally (`core.commit_control = 0`), ALL Commit Control improvements are deleted.

1.3.3.4 Provider load balancing

Provider load balancing is a Commit Control related algorithm, that allows a network operator to evenly balance the traffic over multiple providers, or multiple links with the same provider.

For example, a specific network having an average bandwidth usage of 6Gbps has two separate ISPs. The network operator wants (for performance and/or cost reasons) to evenly push 3Gbps over each provider. In this case, both upstreams are grouped together (see [peer.X.precedence](#)), and the IRP system passively routes traffic for an even traffic distribution. Provider load balancing is enabled by default via parameter `peer.X.group_loadbalance`.




Figure 1.3.4: Provider load balancing

1.3.3.5 Commit control of aggregated groups

Customers can deploy network configurations with many actual links going to a single ISP. The additional links can serve various purposes such as to provision sufficient capacity in case of very large capacity requirements that cannot be fulfilled over a single link, to interconnect different points of presence on either customer (in a multiple routing domain configuration) or provider sides, or for redundancy purposes. Individually all these links are configured in IRP as separate providers. When the customer has an agreement with the ISP that imposes an overall limitation on bandwidth usage, these providers will be grouped together in IRP so that it can optimize the whole group.

The rationale of this feature as illustrated in the figure below is that if in the group overusages on one provider are compensated by underusages on another provider there is no need to take any action since overall the commitments made by the customer to the ISP have not been violated. Commit control algorithm will take action only when the sum of bandwidth usage on all providers in the group exceed the sum of bandwidth limits for the same group of providers.




Figure 1.3.5: Commit control of aggregated groups

The image above highlights that many overusages on the green line are compensated by purple line underusages so that the group usage is below group total limits. Only when the traffic on the purple line increases significantly and there are no sufficient underusages on the other providers in the group to compensate the overusages, Commit Control identifies overusages (highlighted with a red x on the drawing above) and takes action by rerouting some traffic towards providers outside the group.

⚠ It is important to note that in order for this feature to be effective there must be providers configured in IRP that are not part of this group. This way when candidate improvements are considered there are alternative routes via those providers that the traffic can be rerouted to.

In order to configure providers that are optimized as an aggregated group 1) first the providers will be configured with the same precedence in order to form a group; 2) the overall 95th limitation will be distributed across providers in the group as appropriate; 3) and finally load balancing for the group will be Disabled by parameter `peer.X.group_loadbalance`.

1.3.3.6 95th calculation modes

Commit control uses the 95th centile to determine whether bandwidth is below or above commitments. There are different ways to account for Outbound and Inbound traffic when determining the 95th value. IRP supports the following 95th calculation modes:

- Separate 95th for in/out: The 95th value for inbound and outbound traffic are independent and consequently bandwidth control for each is performed independently of each other. For this 95th calculation modes IRP monitors two different 95th for each inbound and outbound traffic levels.
- 95th from greater of in, out: At each time-point the greater of inbound or outbound bandwidth usage value is used to determine 95th.
- Greater of separate in/out 95th: 95th are determined separately for inbound and outbound traffic and the larger value is used to verify if commitments have been met.

Refer 5.14.5.

1.3.3.7 Other commit control features

Current improvements report has been improved to keep CC improvements applied during an algorithm cycle in a single sorted batch. This way the report no longer confuses by intermingling the improvements.








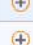
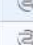



30	08:47:04	91.237.164.0/22	199540	B	A				Rerouted 9.11 Mbps from B[3] (44.61 Mbps, 109%) to A[2] (6.74 Mbps, 11%)
31	08:47:04	85.177.192.0/10	13184	B	A				Rerouted 6.09 Mbps from B[3] (35.49 Mbps, 87%) to A[2] (15.86 Mbps, 26%)
32	08:47:04	146.0.80.0/21	35362	B	A				Rerouted 5.70 Mbps from B[3] (29.40 Mbps, 72%) to A[2] (21.95 Mbps, 37%)
33	08:47:04	128.73.88.0/24	8402	B	A				Rerouted 4.26 Mbps from B[3] (23.71 Mbps, 58%) to A[2] (27.65 Mbps, 46%)

Figure 1.3.6: Sorted Commit Control improvements

The excerpt above highlights the new commit control improvements done at the same time. They unload 9.11, 6.09, 5.70, 4.26 Mbps from provider B to A thus unloading B from 109% to 58% and increasing the load on A from 11% to 46%. The data would have been more confusing if the improvements in that single batch were intermingled.

During retry probing IRP will delete Commit Control improvements with current bandwidth less than the configured threshold given by `core.commit_control.agg_bw_min`.

1.4 What is IRP Global Management Interface

IRP Global Management Interface (GMI) is a pane of glass interface that allows administrators to manage multiple Noction Intelligent Routing Platform instances and get access to various data and statistics for those instances from one, easy to access application.

GMI allows administrators to:

- Monitor multiple IRP instances and their performance
- Get comprehensive network performance analytics
- Facilitate network troubleshooting
- Automatically manage bandwidth levels for provider groups from various points of presence using the Global Commit feature

1.4.1 GMI Operating Details

GMI act as a standard web application, all user interactions are made via a web browser. The list of supported browsers includes:

- Google Chrome / Chromium version 42 and higher
- Safari version 10.1 and higher
- Firefox 64 and higher
- Edge 45 and higher
- Opera 42 and higher

1.4.2 GMI Technical Requirements

To deploy GMI, a series of hardware/software requirements need to be met

1.4.2.1 Hardware requirements

1. CPU

- Recommended Intel® i5 latest 10th, 9th, 8th, 7th, and 6th Generation Products with 5 threads

2. RAM

- Recommended size of at least 4 GB

3. HDD

- Recommended size of at least 4 GB SSD

1.4.2.2 Software requirements

Clean Linux system, with the latest CentOS 7, Red Hat Enterprise Linux 8 or Ubuntu Server LTS of x86_64 architecture installed on the server.

Chapter 2

Configuration

2.1 IRP software management

Software repository

IRP packages (current, new and old versions) are available in Noction's repositories hosted at repo.noction.com. Coordinate getting access to the repository with Noction's representatives.

Software installation and upgrade

Once the IRP packages repository is configured standard OS package management tools (yum/dnf/apt) are used to install, upgrade or downgrade IRP.

Installation:CentOS 7:


```
yum install irp
```

Red Hat Enterprise Linux 8:

```
dnf config-manager --enable repo PowerTools
dnf install irp
```

Ubuntu Server:

```
apt-get update
apt-get install irp
```

 Optional package *irp-documentation* containing this documentation in PDF format and API reference should be installed separately.

Upgrade to latest version:

CentOS:

```
yum upgrade "irp*"
```

Red Hat Enterprise Linux 8:

```
dnf upgrade "irp*"
```

Ubuntu:

```
apt-get update
apt-get install --only-upgrade irp\*
```

Downgrade to a specific version:


CentOS:

```
yum downgrade "irp*3.7*"
```

Red Hat Enterprise Linux 8:

```
dnf downgrade "irp*3.7*"
```

Ubuntu:

 IRP added support for Ubuntu Server starting with version 3.9. IRP versions prior to 3.9 will not be offered for Ubuntu Server.

Check for available IRP versions:

```
apt-cache policy irp
```

Edit file `/etc/apt/preferences.d/irp` (refer to `apt_preferences` man page for package pinning) to pin desired version:

```
Package: irp irp-*
Pin: version 3.9.0-RELEASE~build11806~trusty
Pin-Priority: 1001
```

```
apt-get update
apt-get upgrade irp
```

2.2 Configuration files

IRP is currently configured using a set of textual Unix-style configuration files, located under the `/etc/noction/` directory. In-line comments can be added to any line, using the `"#"` symbol.

Several separate configuration files are presented as follows:

- `/etc/noction/irp.conf` - the main IRP configuration file contains configuration parameters for all IRP components, including algorithm parameters, optimization modes definitions and providers settings.
- `/etc/noction/db.global.conf` - database configuration file for all IRP components, except the Frontend component.
- `/etc/noction/db.frontend.conf` - database configuration file for the Frontend component.
- `/etc/noction/exchanges.conf` - Exchanges configuration file (3.12.7.2).
- `/etc/noction/inbound.conf` - Inbound prefixes configuration file (3.12.12).
- `/etc/noction/policies.conf` - Routing Policies configuration file (1.2.9).
- `/etc/noction/user_directories.conf` - User Directories configuration file (3.12.10.2).

Additional configuration files can be used for several core, global and explorer preferences, as described in sections 5.2.24, 5.7.19 and 5.9.4.

A comprehensive list of all the IRP parameters along with their description can be found in the [Configuration parameters reference](#) chapter.

2.3 IRP Global Management Interface Configuration

2.3.1 GMI Initial Configuration

After the installation, GMI service will start automatically. On startup, GMI service starts listening on the following network ports: 80, 443.

When installed on the same server with an IRP instance, GMI service listens on the following network ports: 1443, 1080.

The GMI initial configuration can be performed by using the Initial Setup wizard (Initial Setup) which can be accessed via the main platform IP address over HTTP or HTTPS protocol.

2.3.2 GMI software management

Software repository

GMI packages (current and new) are available in Noction's repositories hosted at repo.noction.com. Coordinate getting access to the repository with Noction's representatives.

Software installation and upgrade

Once the GMI packages repository is configured, follow the instructions below to install or upgrade GMI.

Installation

CentOS 7

```
yum install irp-gmi
```

RedHat Enterprise Linux 8 Install irp-gmi package:

```
dnf install irp-gmi
```

Ubuntu

```
apt-get update  
apt-get install irp-gmi
```

Upgrade to the latest version

CentOS 7

```
yum upgrade irp-gmi
```

RedHat Enterprise Linux 8


```
dnf upgrade irp-gmi
```

Ubuntu

```
apt-get update  
apt-get upgrade
```

2.3.3 Starting, stopping and getting status of GMI backend

GMI acts as a Linux service, and can be started, stopped or restarted from the Linux terminal, using the commands listed below:

 Critical errors preventing startup of a particular component are logged to journalctl. Details can be obtained from logs stored in journalctl using the following command

```
journalctl -fu irp-gmi
```

Starting backend:

```
systemctl start irp-gmi
```

Stopping backend:

```
systemctl stop irp-gmi
```

Restarting backend:


```
systemctl restart irp-gmi
```

Backend status:

```
systemctl status irp-gmi
```

2.4 Starting, stopping and getting status of IRP components

In order to determine the status of IRP components or to start, stop and restart them standard Linux utilities are used.

 Critical errors preventing startup of a particular component are logged to console. Details can be obtained from logs stored in `/var/log/irp/` directory.

Managing software components in OS with systemd

Starting single component:

```
systemctl start explorer
```

 Multiple components separated by space can be listed


Stopping single component:

```
systemctl stop explorer
```

 Multiple components separated by space can be listed

Starting all components:

```
systemctl start irp.target
```

 This also starts all prerequisites

Stopping all components:

```
systemctl start irp-shutdown.target
```

Stopping all components except bgpd:


```
systemctl start irp-shutdown-except-bgpd.target
```

Restarting all components:

```
systemctl start irp-shutdown.target  
systemctl start irp.target
```

Obtaining overall status of all components:

```
systemctl list-dependencies irp.target
```

 Individual statuses can be checked by executing command *systemctl status component_name*

Managing software components in OS with RC-files

Starting single component:

```
service explorer start
```

Stopping single component:

```
service explorer stop
```

Starting all components:

```
service irp start
```

Stopping all components:

```
service irp stop
```

Stopping all components except bgpd:

```
service irp stop nobgpd
```

Restarting all components:

```
service irp restart
```

Obtaining status of all components:

```
service irp status
```

2.5 IRP Initial Configuration

The platform initial configuration can be performed by using the Initial Setup wizard ([Initial Setup](#)) which can be accessed via the main platform IP address over HTTP or HTTPS protocol.

Example: `https://10.11.12.13`

2.6 Global and Core Configuration

The default values, specified for the Core service are sufficient for a proper system start-up. Several parameters can be adjusted during the initial system deployment. For a comprehensive list please see the [Global parameters](#) and [Core settings](#) sections.

During the initial setup and configuration stage, one must pay attention to the following configuration parameters:

- [global.nonintrusive_bgp](#) - must be set to “1” until the configuration and route propagation tests are completed.
- [global.improve_mode](#) - must be configured according to the specific network operator policies. See also: [IRP Optimization modes](#)
- [global.aggregate](#) - in most cases, it is recommended to enable the aggregates, in order to reduce the number of prefixes advertised by the IRP. Please consult the network infrastructure and configuration.
- [core.commit_control](#) - should be configured according to the specific network operator policies, see also: [Commit Control](#)
- [core.outage_detection](#) - in most cases it must be enabled. For more details see [Outage detection](#)

2.7 Collector Configuration

Depending on the preferred method of traffic data collection, one or both collector components should be configured. As specified in the [IRP Components](#) section, IRP can gather traffic data using the Flow (from now on: `irpflowd`) and Span collector (`irpspan`).


First of all, specific configuration to each collector will be described, along with the required router configuration changes.

2.7.1 Irpflowd Configuration

`Irpflowd` is a NetFlow/sFlow collector that receives and analyzes network traffic information, generated by your router(s).

NetFlow - an IP network statistics protocol developed by Cisco Systems, Inc. for collecting and transferring statistics regarding IP traffic information from network devices such as switches/routers to network analysis applications. `Irpflowd` currently supports the following **NetFlow** versions: v1, v5, v9.


sFlow - is a protocol designed for monitoring network, wireless and host devices. Developed by the [sFlow.org Consortium](#), it is supported by a wide range of network devices, as well as software routing and network solutions.

 Flow collector use is mandatory for Multiple Routing Domain networks since SPAN does not carry attributes to distinguish traffic between different providers.

2.7.1.1 Flow agents

Multi-router networks usually simultaneously carry traffic to a prefix over multiple providers. Flow collector needs to know the exact details of such a configuration in order to correctly determine the overall provider volume and active flows. Each provider configured in IRP can be explicitly set to match Flow statistics to specific Flow agents and help IRP Flow collector assign accurate statistics for each provider.

Flow agents have been added in order to support [Optimization for multiple Routing Domains](#) but when available they are used to enhance IRP capabilities in other areas too. For example a correct set of Flow agents for all providers enables IRP to accurately determine a prefix's current route. IRP components, especially Core during decision making can spot the latest prefix statistics matching a given Flow agent(s) and infer amount of traffic is sent over individual Providers. Collected data is later used to make Performance and Bandwidth decisions with knowing about multiple best routes.

 In case Flow agent data is missing IRP relies on past probing data to determine the current route of a prefix.

Flow agents are specified in the form of:

```
IPv4/interfaceID
```

where IPv4 is the source IP address of the packets coming from the agent and interfaceID is a numerical identifier of the interface in the range 1- 4294967295. The interface ID is usually its SNMP Interface ID on the router.

A collection of such values is assigned when multiple physical interfaces are used. For example:

```
peer.X.flow_agents = 8.8.8.8/1 8.8.8.8/2 8.8.8.8/3 8.8.8.8/4
```

The value is set via parameter `peer.X.flow_agents` or under Providers and Peers configuration in Frontend. The Frontend will also retrieve and suggest a list of available values that can be matched with the provider.

2.7.1.2 Configuration

To use the `irpflowd` collector, the following steps must be completed:

1. NetFlow/sFlow/jFlow must be configured on the router(s), which must send traffic flow information to the main IRP server IP (Figure 2.7.1). See (2.7.1.3) for specific network device configuration instructions




Figure 2.7.1: Flow export configuration

2. `irpflowd` must be enabled (by setting the `collector.flow.enabled` parameter):

```
collector.flow.enabled = 1
```

3. A list of all the networks, advertised by the edge routers that IRP will optimize, should be added to the configuration. This information should be specified in the `collector.ournets` parameter.
4. For security reasons, the list of valid Flow sending IP addresses must be configured in the `collector.flow.sources`, to protect `irpflowd` from unauthorized devices sending Flow data.

Example:

```
collector.flow.sources = 10.0.0.0/29
```

5. In case the Flow exporters are configured to use non-standard port numbers (2055 for NetFlow/jFlow and 6343 for sFlow), then `collector.flow.listen.nf` and `collector.flow.listen.sf` must be adjusted accordingly:

```
collector.flow.listen.nf = 2055
collector.flow.listen.sf = 6343
```

2.7.1.3 Vendor-specific NetFlow configuration examples

The following sections contain vendor-specific configuration examples, along with possible pitfalls and version-specific issues

NetFlow configuration on Cisco 7600/6500 series routers

 Refer also to the [Cisco NetFlow Software Configuration Guide](#)

Listing 2.1: Global MLS settings configuration

```
(config)# mls netflow
(config)# mls flow ip interface-full
(config)# mls flow ipv6 interface-full
(config)# mls sampling packet-based 512 8192
(config)# mls nde sender version 7
```

i Please replace the IP address and port with the actual IRP host IP and the collector UDP port (2055 by default)

Listing 2.2: Global NetFlow settings and export configuration

```
(config)# ip flow-cache entries 524288
(config)# ip flow-cache timeout inactive 60
(config)# ip flow-cache timeout active 1
(config)# ip flow-export version 9
(config)# ip flow-export destination 10.11.12.14 2055
```

i Ingress flow collection must be enabled on all interfaces facing the internal network. MLS NetFlow sampling must be enabled to preserve router resources.

Listing 2.3: Per-interface NetFlow settings configuration

```
(config)# int GigabitEthernet 3/6
(config-if)# mls netflow sampling
(config-if)# ip flow ingress
```

Flexible NetFlow configuration on Cisco 6500 series routers IOS 15.0SY series

Listing 2.4: Flexible NetFlow monitor configuration

```
(config)# flow monitor IRP-FLOW-MONITOR
(config-flow-monitor)# record platform-original ipv4 full
(config-flow-monitor)# exporter IRP-FLOW-EXPORTER
(config-flow-monitor)# cache timeout inactive 60
(config-flow-monitor)# cache timeout active 60
(config-flow-monitor)# cache entries 1048576
```

Listing 2.5: Flexible NetFlow exporter configuration

```
(config)# flow exporter IRP-FLOW-EXPORTER
(config-flow-exporter)# destination 10.11.12.14
(config-flow-exporter)# source Loopback0
(config-flow-exporter)# transport udp 2055
(config-flow-exporter)# template data timeout 120
```

i Please replace the IP address and port with the actual IRP host IP and the collector UDP port (2055 by default). Also replace the source interface with the actual one.


Listing 2.6: Flexible NetFlow sampler configuration


```
(config)# sampler flow-sampler
(config-sampler)# mode random 1 out-of 1024
```

Listing 2.7: Per-interface Flexible NetFlow settings configuration

```
(config)# interface FastEthernet0/0
(config-if)# ip flow monitor IRP-FLOW-MONITOR sampler flow-sampler input
(config-if)# ip flow monitor IRP-FLOW-MONITOR sampler flow-sampler output
```

NetFlow configuration on Cisco 7200/3600 series routers


 Please replace the IP address and port with the actual IRP host IP and the collector UDP port (2055 by default)

 Do not attempt to configure 7600/6500 series routers according to 7200/3600 router's configuration guide.

Listing 2.8: NetFlow configuration on Cisco 7200/3600 series routers

```
Router(config)# ip flow-cache entries 524288
Router(config)# ip flow-cache timeout inactive 60
Router(config)# ip flow-cache timeout active 1
Router(config)# ip flow-export version 9
Router(config)# ip flow-export destination 10.11.12.14 2055
```

Ingress/egress flow export configuration on peering interfaces


 According to [Cisco IOS NetFlow Command Reference](#) regarding the "ip flow" command history in IOS releases, this feature was introduced in IOS 12.3(11)T, 12.2(31)SB2, 12.2(18)SXE, 12.2(33)SRA.

NetFlow exporting must be configured on each peering interface which is used to send and/or receive traffic:

Listing 2.9: Ingress/egress flow export configuration on peering interfaces

```
Router(config)# interface FastEthernet 1/0
Router(config-if)# ip flow ingress
Router(config-if)# ip flow egress
```


Ingress flow export configuration (earlier IOS releases)

 Ingress flow export must be enabled on all interfaces facing the internal network. Prior to IOS 12.3(11)T, 12.2(31)SB2, 12.2(18)SXE, 12.2(33)SRA, flow export can be enabled only for ingress traffic, therefore it must be enabled on each interface that transmits/receives traffic from/to networks that must be improved

Listing 2.10: Ingress flow export configuration

```
Router(config)# interface FastEthernet 1/0
Router(config-if)# ip route-cache flow
```

NetFlow/sFlow configuration examples for Vyatta routers (VC 6.3)

 Do not configure both NetFlow and sFlow export for the same router interface. It will lead to traffic statistics distortion.

 Sampling interval must be set to at least 2000 for 10G links and 1000 for 1G links in order to save resources.

Listing 2.11: Configuring NetFlow export on Vyatta


```
vyatta@vyatta# set system flow-accounting netflow server 10.11.12.14 port
2055
vyatta@vyatta# set system flow-accounting netflow version 5
```

Listing 2.12: Configuration of an interface for the flow accounting

```
vyatta@vyatta# set system flow-accounting interface eth0
vyatta@vyatta# commit
```

jFlow export configuration for Juniper routers

Routing Engine-based sampling supports up to eight flow servers for both version 5 and version 8 configurations. The total number of collectors is limited to eight, regardless of how many are configured for version 5 or version 8. During the sampling configuration, the export packets are replicated to all the collectors, configured to receive them. If two collectors are configured to receive version 5 records, then both collectors will receive records for a specified flow.

 Default export interval for active/inactive flows on some Juniper routers is 1800 seconds. IRP requires significantly more frequent updates. The export interval recommended by IRP is 60 seconds. Refer your JunOS documentation on how to set export intervals. These parameters are named flow-active-timeout and flow-inactive-timeout.

Listing 2.13: Juniper flow export configuration

```
forwarding-options {
    sampling {
        input {
            family inet {
                rate 1000;
            }
        }
        family inet {
            output {
                flow-server 10.10.3.2 {
                    port 2055;
                    version 5;
                    source-address 10.255.255.1;
                }
            }
        }
    }
}
```

NetFlow export must be configured on all the interfaces facing the providers. In some cases, it may also be necessary to enable NetFlow forwarding on the interfaces facing the internal network.

Listing 2.14: Per-interface NetFlow sampling configuration

```
interfaces {
    xe-0/0/0 {
        unit 0 {
            family inet {
                sampling {
                    input output;
                }
            }
        }
    }
}
```

2.7.2 Irpspand Configuration

Irpspand acts like a passive network sniffer, analyzing the traffic that is provided to one or more dedicated network interfaces on the IRP server (defined in `collector.span.interfaces`) from a mirrored

port on your router or switch. Irspand looks up the IP header in the mirrored traffic. When the link level traffic is VLAN tagged as for example in IEEE 802.1Q or 802.1ad for Q-in-Q datagrams, Irspand will advance its IP packet sniffer past VLAN tags. For better results (higher performance and more analyzed traffic), as specified in the [IRP Technical Requirements](#) section, we recommend using Myricom 10Gbps NICs, with Sniffer10G license enabled.

⚠ Enable `collector.span.size_from_ip_header` configuration parameter if packets are stripped before forwarding to IRP's SPAN port.

To use the `irspand` collector, the following steps must be completed:

1. Configure port mirroring on your router or switch, as shown in figures (2.7.2) and (2.7.3).




Figure 2.7.2: Span port configuration (on the router)

or:




Figure 2.7.3: Span port configuration (on the switch)

2. Enable the span collector by setting the `collector.span.enabled` parameter in the configuration file:

```
collector.span.enabled = 1
```

3. Define the list of network interfaces that receive mirrored traffic, by setting the `collector.span.interfaces` parameter (multiple interfaces separated by space can be specified):

```
collector.span.interfaces = eth1 eth2 eth3
```

4. A list of all the networks advertised by the edge routers that IRP will optimize must be added to the configuration. This information should be specified in the `collector.ournets` parameter.

5. In case blackouts, congestions and excessive delays are to be analyzed by the system, the `collector.span.min_delay` must be turned on as well


```
collector.span.min_delay = 1
```

2.8 Explorer Configuration

As described in the [IRP Components](#) section, the explorer is the actual service that performs active remote network probing and tracing.

Explorer runs active and passive probes through all the available providers in order to determine the best one for the particular prefix. Such metrics as packet loss, latency, link capacity and link load are taken into consideration. In order to run a probe through a particular provider, Explorer needs the following to be configured:

1. An additional IP alias for each provider should be assigned and configured on the IRP server. This IP will be used as a source address during the probing process.

 It is recommended to configure reverse DNS records for each IP using the following template:
`performance-check-via-<PROVIDER-NAME> .`
`HARMLESS-NOCTION-IRP-PROBING.<YOUR-DOMAIN-NAME> .`

2. Policy-based routing (PBR) has to be configured on the edge router(s), so that traffic originating from each of these probing IP addresses will exit the network via specific provider. See specific PBR configuration in the [Specific PBR configuration scenarios](#) section.

 If network has Flowspec capabilities then alternatively Flowspec policies can be used instead of PBR. Refer for example [Flowspec policies](#), [global.flowspec.pbr](#).

1. Policy-based routing has to be configured to drop packets rather than routing them through the default route in case that the corresponding Next-Hop does not exist in the routing table.

2.8.1 Specific PBR configuration scenarios

PBRs can be setup in multiple ways, depending on the existing network infrastructure.

We will assume the following IP addresses/networks are used:

10.0.0.0/24 - used on the IRP server as well as the probing VLANs

10.0.0.2/32 - main IRP server IP address

10.0.0.3-10.0.0.5 - probing IP addresses

10.0.0.250-10.0.0.254 - router-side IP addresses for the probing VLANs

10.0.1.0/24 - used for GRE tunnel interfaces, if needed

10.10.0.0/24 - real edge routers IP addresses


10.11.0.0/30 - BGP session with the 1st provider, **10.11.0.1** being the ISP BGP neighbor IP

10.12.0.0/30 - BGP session with the 2nd provider, **10.12.0.1** being the ISP BGP neighbor IP

10.13.0.0/30 - BGP session with the 3rd provider, **10.13.0.1** being the ISP BGP neighbor IP

Vlan 3 - the probing Vlan

eth0 - the probing network interface on the IRP server

 In a production network, please change the IP addresses used in these examples to the actual addresses assigned and used in the network. Same goes for the Vlan interface.

i Brocade routers use Cisco compliant configuration commands. Unless otherwise noted, the Cisco configuration examples will also work on Brocade devices.

Case 1: Single router, two providers, and separate probing Vlan.

i 10.0.0.1/32 is configured on the edge router, on the probing vlan interface (ve3).




Figure 2.8.1: PBR configuration: single router and separate probing Vlan

In this case, a single PBR rule should be enforced on the Edge router for each of the probing IP addresses.

Listing 2.15: Cisco IPv4 PBR configuration: 1 router, 2 providers

```
access-list 1 permit ip host 10.0.0.3
access-list 2 permit ip host 10.0.0.4
!
route-map irp-peer permit 10
match ip address 1
set ip next-hop 10.11.0.1
set interface Null0
!
route-map irp-peer permit 20
match ip address 2
set ip next-hop 10.12.0.1
set interface Null0
!
interface ve 3
ip policy route-map irp-peer
```

- Route-map entries with the destination set to Null0 interface are used for preventing packets to flow through the default route in the case that the corresponding next-hop does not exist in the routing table (typically when a physical interface administrative/operational status is down).

Cisco ASR9000 routers running IOS XR use a different PBR syntax.

Listing 2.16: Cisco IPv4 PBR configuration for IOS XR: 1 router, 2 providers

```
configure
ipv4 access-list irp-peer
  10 permit ipv4 host 10.0.0.3 any nexthop1 ipv4 10.11.0.1 nexthop2 ipv4
    169.254.0.254
  11 permit ipv4 host 10.0.0.4 any nexthop1 ipv4 10.12.0.1 nexthop2 ipv4
    169.254.0.254
end
```

```

router static
  address-family ipv4 unicast
  169.254.0.254 Null0
end

interface FastEthernet1/1
  ipv4 access-group irp-peer ingress
end

```

For Juniper routers, a more complex configuration is required:

Listing 2.17: Juniper IPv4 PBR configuration: 1 router, 2 providers

```

[edit interfaces]
xe-0/0/0 {
  unit 3 {
    family inet {
      filter {
        input IRP-policy;
      }
    }
  }
}
[edit firewall]
family inet {
  filter IRP-policy {
    term irp-peer1 {
      from {
        source-address 10.0.0.3/32;
      }
      then {
        routing-instance irp-ispl-route;
      }
    }
    term irp-peer2 {
      from {
        source-address 10.0.0.4/32;
      }
      then {
        routing-instance irp-isp2-route;
      }
    }
    term default {
      then {
        accept;
      }
    }
  }
}
[edit]
routing-instances {
  irp-ispl-route {
    instance-type forwarding;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop 10.11.0.1;
      }
    }
  }
}

```

```

    irp-isp2-route {
        instance-type forwarding;
        routing-options {
            static {
                route 0.0.0.0/0 next-hop 10.12.0.1;
            }
        }
    }
}
routing-options {
    interface-routes {
        rib-group inet irp-policies;
    }
    rib-groups {
        irp-policies {
            import-rib [ inet.0 irp-isp1-route.inet.0 irp-isp2-route.inet.0 ];
        }
    }
}

```

PBR configuration on Vyatta routers.

Unfortunately, prior to and including VC6.4, Vyatta does not natively support policy-based routing. Thus, the PBR rules should be configured using the standard Linux `ip` toolset. To make these rules persistent, they should be also added to `/etc/rc.local` on the Vyatta system.

Listing 2.18: Vyatta IPv4 PBR configuration example

```

ip route add default via 10.11.0.1 table 101
ip route add default via 10.12.0.1 table 102
ip rule add from 10.0.0.3 table 101 pref 32001
ip rule add from 10.0.0.4 table 102 pref 32002

```

Vyatta versions VC6.5 and up, natively support source-based routing. The following example can be used:

Listing 2.19: Vyatta (VC6.5 and up) IPv4 PBR configuration example

```

# Setup the routing policy:
set policy route IRP-ROUTE
set policy route IRP-ROUTE rule 10 destination address 0.0.0.0/0
set policy route IRP-ROUTE rule 10 source address 10.0.0.3/32
set policy route IRP-ROUTE rule 10 set table 103
set policy route IRP-ROUTE rule 20 destination address 0.0.0.0/0
set policy route IRP-ROUTE rule 20 source address 10.0.0.4/32
set policy route IRP-ROUTE rule 20 set table 104
set policy route IRP-ROUTE rule 30 destination address 0.0.0.0/0
set policy route IRP-ROUTE rule 30 source address 0.0.0.0/0
set policy route IRP-ROUTE rule 30 set table main
commit

# Create static route tables:
set protocols static table 103 route 0.0.0.0/0 nexthop 10.11.0.1
set protocols static table 104 route 0.0.0.0/0 nexthop 10.12.0.1
commit

# Assign policies to specific interfaces, Vlan 3 on eth1 in this example:
set interfaces ethernet eth1.3 policy route IRP-ROUTE

# Verify the configuration:
show policy route IRP-ROUTE

```

```
show protocols static
show interfaces ethernet eth1.3
```

Case 2: Two edge routers, two providers, and a separate probing Vlan.

i Following IP addresses are configured on the routers:

- **10.0.0.251** is configured on **R1, VE3**
- **10.0.0.252** is configured on **R2, VE3**




Figure 2.8.2: PBR configuration: two routers and separate probing Vlan

To reduce the number of PBR rules to one per each router, additional source based routing rules must be configured on the IRP server. This can be done either by using the `ip` command, or by adjusting the standardized `/etc/sysconfig/network-scripts/route-eth*` and `/etc/sysconfig/network-scripts/rule-eth*` configuration files. Both ways will be presented.

Listing 2.20: IRP server source-based routing using the ‘ip’ command

```
ip route add default via 10.0.0.251 table 201
ip route add default via 10.0.0.252 table 202
ip rule add from 10.0.0.3 table 201 pref 32101
ip rule add from 10.0.0.4 table 202 pref 32102
```

Listing 2.21: IRP server source-based routing using standard CentOS configuration files

```
#!/etc/sysconfig/network-scripts/route-eth0:
default via 10.0.0.251 table 201
default via 10.0.0.252 table 202
#!/etc/sysconfig/network-scripts/rule-eth0:
from 10.0.0.3 table 201 pref 32101
from 10.0.0.4 table 202 pref 32102
```

i Refer to OS configuration manual for configuration guidelines.

Router configuration looks similar to the previous case. A Cisco/Brocade example will be provided.

! Some Brocade routers/switches have PBR configuration limitations. Please refer to the “Policy-Based Routing” → “Configuration considerations” section in the Brocade documentation for your router/switch model.

For example, BigIron RX Series of switches do not support more than 6 instances of a route map, more than 6 ACLs in a matching policy of each route map instance, and more than 6 next hops in a set policy of each route map instance.

On the other hand, some Brocade CER/CES routers/switches have these limits raised up to 200 instances (depending on package version).

Listing 2.22: Cisco IPv4 PBR configuration: 2 routers, 2 providers, separate Vlan

```

#Router R1
access-list 1 permit ip host 10.0.0.3
!
route-map irp-peer permit 10
match ip address 1
  set ip next-hop 10.11.0.1
  set interface Null0
!
interface ve 3
ip policy route-map irp-peer

#Router R2
access-list 1 permit ip host 10.0.0.4
!
route-map irp-peer permit 10
match ip address 1
  set ip next-hop 10.12.0.1
  set interface Null0
!
interface ve 3
ip policy route-map irp-peer

```


Case 3: Complex network infrastructure, multiple routers, no probing VLAN

Figure 2.8.3: PBR configuration via GRE tunnels

In specific complex scenarios, traffic from the IRP server should pass multiple routers before getting to the provider. If a separate probing Vlan cannot be configured across all routers, GRE tunnels from IRP to the Edge routers should be configured.

i It is also possible to configure the PBRs without GRE tunnels, by setting PBR rules on each transit router on the IRP↔ provider path.

w Brocade routers do not support PBR set up on GRE tunnel interfaces. In this case the workaround is to configure PBR on each transit interface towards the exit edge router(s) interface(s).

GRE Tunnels configuration

Listing 2.23: IRP-side IPv4 GRE tunnel CLI configuration


```

modprobe ip_gre
ip tunnel add tun0 mode gre remote 10.10.0.1 local 10.0.0.2 ttl 64 dev eth0
ip addr add dev tun0 10.0.1.2/32 peer 10.0.1.1/32
ip link set dev tun0 up

```

Listing 2.24: IRP-side IPv4 GRE tunnel configuration using standard CentOS configs

```
#/etc/sysconfig/network-scripts/ifcfg-tun0
DEVICE=tun0
TYPE=GRE
ONBOOT=yes
MY_INNER_IPADDR=10.0.1.2
MY_OUTER_IPADDR=10.0.0.2
PEER_INNER_IPADDR=10.0.1.1
PEER_OUTER_IPADDR=10.10.0.1
TTL=64
```

 Refer to OS configuration manual for configuration guidelines.

Listing 2.25: Router-side IPv4 GRE tunnel configuration (Vyatta equipment)

```
set interfaces tunnel tun0
set interfaces tunnel tun0 address 10.0.1.1/30
set interfaces tunnel tun0 description "IRP_Tunnel_1"
set interfaces tunnel tun0 encapsulation gre
set interfaces tunnel tun0 local-ip 10.10.0.1
set interfaces tunnel tun0 remote-ip 10.0.0.2
```

Listing 2.26: Router-side IPv4 GRE tunnel configuration (Cisco equipment)

```
interface Tunnel0 routers
ip address 10.0.1.1 255.255.255.252
tunnel mode gre ip
tunnel source Loopback1
tunnel destination 10.0.0.2
```

Listing 2.27: Router-side IPv4 GRE tunnel configuration (Juniper equipment)

```
interfaces {
    gr-0/0/0 {
        unit 0 {
            tunnel {
                source 10.0.0.2;
                destination 10.10.0.1;
            }
            family inet {
                address 10.0.1.1/32;
            }
        }
    }
}
```

The above configuration is for the first edge router (R1). One more GRE tunnel should be configured on the 2nd router (R2).

As soon as the GRE tunnels are configured, the source-based routing and the PBR rules should be configured, similar to the previous section.


Listing 2.28: IRP server IPv4 source-based routing via GRE using the ‘ip’ command

```
ip route add default dev tun0 table 201
ip route add default dev tun1 table 202
ip route add default dev tun2 table 203
ip rule add from 10.0.1.2 table 201 pref 32101
```

```
ip rule add from 10.0.1.6 table 202 pref 32102
ip rule add from 10.0.1.10 table 202 pref 32103
```

Listing 2.29: IRP server IPv4 source-based routing via GRE using standard CentOS configuration files

```
#/etc/sysconfig/network-scripts/route-tun0:
default dev tun0 table 201
default dev tun1 table 202
default dev tun2 table 203
#/etc/sysconfig/network-scripts/rule-tun0:
from 10.0.1.2 table 201 pref 32101
from 10.0.1.6 table 202 pref 32102
from 10.0.1.10 table 203 pref 32103
```

 Refer to OS configuration manual for configuration guidelines.

Router configuration looks similar to the previous cases. A Cisco/Brocade example will be provided.

Listing 2.30: Cisco IPv4 PBR configuration: 2 routers, 2 providers, separate Vlan

```
#Router R1
access-list 1 permit ip host 10.0.1.2
access-list 2 permit ip host 10.0.1.6
!
route-map irp-peer permit 10
 match ip address 1
  set ip next-hop 10.11.0.1
  set interface Null0
!
route-map irp-peer permit 20
 match ip address 2
  set ip next-hop 10.12.0.1
  set interface Null0
!

interface Tunnel0
ip policy route-map irp-peer
interface Tunnell
ip policy route-map irp-peer

#Router R2
access-list 1 permit ip host 10.0.1.10
!
route-map irp-peer permit 10
 match ip address 1
  set ip next-hop 10.13.0.1
  set interface Null0
!
interface Tunnel0
ip policy route-map irp-peer
```

Case 4: Internet Exchanges configuration examples

The PBR rules for Cisco routers/switches are generated by IRP, following the template below.

Listing 2.31: Cisco IPv4 PBR configuration template

```
!--- repeated block for each peering partner
no route-map <ROUTEMAP> permit <ACL>
```

```

no ip access-list extended <ROUTEMAP>--<ACL>

ip access-list extended <ROUTEMAP>--<ACL>
  permit ip host <PROBING_IP> any dscp <PROBING_DSCP>

route-map <ROUTEMAP> permit <ACL>
  match ip address <ROUTEMAP>--<ACL>
  set ip next-hop <NEXT_HOP>
  set interface Null0

!--- block at the end of PBR file
interface <INTERFACE>
  ip policy route-map <ROUTEMAP>

```

The “<>” elements represent variables with the following meaning:

- <ROUTEMAP> represents the name assigned by IRP and equals the value of the Route Map parameter in PBR Generator ("irp-ix" in Figure 4)
- <ACL> represents a counter that identifies individual ACL rules. This variable's initial value is taken from ACL name start field of PBR Generator and is subsequently incremented for each ACL
- <PROBING_IP> one of the configured probing IPs that IRP uses to probe link characteristics via different peering partners. One probing IP is sufficient to cover up to 64 peering partners
- <PROBING_DSCP> an incremented DSCP value assigned by IRP for probing a specific peering partner. This is used in combination with the probing IP
- <NEXT_HOP> represents the IP address identifying the peering partner on the exchange. This parameter is retrieved during autoconfiguration and preserved in Exchange configuration
- <INTERFACE> represents the interface where traffic conforming to the rule will exist the Exchange router. This is populated with the Interface value of PBR Generator

The PBR rules for Brocade non-XMR router/switches are generated by IRP, following the template below.

Listing 2.32: Cisco IPv4 PBR configuration template

```

!--- repeated block for each peering partner
no route-map <ROUTEMAP> permit <ACL>
no ip access-list extended <ROUTEMAP>--<ACL>


ip access-list extended <ROUTEMAP>--<ACL>
  permit ip host <PROBING_IP> any dscp-matching <PROBING_DSCP>

route-map <ROUTEMAP> permit <ACL>
  match ip address <ROUTEMAP>--<ACL>
  set ip next-hop <NEXT_HOP>
  set interface Null0

!--- block at the end of PBR file
interface <INTERFACE>
  ip policy route-map <ROUTEMAP>

```

The “<>” elements represent variables with the same meaning as per Cisco example.

 Brocade XMR routers use keyword “dscp-mapping” instead of “dscp-matching”.

The PBR rules for Juniper routers/switches are generated by IRP, following the template below.

i It is important to note that the different sections of the PBR rules (load replace/merge relative terminal) should be entered independently and not as a single file that is output by IRP. Also take note that the last group includes a 'load merge' combo and not a 'load replace' as the first three groups.

Listing 2.33: Juniper IPv4 PBR configuration template

```
load replace relative terminal
[Type ^D at a new line to end input]

interfaces {
  <INTERFACE> {
    unit <INTERFACE_UNIT> {
      family inet {
        filter {
          replace:
            input <ROUTEMAP>;
        }
      }
    }
  }
}

load replace relative terminal
[Type ^D at a new line to end input]

firewall {
  family inet {
    filter <ROUTEMAP> {
      replace:
        term <ROUTEMAP><ACL> {
          from {
            source-address <PROBING_IP>;
            dscp <PROBING_DSCP>;
          }
          then {
            routing-instance <ROUTEMAP><ACL>-route;
          }
        }
      ...
    }
    replace:
    term default {
      then {
        accept;
      }
    }
  }
}

load replace relative terminal
[Type ^D at a new line to end input]

routing-instances {
  replace:
```

```

    <ROUТЕMAP><ACL>-route {
        instance-type forwarding;
        routing-options {
            static {
                route 0.0.0.0/0 next-hop <NEXT_HOP>;
            }
        }
    }
...
}

load merge relative terminal
[Type ^D at a new line to end input]

routing-options {
    interface-routes {
        replace:
        rib-group inet <ROUТЕMAP>rib;
    }
    rib-groups {
        replace:
        <ROUТЕMAP>rib {
            import-rib [ inet.0 <ROUТЕMAP><ACL>-route.inet.0 ... ];
        }
    }
}

```

The “<>” elements represent variables with the following meaning:

- <INTERFACE> represents the interface where traffic conforming to the rule will exist the Exchange router. This is populated with the Interface value of PBR Generator
- <INTERFACE_UNIT> is the value of the Interface Unit parameter in PBR Generator
- <ROUТЕMAP> represents the name assigned by IRP and equals the value of the Route Map parameter in PBR Generator
- <ACL> represents a combined counter like "00009" that identifies individual ACL rules. This variable's initial value is taken from ACL name start field of PBR Generator and is subsequently incremented for each ACL
- <PROBING_IP> one of the configured probing IPs that IRP uses to probe link characteristics via different peering partners. One probing IP is sufficient to cover up to 64 peering partners
- <PROBING_DSCP> an incremented DSCP value assigned by IRP for probing a specific peering partner. This is used in combination with the probing IP
- <NEXT_HOP> represents the IP address identifying the peering partner on the exchange. This parameter is retrieved during autoconfiguration and preserved in Exchange configuration

i Note that Juniper routers/switches need an additional parameter in order to correctly configure PBRs - Interface unit.

Verifying PBR configuration

To ensure that the PBRs are properly configured on the router(s), the standard *NIX traceroute command can be used, by specifying each probing IP as the source IP address for tracing the route. The route should pass through a different provider (note the ISP BGP neighbor IP).

Listing 2.34: PBR validation using ‘traceroute’

```

root@server ~ $ traceroute -m 5 8.8.8.8 -nns 10.0.0.3
traceroute to 8.8.8.8 (8.8.8.8), 30 hops max, 60 byte packets
 1 10.0.0.1 0.696 ms 0.716 ms 0.783 ms
 2 10.11.0.1 0.689 ms 0.695 ms 0.714 ms
 3 84.116.132.146 14.384 ms 13.882 ms 13.891 ms
 4 72.14.219.9 13.926 ms 14.477 ms 14.473 ms
 5 209.85.240.64 14.397 ms 13.989 ms 14.462 ms

root@server ~ $ traceroute -m 5 8.8.8.8 -nns 10.0.0.4
traceroute to 8.8.8.8 (8.8.8.8), 30 hops max, 60 byte packets
 1 10.0.0.1 0.696 ms 0.516 ms 0.723 ms
 2 10.12.0.1 0.619 ms 0.625 ms 0.864 ms
 3 83.16.126.26 13.324 ms 13.812 ms 13.983 ms
 4 72.14.219.9 15.262 ms 15.347 ms 15.431 ms
 5 209.85.240.64 16.371 ms 16.991 ms 16.162 ms

```

Another useful method for checking that the PBRs are properly configured is to use Explorer self check option. Make sure that the providers are added to IRP configuration before executing Explorer self check.

Listing 2.35: PBR validation using Explorer self check

```

root@server ~ $ /usr/sbin/explorer -s
Starting PBR check

PBR check failed for provider A[2]. Diagnostic hop information: IP
=10.11.0.12 TTL=3
PBR check succeeded for provider B[3]. Diagnostic hop information: IP
=10.12.0.1 TTL=3

```

In order to ensure that the PBRs are properly configured for Internet Exchanges, the method below can be used.

Listing 2.36: PBR validation using ‘iptables’ and ‘traceroute’

```

root@server ~ $ iptables -t mangle -I OUTPUT -d 8.8.8.8 -j DSCP --set-dscp
<PROBING_DSCP>


root@server ~ $ traceroute -m 5 -nns <PROBING_IP> 8.8.8.8
traceroute to 8.8.8.8, 30 hops max, 60 byte packets
 1 ...
 2 ...
 3 <NEXT_HOP> 126.475 ms !X^C

```

where

- <NEXT_HOP> is a Peering Partner’s next-hop IP address in IRP configuration
- <PROBING_DSCP> is a Peering Partner’s DSCP value in IRP configuration
- <PROBING_IP> is a Peering Partner’s probing IP address in IRP configuration

The first IP of the trace leaving client infrastructure should be on the Exchange and the next-hop should belong to the correct Peering Partner.

 iptables rules should be deleted after all tests are done.

2.8.1.1 Current route detection

Current route detection algorithm requires the `explorer.infra_ips` parameter to be configured. All the IP addresses and subnets belonging to the infrastructure should be added here. These IP addresses can be easily determined by using the `tracert` command.

2.8.1.2 Providers configuration

Before Explorer starts probing the prefixes detected by the Collector, all providers should be configured.

i For the complete list of provider settings please see the [Provider](#) section.
Before configuring providers in IRP, the BGP sessions need to be defined, see [Bgp Configuration](#)

- Please note that some Brocade router models do not support SNMP counters per Vlan, therefore physical interfaces should be used for gathering traffic statistics.

For the sample configuration below, let's assume the following:




Figure 2.8.4: PBR configuration: single router and separate probing Vlan

ISP1 - the provider's name

10.0.0.1 - Router IP configured on the probing Vlan

10.0.0.3 - Probing IP for ISP1, configured on the IRP server

10.11.0.1, 10.11.0.2 - IP addresses used for the EBGp session with the ISP, 10.11.0.2 being configured on the router

400Mbps - the agreed bandwidth

1Gbps - the physical interface throughput

'public' - read-only SNMP community configured on R1

GigabitEthernet2/1 - the physical interface that connects R1 to ISP1

As presented in Listing 2.37, all the parameters are pretty self-explanatory. Please check the [BGP Monitoring](#) section as well.

Listing 2.37: Sample provider configuration

```
peer.1.95th = 400
peer.1.95th.bill_day = 1
peer.1.bgp_peer = R1
peer.1.cost = 6
peer.1.description = ISP1
peer.1.ipv4.next_hop = 10.11.0.1
peer.1.ipv4.probing = 10.0.0.3
```



```

peer.1.ipv4.diag_hop = 10.11.0.1
peer.1.ipv4.mon = 10.11.0.1 10.11.0.2
peer.1.limit_load = 1000
peer.1.shortname = ISP1
peer.1.snmp.interfaces = 1:GigabitEthernet2/1
peer.1.mon.ipv4.bgp_peer = 10.11.0.1

snmp.1.name = Host1
snmp.1.ip = 10.0.0.1
snmp.1.community = public

```

SNMP parameters validation

To make sure that the SNMP parameters are correct, the ‘snmpwalk’ tool can be run on the IRP server:

Listing 2.38: SNMP parameters validation

```

root@server ~ $ snmpwalk -v2c -c irp-public 10.0.0.1 ifDescr
IF-MIB::ifDescr.1 = STRING: GigabitEthernet1/1
IF-MIB::ifDescr.2 = STRING: GigabitEthernet1/2
IF-MIB::ifDescr.3 = STRING: GigabitEthernet2/1
IF-MIB::ifDescr.4 = STRING: GigabitEthernet2/2
IF-MIB::ifDescr.5 = STRING: GigabitEthernet2/3
IF-MIB::ifDescr.6 = STRING: GigabitEthernet2/4

root@server ~ $ snmpwalk -v2c -c irp-public 10.0.0.1 ifIndex
IF-MIB::ifIndex.1 = INTEGER: 1
IF-MIB::ifIndex.2 = INTEGER: 2
IF-MIB::ifIndex.3 = INTEGER: 3
IF-MIB::ifIndex.4 = INTEGER: 4
IF-MIB::ifIndex.5 = INTEGER: 5
IF-MIB::ifIndex.6 = INTEGER: 6


```

2.8.2 Flowspec PBR

In networks that have Flowspec and Redirect to IP capabilities, PBR can be implemented by means of Flowspec policies. Refer details in [Flowspec policies](#).

In order to use assigned providers or peers on Internet Exchanges in IRP the same set of source IP addresses and DSCP values (for Internet Exchanges) are assigned to individual providers/peers. These values are used to automatically generate Flowspec policies that redirect IRP probes to designated next-hops.

In order to use this feature `global.flowspec`, `global.flowspec.pbr` and `bgpd.peer.X.flowspec` must be enabled.

 Flowspec PBR cannot be used during non-intrusive (`global.nonintrusive_bgp`) mode.


2.9 Bgpd Configuration

For IRP to inject the improved prefixes to the routing tables, proper iBGP sessions have to be configured between the edge routers and the IRP. The IRP BGP daemon acts similarly to a *route-reflector-client*. Following criteria need to be met:


1. An internal BGP session using the same autonomous system number (ASN) must be configured between each edge router and the IRP. BGP sessions must not be configured with *next-hop-self* (route reflectors can't be used to inject routes with modified *next-hop*) - the *next-hop* parameter advertised by IRP Bgpd should be distributed to other iBGP neighbors.
2. *route-reflector-client* must be enabled for the routes advertised by IRP Bgpd to be distributed to all non-client neighbors.
3. Routes advertised by IRP Bgpd must have a higher preference over routes received from external BGP neighbors.

This can be done by different means, on the IRP or on the router side:

- *Local-pref* can be set to a reasonably high value in the Bgpd configuration
- *Communities* can be appended to prefixes advertised by Bgpd


 Avoid collisions of localpref or communities values assigned to IRP within both its configuration and/or on customer's network.


- Multi-exit-discriminator (*MED*) can be changed to affect the best-path selection algorithm
- *Origin* of the advertised route can be left unchanged or overridden to a specific value (*incomplete*, *IGP*, *EGP*)

 LocalPref, MED and Origin attribute values are set with the first nonempty value in this order: 1) value from configuration or 2) value taken from incoming aggregate or 3) default value specified in RFC4271.

Communities attribute value concatenates the value taken from incoming aggregate with configuration value. The router should be configured to send no Communities attribute in case it is required that IRP announces Communities attribute that contain only the configured value.

4. BGP *next-hop* must be configured for each provider configured in IRP (please refer to [Providers configuration](#) and [Provider](#))

 Special attention is required if edge routers are configured to announce prefixes with empty AS-Path. In some cases, improvements announced by Bgpd may have an empty AS-Path. Thus, when edge router does not have any prefix-list filtering enforced, all the current improvements can be improperly advertised to routers - this may lead to policy violations and session reset. Refer to [AS-Path behavior in IRP Bgpd](#) regarding the way to disallow announcements with empty AS-Path. None of the improvements advertised by IRP should be advertised to your external peers (refer to [bgpd.no_export](#)).

 We recommend the routes to be injected into the edge router which runs the BGP session with the provider. This ensures that the routes are properly redistributed across the network.

For example, there are two routers: R1 and R2. R1 runs a BGP session with Level3 and R2 with Cogent. The current route is x.x.x.x/24 with the next-hop set to Level3 and all the routers learn this route via R1. The system injects the x.x.x.x/24 route to R2 with the next-hop updated to Cogent.

In this case the new route is installed on the routing table and it will be properly propagated to R1 (and other routers) via iBGP.

However if the system injects the new route to R1 instead of R2, the route's next-hop will point to R2 while R2 will have the next-hop pointing to R1 as long as the injected route is propagated over iBGP to other routers. In this case a routing loop will occur.

In the following BGP session configuration example, the IRP server with IP 10.0.0.2 establishes an iBGP session to the edge router (IP: 10.0.0.1). The *local-pref* parameter for the prefixes advertised by IRP is set to 190. BGP monitoring (see [BGP Monitoring](#), [Bgpd settings](#)) is enabled.

Listing 2.39: iBGP session configuration example

```
bgpd.peer.R1.as           = 65501
bgpd.peer.R1.our_ip       = 10.0.0.2
bgpd.peer.R1.master_peer_ip = 10.0.0.1
bgpd.peer.R1.listen      = 1
bgpd.peer.R1.localpref    = 190
bgpd.peer.R1.shutdown    = 0
```

Vendor-specific router-side iBGP session configuration examples:

Vyatta routers:

Listing 2.40: Vyatta IPv4 iBGP session configuration example

```
set protocols bgp 65501 neighbor 10.0.0.2 remote-as '65501'
set protocols bgp 65501 neighbor 10.0.0.2 route-reflector-client
set protocols bgp 65501 parameters router-id '10.0.0.1'
```

Listing 2.41: Vyatta IPv6 iBGP session configuration example

```
delete system ipv6 disable-forwarding
commit
set protocols bgp 65501 neighbor 2001:db8:2::2 remote-as '65501'
set protocols bgp 65501 neighbor 2001:db8:2::2 route-reflector-client
set protocols bgp 65501 neighbor 2001:db8:2::2 address-family 'ipv6-unicast'
set protocols bgp 65501 parameters router-id '10.0.0.1'
```

Listing 2.42: Vyatta IPv4 route-map for setting local-pref on the router

```
set protocols bgp 65501 neighbor 10.0.0.2 route-map import 'RM-IRP-IN'
set policy route-map RM-IRP-IN rule 10 action 'permit'
set policy route-map RM-IRP-IN rule 10 set local-preference '190'
```

Listing 2.43: Vyatta IPv6 route-map for setting local-pref on the router

```
set protocols bgp 65501 neighbor 2001:db8:2::2 route-map import 'RM-IRP-IN'
set policy route-map RM-IRP-IN rule 10 action 'permit'
set policy route-map RM-IRP-IN rule 10 set local-preference '190'
```

Cisco routers:

Listing 2.44: Cisco IPv4 iBGP session configuration example

```
router bgp 65501
neighbor 10.0.0.2 remote-as 65501
neighbor 10.0.0.2 send-community
neighbor 10.0.0.2 route-reflector-client
```

Listing 2.45: Cisco IPv6 iBGP session configuration example

```
router bgp 65501
neighbor 2001:db8:2::2 remote-as 65501
neighbor 2001:db8:2::2 send-community
neighbor 2001:db8:2::2 route-reflector-client
```

or

```
router bgp 65501
neighbor 2001:db8:2::2 remote-as 65501
no neighbor 2001:db8:2::2 activate
address-family ipv6
neighbor 2001:db8:2::2 activate
neighbor 2001:db8:2::2 send-community
neighbor 2001:db8:2::2 route-reflector-client
```

Listing 2.46: Cisco IPv4 route-map for setting local-pref on the router

```
router bgp 65501
neighbor 10.0.0.2 route-map RM-IRP-IN input
route-map RM-IRP-IN permit 10
  set local-preference 190
```

Listing 2.47: Cisco IPv6 route-map for setting local-pref on the router

```
router bgp 65501
neighbor 2001:db8:2::2 route-map RM-IRP-IN input
route-map RM-IRP-IN permit 10
  set local-preference 190
```


Listing 2.48: Limiting the number of received prefixes for an IPv4 neighbor on Cisco

```
router bgp 65501
neighbor 10.0.0.2 maximum-prefix 10000
```

Listing 2.49: Limiting the number of received prefixes for an IPv6 neighbor on Cisco

```
router bgp 65501
neighbor 2001:db8:2::2 maximum-prefix 10000
```

Juniper equipment:

 The Cluster ID must be unique in multi-router configuration. Otherwise improvements will not be redistributed properly. Cluster ID is optional for single router networks.

Listing 2.50: Juniper IPv4 iBGP session configuration example

```
[edit]
routing-options {
  autonomous-system 65501;
  router-id 10.0.0.1;
}
protocols {
  bgp {
    group 65501 {
      type internal;
      cluster 0.0.0.1;
```

```

        family inet {
            unicast;
        }
        peer-as 65501;
        neighbor 10.0.0.2;
    }
}

```

Listing 2.51: Juniper IPv6 iBGP session configuration example

```

[edit]
routing-options {
    autonomous-system 65501;
    router-id 10.0.0.1;
}
protocols {
    bgp {
        group 65501 {
            type internal;
            cluster 0.0.0.1;
            family inet6 {
                any;
            }
            peer-as 65501;
            neighbor 2001:db8:2::2;
        }
    }
}

```

Listing 2.52: Juniper IPv4 route-map for setting local-pref on the router

```

[edit]
routing-options {
    autonomous-system 65501;
    router-id 10.0.0.1;
}
protocols {
    bgp {
        group 65501 {
            type internal;
            peer-as 65501;
            neighbor 10.0.0.2 {
                preference 190;
            }
        }
    }
}

```

Listing 2.53: Limiting the number of received prefixes for an IPv4 neighbor on Juniper

```

protocols {
    bgp {
        group 65501 {
            neighbor 10.0.0.2 {
                family inet {
                    any {
                        prefix-limit {
                            maximum 10000;
                        }
                    }
                }
            }
        }
    }
}

```


2.9.3 BGP Additional paths

Typically router sends only the best routes and that behaviour hides alternatives.

Enabling sending of additional paths ([RFC 7911](#)) allows IRP to automatically configure complete list of routes for partial routing providers and internet exchanges peering partners.

Listing 2.55: Cisco: Configure BGP Additional Paths

```
router bgp 65500
neighbor 10.1.1.2 remote-as 65500
!
address-family ipv4
  bgp additional-paths select all
  neighbor 10.1.1.2 activate
  neighbor 10.1.1.2 additional-paths send
  neighbor 10.1.1.2 advertise additional-paths all
exit-address-family
```


Listing 2.56: Juniper: Configure BGP Additional Paths

```
[edit protocols bgp group group-name family family]
add-path {
  send {
    path-count 20;
  }
}
```

2.10 BMP Configuration


IRP's BMP monitoring station passively listens on a designated TCP port for monitoring routers to establish BMP sessions. Further BMP setup is performed on monitoring router where

- BMP monitoring station's IP address and port are set and also
- filtering rules regarding what BMP data is sent to IRP are applied if needed.

 By default BMP implementations do not filter routing data sent to a monitoring station (refer to [peer.X.bmp.check_routes](#)).

BMP related features are configured individually for example:

- BMP monitoring station: [BMP monitoring station settings](#).
- primary source of data for current route re-construction: [bgpd.as_path](#).

 It is recommended that BMP is set as the primary source for current route reconstruction only when BMP data is available for all providers.

- improvement old and new provider re-probing on AS Path changes: [bgpd.retry_probing.new.bmp_path_change](#), [bgpd.retry_probing.old.bmp_path_change](#).

2.11 Threat Mitigation Configuration

2.11.1 BGP Blackholing

BGP Blackholing requires the following parameters to be configured:

- Blackholing next-hop `bgpd.peer.X.blackholing.ipv4.next_hop` & `bgpd.peer.X.blackholing.ipv6.next_hop`
- Each eligible provider should have a configured Blackholing BGP session `peer.X.blackholing.bgp_peer` and Blackholing community `peer.X.blackholing.community` (formerly offered by a provider)
- Default BGP reaction (`irpdetectd.bgp.reaction`) should be set to 0 (Drop)
- Moderated/Automated Mode (`irpdetectd.mode`) requires thresholds to be set to operate properly (`irpdetectd.blackhole.threshold.kpps` and/or `irpdetectd.blackhole.threshold.mbps`).

2.11.2 BGP Redirect

BGP Redirect requires the following parameters to be configured:

- Blackholing next-hop `bgpd.peer.X.blackholing.ipv4.next_hop` & `bgpd.peer.X.blackholing.ipv6.next_hop`
- BGP redirect routers (`irpdetectd.bgp.redirect.bgp_peers`) and BGP redirect communities (`irpdetectd.bgp.redirect.communities`)
- Default BGP reaction (`irpdetectd.bgp.reaction`) should be set to 1 (Redirect)
- Moderated/Automated Mode (`irpdetectd.mode`) requires thresholds to be set to operate properly (`irpdetectd.blackhole.threshold.kpps` and/or `irpdetectd.blackhole.threshold.mbps`).

2.11.3 FlowSpec Drop

FlowSpec Drop requires the following parameters to be configured:

- Configure FlowSpec globally and enable on each required BGP session (`bgpd.peer.X.flowspec`)
- Default FlowSpec reaction (`irpdetectd.flowspec.reaction`) should be set to 0 (Drop)
- Moderated/Automated Mode (`irpdetectd.mode`) requires thresholds to be set to operate properly (`irpdetectd.flowspec.threshold.kpps` and/or `irpdetectd.flowspec.threshold.mbps`).

2.11.4 FlowSpec Redirect

FlowSpec Redirect requires the following parameters to be configured:

- Configure FlowSpec globally and enable on each required BGP session (`bgpd.peer.X.flowspec`)
- FlowSpec redirect IPv4 (`irpdetectd.flowspec.ipv4.redirect`) & FlowSpec redirect IPv6 (`irpdetectd.flowspec.ipv6.redirect`)
- Default FlowSpec reaction (`irpdetectd.flowspec.reaction`) should be set to 1 (Redirect)
- Moderated/Automated Mode (`irpdetectd.mode`) requires thresholds to be set to operate properly (`irpdetectd.flowspec.threshold.kpps` and/or `irpdetectd.flowspec.threshold.mbps`).

2.11.5 Mitigation prefix size (IPv4 / IPv6)

Threat Mitigation can be configured to perform an action against the attacked host or subnet.

The size of the subnet mask for the announced IPv4 prefixes (`irpdetectd.ipv4.prefix_size`) and IPv6 prefixes (`irpdetectd.ipv6.prefix_size`) can be adjusted to the desired value.

2.12 Administrative Components


There are two additional components in the IRP system: **dbcron** and **irpmng**. For detailed configuration parameters description, please see the [Administrative settings](#) section.

Dbcron handles periodic database maintenance tasks and statistics processing.

Irpmng performs various management tasks and provides a CLI interface to query certain IRP data.

2.13 Failover Configuration

Failover relies on many operating system, networking and database components to work together. All these must be planned in advance and require careful implementation.

 We recommend that you request Noction's engineers to assist with failover configuration.


Subsequent sections should be considered in presented order when configuring a fresh IRP failover setup. Refer to them when troubleshooting or restoring services.


2.13.1 Initial failover configuration

Prerequisites

Before proceeding with failover configuration the following prerequisites must be met:

- One IRP node is configured and fully functional. We will refer to this node as \$IRPMaster.
- Second IRP node is installed with the same version of IRP as on \$IRPMaster. We will refer to this node as \$IRPSlave.

 Second IRP node MUST run the same operating system as \$IRPMaster.

 When troubleshooting problems, besides checking matching IRP versions ensure the same versions of irp MySQL databases are installed on both failover nodes.

- IRP services, MySQL and HTTP daemons are stopped on \$IRPSlave node.
- Network operator can SSH to both \$IRPMaster and \$IRPSlave and subsequent commands are assumed to be run from a \$IRPMaster console.

 \$IRPMaster and \$IRPSlave must have different hostnames.

Configure communication channel from \$IRPMaster to \$IRPSlave

This channel is used during failover configuration and subsequently \$IRPMaster uses it to synchronize IRP configuration changes when these occur. It uses key-based authentication without a passphrase to allow automated logins on \$IRPSlave by \$IRPMaster processes.

 Adjust firewalls if any so that \$IRPMaster node can access \$IRPSlave via SSH.

Generate SSH keys pair WITHOUT passphrase:

Listing 2.57: Generate keys on \$IRPMaster

```
root@IRPMaster ~ # ssh-keygen -t rsa -b 2048 -f ~/.ssh/id_rsa -C "
failover@noction"
```

i Default keys files are used. In case your system needs additional keys for other purposes we advise that those keys are assigned a different name. If this is not possible then keys file name designated for failover use should be also specified in IRP configuration parameter `global.failover_identity_file`.

Copy public SSH key from master to slave instance:

Listing 2.58: Install public key on \$IRPSlave

```
root@IRPMaster ~ # cat ~/.ssh/id_rsa.pub | while read key; do ssh $IRPSlave
"echo $key >> ~/.ssh/authorized_keys"; done
```

Check if \$IRPMaster can login to \$IRPSlave without using a password:

Listing 2.59: Check SSH certificate-based authentication works

```
root@IRPMaster ~ # ssh $IRPSlave
```

Install certificate and keys for MySQL Multi-Master replication between \$IRPMaster and \$IRPSlave

MySQL Multi-Master replication uses separate communication channels that also require authentication. IRP failover uses key-based authentication for these channels too.

i Adjust firewalls if any so that \$IRPMaster and \$IRPSlave can communicate with each other bidirectionally.

Create Certificate Authority and server/client certificates and keys. Commands must be run on both \$IRPMaster and \$IRPSlave nodes:

Listing 2.60: Generate CA and certificates

```
# cd && rm -rvf irp-certs && mkdir -p irp-certs && cd irp-certs

# openssl genrsa 2048 > $(hostname -s)-ca-key.pem

# openssl req -new -x509 -nodes -days 3600 -subj "/C=US/ST=CA/L=Palo Alto/O=
=Noction/OU=Intelligent Routing Platform/CN=$(/bin/hostname) CA/
emailAddress=support@noction.com" -key $(hostname -s)-ca-key.pem -out $(
hostname -s)-ca-cert.pem

# openssl req -newkey rsa:2048 -days 3600 -subj "/C=US/ST=CA/L=Palo Alto/O=
Noction/OU=Intelligent Routing Platform/CN=$(/bin/hostname) server/
emailAddress=support@noction.com" -nodes -keyout $(hostname -s)-server-
key.pem -out $(hostname -s)-server-req.pem

# openssl rsa -in $(hostname -s)-server-key.pem -out $(hostname -s)-server-
key.pem
```

```
# openssl x509 -req -in $(hostname -s)-server-req.pem -days 3600 -CA $(
    hostname -s)-ca-cert.pem -CAkey $(hostname -s)-ca-key.pem -set_serial 01
    -out $(hostname -s)-server-cert.pem

# openssl req -newkey rsa:2048 -days 3600 -subj "/C=US/ST=CA/L=Palo Alto/O=
    Noction/OU=Intelligent Routing Platform/CN=$(bin/hostname) client/
    emailAddress=support@noction.com" -nodes -keyout $(hostname -s)-client-
    key.pem -out $(hostname -s)-client-req.pem

# openssl rsa -in $(hostname -s)-client-key.pem -out $(hostname -s)-client-
    key.pem

# openssl x509 -req -in $(hostname -s)-client-req.pem -days 3600 -CA $(
    hostname -s)-ca-cert.pem -CAkey $(hostname -s)-ca-key.pem -set_serial 01
    -out $(hostname -s)-client-cert.pem
```

Verify certificates. Commands must be run on both \$IRPMaster and \$IRPSlave nodes:

Listing 2.61: Verify certificates

```
# openssl verify -CAfile $(hostname -s)-ca-cert.pem $(hostname -s)-server-
    cert.pem $(hostname -s)-client-cert.pem

server-cert.pem: OK
client-cert.pem: OK
```

Install certificates in designated directories. Commands must be run on both \$IRPMaster and \$IRPSlave nodes:

Listing 2.62: Install certificates in designated directories

```
# mkdir -p /etc/pki/tls/certs/mysql/server/ /etc/pki/tls/certs/mysql/client
    / /etc/pki/tls/private/mysql/server/ /etc/pki/tls/private/mysql/client/

# cp $(hostname -s)-ca-cert.pem $(hostname -s)-server-cert.pem /etc/pki/tls
    /certs/mysql/server/
# cp $(hostname -s)-ca-key.pem $(hostname -s)-server-key.pem /etc/pki/tls/
    private/mysql/server/
# cp $(hostname -s)-client-cert.pem /etc/pki/tls/certs/mysql/client/
# cp $(hostname -s)-client-key.pem /etc/pki/tls/private/mysql/client/

# cd && rm -rvf irp-certs
```

Cross copy client key and certificates:

Listing 2.63: Cross copy client key and certificates

```
root@IRPMaster ~# scp "/etc/pki/tls/certs/mysql/server/$IRPMaster-ca-cert.
    pem" "$IRPSlave:/etc/pki/tls/certs/mysql/client/"
root@IRPMaster ~# scp "/etc/pki/tls/certs/mysql/client/$IRPMaster-client-
    cert.pem" "$IRPSlave:/etc/pki/tls/certs/mysql/client/"
root@IRPMaster ~# scp "/etc/pki/tls/private/mysql/client/$IRPMaster-client-
    key.pem" "$IRPSlave:/etc/pki/tls/private/mysql/client/"

root@IRPMaster ~# scp "$IRPSlave:/etc/pki/tls/certs/mysql/server/$IRPSlave-
    ca-cert.pem" "/etc/pki/tls/certs/mysql/client/"
root@IRPMaster ~# scp "$IRPSlave:/etc/pki/tls/certs/mysql/client/$IRPSlave-
    client-cert.pem" "/etc/pki/tls/certs/mysql/client/"
```

```
root@IRPMaster ~# scp "$IRPSLAVE:/etc/pki/tls/private/mysql/client/
$IRPSLAVE-client-key.pem" "/etc/pki/tls/private/mysql/client/"
```

Adjust file permissions. Commands must be run on both \$IRPMaster and \$IRPSLAVE nodes:

Listing 2.64: Set file permissions for keys and certificates

```
# chown -R mysql:mysql /etc/pki/tls/certs/mysql/ /etc/pki/tls/private/mysql/
/
# chmod 0600 /etc/pki/tls/private/mysql/server/* /etc/pki/tls/private/mysql/
/client/*
```

Configure MySQL replication on \$IRPSLAVE

Each node of MySQL Multi-Master replication is assigned its own identifier and references previously configured certificates. There are also configuration parameters such as binary/relay log file names and auto-increment and auto-increment-offset values.

IRP includes a template config file `/usr/share/doc/irp/irp.my_repl_slave.cnf.template`. The template designates \$IRPSLAVE as second server of the Multi-Master replication and includes references to `$(hostname -s)` that need to be replaced with the actual hostname of \$IRPSLAVE before installing. Apply the changes and review the configuration file. Alternatively a command like the example below can be used to create \$IRPSLAVE config file from template. Make sure to use the actual short hostname instead of the provided variable:

Listing 2.65: Example \$IRPSLAVE configuration from template

```
root@IRPSLAVE ~# sed 's|$(hostname -s)|$IRPSLAVE|' < /usr/share/doc/irp/irp
.my_repl_slave.cnf.template > /etc/noction/mysql/irp.my_repl_slave.cnf
```

The config file created above must be included into \$IRPSLAVE node's MySQL config `my.cnf`. It is recommended that line `"include /etc/noction/mysql/irp.my_repl_slave.cnf"` is un-commented.

Check Multi-Master configuration on \$IRPSLAVE:

Listing 2.66: Check MySQL on \$IRPSLAVE works correctly

```
root@IRPSLAVE ~# service mysqld start
root@IRPSLAVE ~# tail -f /var/log/mysqld.log
root@IRPSLAVE ~# mysql irp -e "show master status \G"
root@IRPSLAVE ~# service mysqld stop
```

Configure MySQL replication on \$IRPMaster

Configuring \$IRPMaster is done with running services.



Configuring MySQL Multi-Master replication on \$IRPMaster should only be done after confirming it works on \$IRPSLAVE.

Similarly to \$IRPSLAVE above IRP comes with a template configuration file for \$IRPMaster - `/usr/share/doc/irp/irp.my_repl_master.cnf.template`. The template designates \$IRPMaster as first server of the Multi-Master replication and includes references to `$(hostname -s)` that need to be replaced with the actual hostname of \$IRPMaster before installing. Apply the changes and review the resulting configuration file.

Alternatively a command like the example below can be used to create \$IRPMaster config file from template. Make sure to use the actual hostname instead of the provided variable:


Listing 2.67: Set \$IRPMaster as a first node for Multi-Master replication

```
root@IRPMaster ~# sed 's|$(hostname -s)|$IRPMaster|' < /usr/share/doc/irp/
    irp.my_repl_master.cnf.template > /etc/noctio/mysql/irp.my_repl_master.
    cnf
```

Again, the config file created above must be included into \$IRPMaster node's MySQL config my.cnf. It is recommended that line "include /etc/noctio/mysql/irp.my_repl_master.cnf" is un-commented. Check MySQL Multi-Master configuration on \$IRPMaster:


Listing 2.68: Check MySQL on \$IRPMaster works correctly

```
root@IRPMaster ~# service mysqld restart
root@IRPMaster ~# tail -f /var/log/mysqld.log
root@IRPMaster ~# mysql irp -e "show master status \G"
```

 If Multi-Master configuration on \$IRPMaster fails or causes unrecoverable errors, a first troubleshooting step is to comment back the included line in /etc/my.cnf on master and slave and restart mysqld service to revert to previous good configuration.

Create replication grants on \$IRPMaster

Replication requires a MySQL user with corresponding access rights to replicated databases. The user must be assigned a password and is designated as connecting correspondingly from \$IRPMaster or \$IRPSLAVE. Unfortunately hostnames cannot be used for this and the exact IP addresses of the corresponding nodes must be specified.

 The user is created only once in our procedure since after being created the database on \$IRPMaster is manually transferred to \$IRPSLAVE and the user will be copied as part of this process.

Grant replication access to replication user:

Listing 2.69: Replication user and grants

```
mysql> CREATE USER 'irprepl'@'<mysql_slave1_ip_address>' IDENTIFIED BY '<
    replication_user_password>';
mysql> GRANT REPLICATION SLAVE ON *.* TO 'irprepl'@'<
    mysql_masterslave1_ip_address>' REQUIRE CIPHER 'DHE-RSA-AES256-SHA';
mysql> CREATE USER 'irprepl'@'<mysql_master2_ip_address>' IDENTIFIED BY '<
    replication_user_password>';
mysql> GRANT REPLICATION SLAVE ON *.* TO 'irprepl'@'<
    mysql_slave2_ip_address>' REQUIRE CIPHER 'DHE-RSA-AES256-SHA';
```

Copy IRP database configuration and database to \$IRPSLAVE

IRP failover copies the irp database and corresponding configuration file from \$IRPMaster to \$IRPSLAVE before activating second server. This avoids the need to manually verify and synchronize identifiers.

Copy root's .my.cnf config file if exists:

Listing 2.70: Copy database root user configuration file

```
root@IRPMaster ~# scp /root/.my.cnf $IRPSLAVE:/root/
```

Copy config files:


Listing 2.71: Copy database configuration files

```
root@IRPMaster ~# scp /etc/noction/db.global.conf /etc/noction/db.frontend.conf $IRPSLAVE:/etc/noction/
root@IRPMaster ~# scp /etc/noction/clickhouse/users.xml $IRPSLAVE:/etc/noction/clickhouse/
```

Preliminary copy database files:

Listing 2.72: Copy database data files

```
root@IRPMaster ~# rsync -av --progress --delete --delete-after --exclude="master.info" --exclude="relay-log.info" --exclude="*-bin.*" --exclude="*-relay.*" /var/lib/mysql/ $IRPSLAVE:/var/lib/mysql/
```

 Preliminary copy ensures that large files that take a long time to copy are synced to \$IRPSLAVE without stopping MySQL daemon on \$IRPMaster and only a reduced number of differences will need to be synced while MySQL is stopped. This operation can be rerun one more time to reduce the duration of the downtime on \$IRPMaster even more.


Finish copy of database files (OS with Systemd):

Listing 2.73: Copy differences of database files (OS with Systemd)

```
root@IRPMaster ~# systemctl stop httpd24-httpd mariadb clickhouse-server # CentOS
root@IRPMaster ~# systemctl stop apache2 mysql clickhouse-server # Ubuntu
root@IRPMaster ~# systemctl start irp-stop-nobgpd.target

systemctl start irp-shutdown-except-bgpd.target
systemctl start irp-shutdown.target

root@IRPMaster ~# cd /var/lib/mysql && rm -vf ./master.info ./relay-log.info ./*-bin.* ./*-relay.*
root@IRPMaster ~# rsync -av --progress --delete --delete-after /var/lib/mysql/ $IRPSLAVE:/var/lib/mysql/
```

 The procedure above tries to reduce the downtime of MySQL daemon on \$IRPMaster. During this time Bgpd preserves IRP Improvements. Make sure this action takes less than `bgpd.db.timeout.withdraw`.


Finish copy of database files (OS with RC-files):

Listing 2.74: Copy differences of database files (OS with RC-files)

```
root@IRPMaster ~# service irp stop nobgpd

root@IRPMaster ~# service httpd24-httpd stop
root@IRPMaster ~# service mysqld stop
root@IRPMaster ~# service clickhouse-server stop
```

```
root@IRPMaster ~# cd /var/lib/mysql && rm -vf ./master.info ./relay-log.
info ./*-bin.* ./*-relay.*
root@IRPMaster ~# rsync -av --progress --delete --delete-after /var/lib/
mysql/ $IRPSLAVE:/var/lib/mysql/
```

 The procedure above tries to reduce the downtime of MySQL daemon on \$IRPMaster. During this time Bgpd preserves IRP Improvements. Make sure this action takes less than `bgpd.db.timeout.withdraw`.

Start MySQL on \$IRPSLAVE and check if there are no errors at `/var/log/mysqld.log` on CentOS or `/var/log/mysql/error.log` on Ubuntu.


 First \$IRPSLAVE must be checked.

Start MySQL on IRP master and check if there are no errors in MySQL logs as above.

Start replication (Slaves) on both \$IRPMaster and \$IRPSLAVE

MySQL Multi-Master replication uses a replication scheme where a replicating master acts as both replication master and replication slave. The steps above configured IRP nodes to be capable of acting as replication masters. Here we ensure that both nodes are capable of acting as replication slaves.

IRP provides a template command that needs to be adjusted for each \$IRPMaster and \$IRPSLAVE and will instruct MySQL daemon to take the replication slave role. The template is `/usr/share/doc/irp/changemasterto.template`.

 The template generates a different command for each \$IRPMaster and \$IRPSLAVE nodes and requires multiple values to be reused from configuration settings described above. The command that is run on one node points to the other node as its master.

Make \$IRPMaster a replication slave for \$IRPSLAVE:

Listing 2.75: Set \$IRPMaster as replication slave

```
$IRPMaster-mysql>
CHANGE MASTER TO
MASTER_HOST=' $IRPSLAVE-ip-address',
MASTER_USER='irprepl',
MASTER_PASSWORD=' $IRPSLAVE-password>',
MASTER_PORT=3306,
MASTER_LOG_FILE= ' $IRPSLAVE--bin.000001',
MASTER_LOG_POS= <$IRPSLAVE-bin-log-position>,
MASTER_CONNECT_RETRY=10,
MASTER_SSL=1,
MASTER_SSL_CAPATH='/etc/pki/tls/certs/mysql/client/',
MASTER_SSL_CA='/etc/pki/tls/certs/mysql/client/$IRPSLAVE-ca-cert.pem',
MASTER_SSL_CERT='/etc/pki/tls/certs/mysql/client/$IRPSLAVE-client-cert.pem'
,
MASTER_SSL_KEY='/etc/pki/tls/private/mysql/client/$IRPSLAVE-client-key.pem'
,
MASTER_SSL_CIPHER='DHE-RSA-AES256-SHA';
```

i You must manually check what values to assign to \$IRPSLAVE-bin.000001 and <\$IRPSLAVE-bin-log-position> by running the following MySQL command on \$IRPSLAVE

```
mysql> show master status
```

For the initial configuration the values for \$IRPSLAVE-bin.000001 and <\$IRPSLAVE-bin-log-position> must be as follows:

```
Binlog file: $IRPSLAVE-bin.000001
Binlog position: 106
```

Run the commands below to run the replication and check the slave status:

Listing 2.76: Starting replication slave on \$IRPMaster

```
mysql> START SLAVE \G
mysql> show slave status \G
```

i Check the Slave_IO_State, Last_IO_Errno, Last_IO_Error, Last_SQL_Errno, Last_SQL_Error values for errors. Make sure there are no errors.

Make \$IRPSLAVE a replication slave for \$IRPMaster:

Listing 2.77: Set \$IRPSLAVE as replication slave

```
$IRPSLAVE-mysql>
CHANGE MASTER TO
MASTER_HOST=' $IRPMaster-ip-address',
MASTER_USER=' irprepl',
MASTER_PASSWORD=' $IRPMaster-password',
MASTER_PORT=3306,
MASTER_LOG_FILE= ' $IRPMaster-bin.000001',
MASTER_LOG_POS= <$IRPMaster-bin-log-position>,
MASTER_CONNECT_RETRY=10,
MASTER_SSL=1,
MASTER_SSL_CAPATH='/etc/pki/tls/certs/mysql/client/',
MASTER_SSL_CA='/etc/pki/tls/certs/mysql/client/$IRPMaster-ca-cert.pem',
MASTER_SSL_CERT='/etc/pki/tls/certs/mysql/client/$IRPMaster-client-cert.pem',
MASTER_SSL_KEY='/etc/pki/tls/private/mysql/client/$IRPMaster-client-key.pem',
MASTER_SSL_CIPHER=' DHE-RSA-AES256-SHA' ;
```

i You must manually check what values to assign to \$IRPMaster-bin.000001 and <\$IRPMaster-bin-log-position> by running the following MySQL command on \$IRPMaster

```
mysql> show master status
```

For the initial configuration the values for \$IRPMaster-bin.000001 and <\$IRPMaster-bin-log-position> must be as follows:

```
Binlog file: $IRPMaster-bin.000001
Binlog position: 106
```


Run the commands below to run the replication and check the slave status:

Listing 2.78: Starting replication slave

```
mysql> START SLAVE \G
mysql> show slave status \G
```

Run IRP and HTTPD services on \$IRPMaster and \$IRPSLAVE (OS with Systemd):


Listing 2.79: Starting IRP services and Frontend (OS with Systemd)

```
# systemctl start irp.target
```

Run IRP and HTTPD services on \$IRPMaster and \$IRPSLAVE (OS with RC-files):

Listing 2.80: Starting IRP services and Frontend (OS with RC-files)

```
# service httpd24-httpd start
# service irp start
```


 Start services on \$IRPMaster first if the actions above took very long in order to shorten MySQL downtime.

Configure Failover using Wizard on \$IRPMaster

These steps should only be performed once SSH communication channel and MySQL Multi-Master irp database replication have been setup and are verified to work. Refer subsections of [Failover Configuration](#) above.


Run failover wizard:

Login into IRP's Frontend on \$IRPMaster node and run Configuration -> Setup wizard -> Setup Failover.

 A valid failover license should be acquired before failover configuration settings become available.

Configure IRP failover:

Follow Setup Failover wizard steps and provide the required details. Refer [Setup Failover wizard](#) for details. Once your configuration changes are submitted IRP will validate configuration and if it is valid the configuration will be synced to \$IRPSLAVE.

 Only after this synchronization step takes place will \$IRPSLAVE node know what is its role in this setup.

Apply configuration changes to edge routers:

Ensure designated probing IPs, BGP session(s) are also setup on edge routers. Refer to corresponding sections of this document for details.

Enable failover:

Use IRP's Frontend Configuration -> Global -> Failover to configure IRP failover on both nodes. Monitor both IRP nodes and ensure everything is within expectations.

i It is recommended that after finishing the preparatory steps above both IRP master and slave nodes run with disabled failover for a short period of time (less than 1 hour) in order to verify that all IRP components and MySQL Multi-Master replication work as expected. Keep this time interval short to avoid a split-brain situation when the two nodes make contradictory decisions.

Synchronize RRD statistics to \$IRPSLAVE

RRD statistics is collected by both IRP failover nodes and needs not be synchronized during normal operation. Still, when IRP failover nodes are setup at a large time interval between them it is recommended that RRD files are synchronized too. This will ensure that past statistics is available on both IRP nodes. During short downtimes of each of the nodes synchronization is less useful but can be performed in order to cover the short downtime gaps in RRD based graphs.

Sample command to synchronize IRP's RRD statistics:

Listing 2.81: Synchronize RRD

```
root@IRPMASTER ~ # rsync -av /var/spool/irp/ $IRPSLAVE:/var/spool/irp
```

2.13.2 Re-purpose operational IRP node into an IRP failover slave

Sometimes an existing fully configured IRP node is designated to be re-purposed as a failover slave. Take the following actions to make a node as an IRP slave:

1. upgrade IRP to the version matching IRP failover master node
2. create a backup copy of your configuration
3. delete the following configuration files: `/etc/noction/irp.conf`, `/etc/noction/exchanges.conf`, `/etc/noction/policies.conf`, `/etc/noction/inbound.conf`, `/etc/noction/user_directories.conf`
4. proceed with configuration as detailed in [Initial failover configuration](#)

2.13.3 Re-purpose operational IRP failover slave into a new master

If a master node fails catastrophically and cannot be recovered an operational IRP slave can be re-purposed to become the new master. Once re-purposing is finished a new slave node can be configured as detailed in [Initial failover configuration](#).

Failover status bar (3.2.2) includes a context menu with functions to switch IRP failover node role. Simply click on the failover status bar's 'Slave node' icon and when prompted if you want to switch role confirm it.

i If you switched a node's role by error, simply change it one more time to revert to previous configuration.

2.13.4 Recover prolonged node downtime or corrupted replication

Multi-Master replication used by IRP failover is able to cope with downtime of less than 24 hours. Replication will not be able to continue in case of prolonged downtime or corrupt replications. Recovery in this case might use either:

- new server: follow configuration steps as detailed in [Initial failover configuration](#).
- same server: follow recovery steps below.

MySQL Multi-Master recovery prerequisites

Before proceeding with recovery of replication the following prerequisites must be met:

- Currently active IRP node is designated as MySQL sync 'origin'. This node currently stores reference configuration parameters and data. These will be synced to the node being recovered and we designate it as 'destination'.
- Recovery should be scheduled during non-peak hours.
- Recovery must finish before `bgpd.db.timeout.withdraw` (default 4h) expires. If recovery can not be completed in time it is required to start MySQL on the active node.

MySQL Multi-Master recovery procedure

It shall be noted that recovery closely follows actions described in and with the clarification that data and files from origin node target destination node (instead of IRP failover master files targeting IRP failover slave). Refer sections , for details about the steps below.

The steps are as follows:

1. destination: stop httpd, irp, mysqld
2. origin: sync /etc/noction/db.* to slave:/etc/noction/
3. origin: sync /root/.my.cnf to slave:/root/.my.cnf
4. origin: sync /var/lib/mysql/ to slave:/var/lib/mysql/
exclude files:
master.info relay-log.info -bin.* -relay.*

wait until sync at (4) succeeds and continue with:
5. origin: stop httpd, irp (except bgpd), mysqld
6. origin: delete files master.info relay-log.info -bin.* -relay.*
7. origin: sync /var/lib/mysql/ to slave:/var/lib/mysql/
8. destination: start mysqld and check /var/log/mysqld.log for errors
9. origin: start mysqld and check /var/log/mysqld.log for errors
10. origin: run CHANGE MASTER TO from the /usr/share/doc/irp/changemasterto template
11. destination: run CHANGE MASTER TO from the /usr/share/doc/irp/changemasterto template
12. destination: show slave status \G
13. origin: show slave status \G
14. origin: start IRP (bgpd should be already running), httpd
15. destination: start IRP, httpd

Chapter 3

Using IRP

3.1 Getting started

The **Frontend** is a web interface with a comprehensive set of reports, graphs and diagnostic information, which can reflect the current and historical network state, as well as the benefits of IRP network optimization.

3.1.1 Accessing the system

The system frontend is accessible over HTTPS. IRP's Frontend can be accessed using any major browser, via the main IRP server IP address, or FQDN. E.g: <http://irp-system.notion.com/>. When opening the Frontend's URL, a login form will pop up:

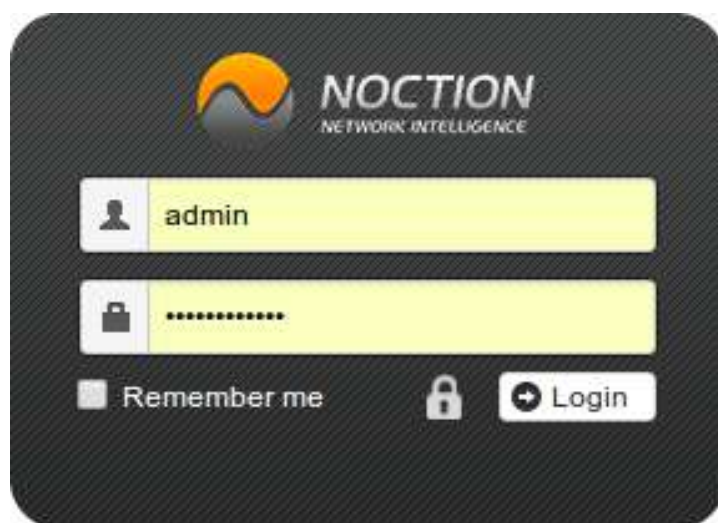


Figure 3.1.1: Frontend login form

The username and password provided by the system administrator must be used. The “Remember me” check-box can be used as well.

3.1.2 System dashboard

After a successful login, the default Dashboard will become available. It contains a standard set of gadgets - elements, which can contain a reduced version of a report or a graph.

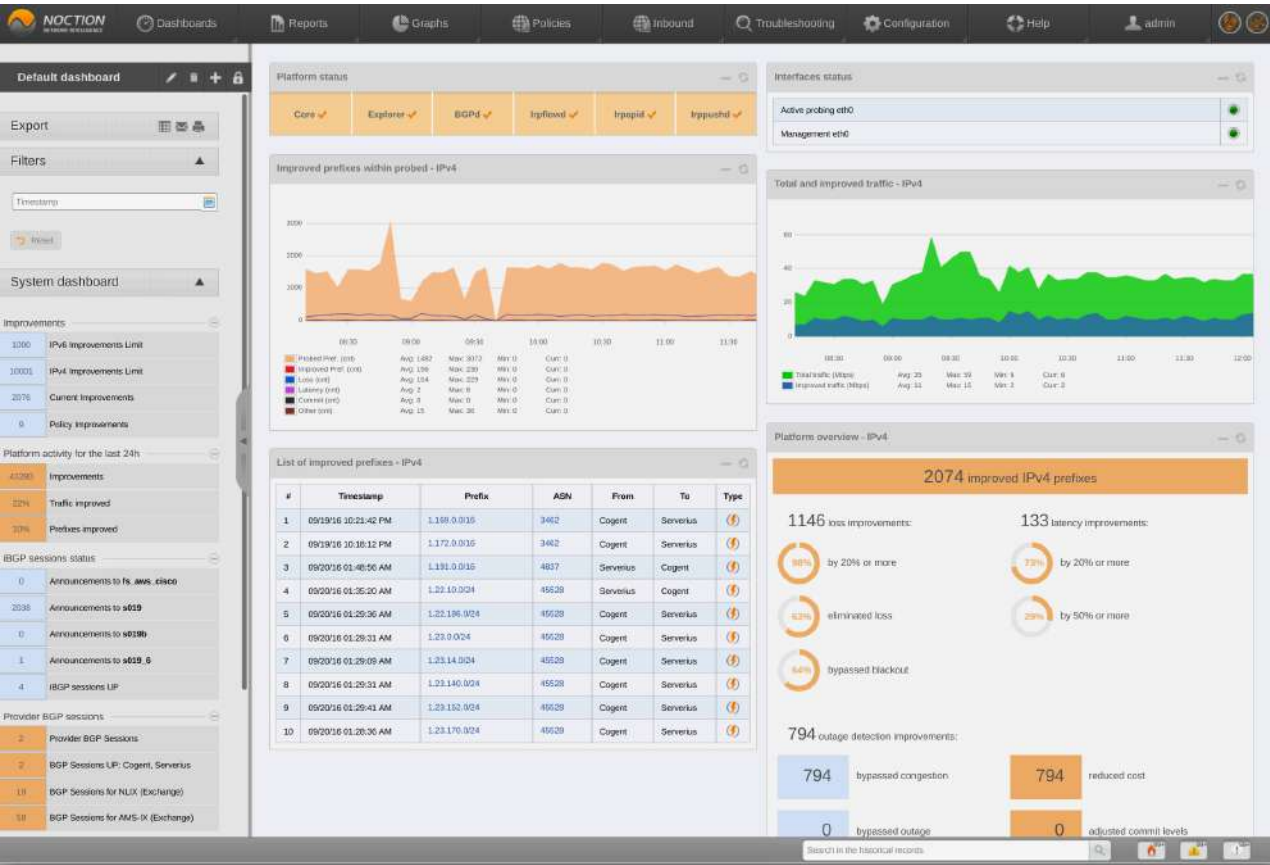


Figure 3.1.2: Default Dashboard

3.2 Frontend structure

The Frontend has several components that operate together and allow a regular user or administrator to access various reports, graphs, or configuration pages.

1. The menu bar (see Figure 3.2.1) allows quick access to any part of the system via its menu items.



Figure 3.2.1: Menu bar

2. The sidebar (See section 3.2.1) is present on the Dashboard or on any graph or report. Its contents depend on the current page, and can include filters, export functions, etc., as explained in the next section.
3. The status bar (see Figure 3.2.2), allows one to access the global search function (Section 3.2.3), as well as system events and notifications (See the [Events](#) section).



Figure 3.2.2: Status bar

3.2.1 Sidebar

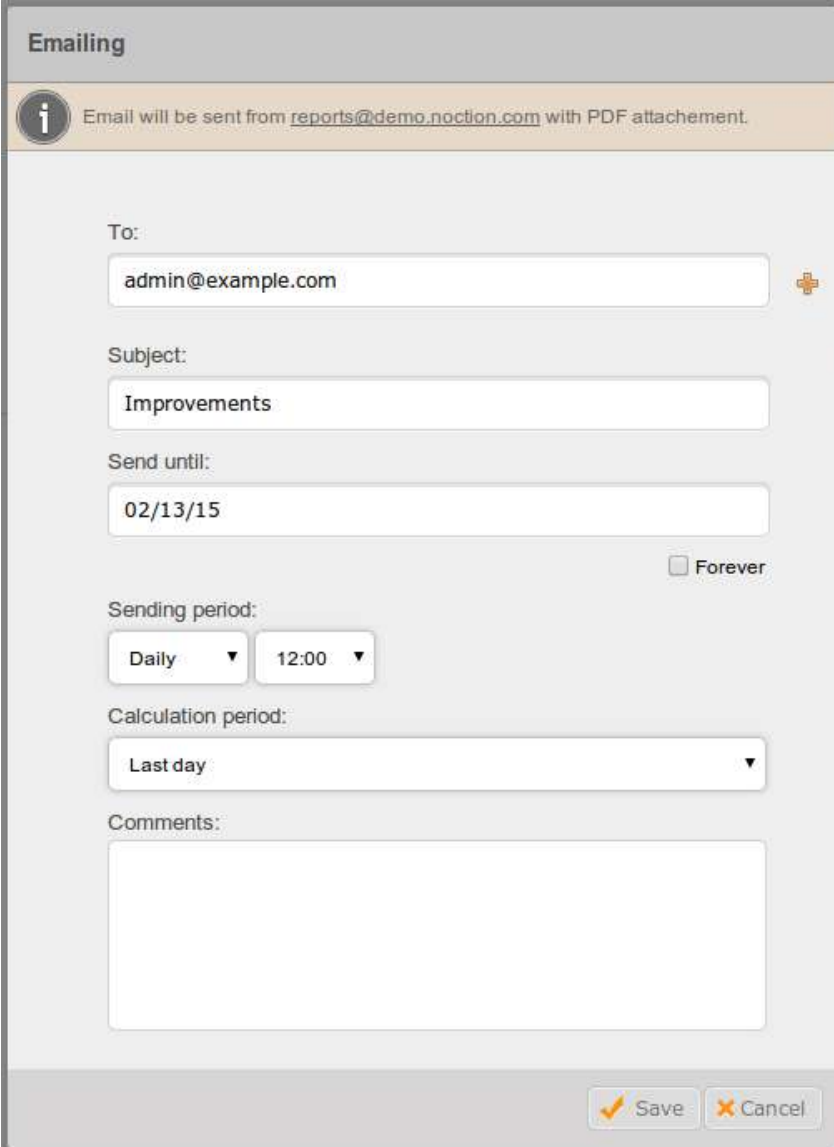
All dashboards, reports, graphs and events have a sidebar, located at the left side of the screen. Depending on the current page, the sidebar has a different set of controls, allowing users to manage the dashboard components or to filter the report and graph results.

Every sidebar contains Export, Email and Print buttons. Reports can be exported in .xls, .csv, .pdf formats, while the graphs and dashboard - only as .pdf .



Figure 3.2.3: Common sidebar buttons

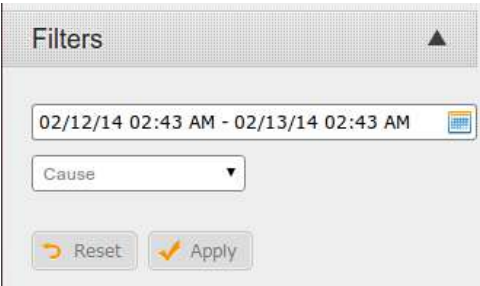
The “Email” button allows scheduling email delivery of several reports or graphs, or even an entire dashboard (Figure 3.2.4). Multiple destination email addresses can be used. Most elements in this form have informative tool-tips, that can be seen when hovering over a specific control. Note that Email Reports can be removed from <Username>→Emailing from the top panel.



The 'Emailing' configuration window has a grey header with the title 'Emailing'. Below the header is an orange banner with an information icon and the text 'Email will be sent from `reports@demo.noction.com` with PDF attachment.' The main area is light grey and contains several form fields: 'To:' with a text box containing 'admin@example.com' and a plus icon; 'Subject:' with a text box containing 'Improvements'; 'Send until:' with a text box containing '02/13/15' and a 'Forever' checkbox; 'Sending period:' with two dropdown menus set to 'Daily' and '12:00'; 'Calculation period:' with a dropdown menu set to 'Last day'; and 'Comments:' with a large empty text area. At the bottom right are 'Save' and 'Cancel' buttons with checkmark and X icons respectively.

Figure 3.2.4: Email schedule configuration

In the case of a report, custom filters can be applied, so that only specific information will be displayed in the report.



The 'Filters' configuration window has a grey header with the title 'Filters' and a collapse icon. Below the header is a text box showing the time range '02/12/14 02:43 AM - 02/13/14 02:43 AM' with a calendar icon. Below this is a dropdown menu labeled 'Cause'. At the bottom are 'Reset' and 'Apply' buttons with circular arrow and checkmark icons respectively.

Figure 3.2.5: Report filters

On graphs, reports and dashboard, the time period which the information is being provided for, can be selected by using the “Timestamp” button. An intuitive control will open, and the desired time period can be selected.

02/12/14 13:00 - 02/13/14 13:00

January, 2014

S	M	T	W	T	F	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

February, 2014

S	M	T	W	T	F	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	

March, 2014

S	M	T	W	T	F	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31					

Date range: 02/02/14 00:00 02/08/14 23:59

Custom

Cancel Save

Figure 3.2.6: Time period selection

The number of results can be adjusted on a report page, using the “Results per page” slider.

Results per page: 50

20 50 100 200 300

Figure 3.2.7: Adjusting the number of results per page

Dashboards have a different sidebar layout.

⚠ The Delete button cannot be found in the default dashboard. The default dashboard cannot be deleted.

System dashboard represents the overall system information which includes:

- Configured maximum number of improvements (Max IPv6 Improvements will be shown if IPv6 is enabled)
- Number of Routing Policies improvements
- Number of improvements performed today
- The amount of traffic improved today
- Percentage of prefixes improved today
- Number of operating iBGP sessions
- Number of announcements to the edge router
- Announcements to each from configured routers
- Number of BGP sessions with the providers
- Number of operating BGP sessions with the providers

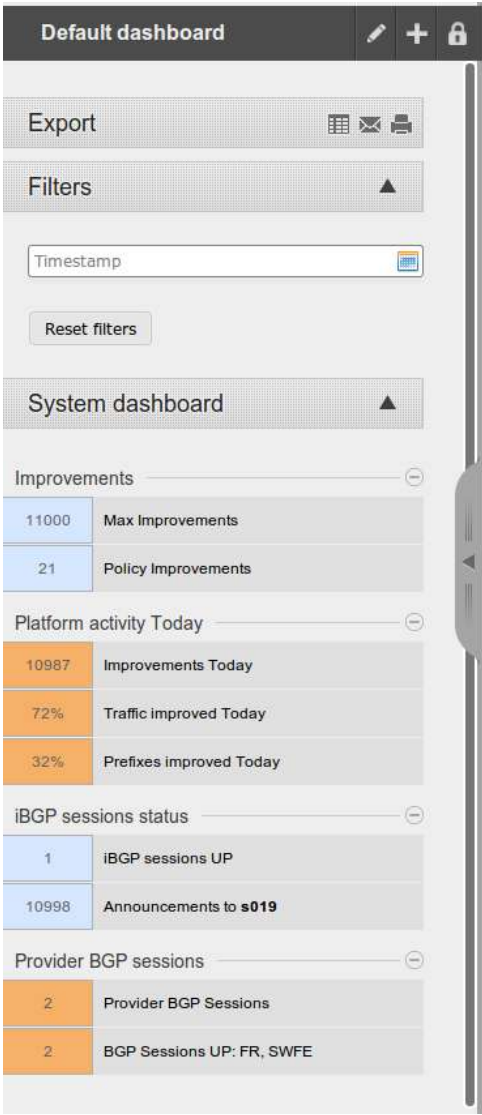


Figure 3.2.8: Dashboard buttons

See also: [Customizing the dashboard](#).

3.2.2 Failover status

When failover is configured and operational the status of the components is highlighted below system dashboard for each master and slave nodes.

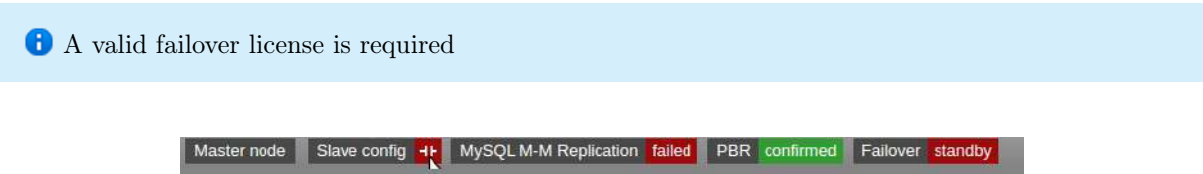


Figure 3.2.9: Failover status of master node

Slave node status shall be considered together with component status in the dashboard. During normal operation components Core, Explorer and Irppushd are stopped.



Figure 3.2.10: Failover status of slave node

3.2.3 Global search

The Frontend allows to search for historical information on optimized prefixes, AS Numbers and AS Names. The search function can be accessed using the Search box in the status bar.



Figure 3.2.11: Search box

Enter the needed prefix, AS number or AS name in the search box, then press Enter or click on the magnifying glass button.

While entering the search terms, results are dynamically loaded in a modal window (Figure 3.2.12).



Figure 3.2.12: Dynamical global search results

i When displaying the complete results, specific prefixes can be manually removed, or scheduled for re-probing, by using the buttons next to each result (figure 3.2.13).

Search results for: 46.

Improvements		
46.26.165.0/24 Old loss 0%, avg rt 128 ms, new loss 0%, avg rt 110 ms, ASN: 39824, Cause: Performance Reports > Improvements		30/10/10 01:00:00 PM
46.176.192.0/19 Old loss 0%, avg rt 99 ms, new loss 0%, avg rt 94 ms, ASN: 3929, Cause: Cost Reports > Improvements		30/10/10 01:00:00 PM
46.261.64.0/19 Old loss 0%, avg rt 74 ms, new loss 0%, avg rt 80 ms, ASN: 6548, Cause: Cost Reports > Improvements		30/10/10 01:00:00 PM
46.268.68.0/22 Old loss 0%, avg rt 120 ms, new loss 0%, avg rt 109 ms, ASN: 41822, Cause: Cost Reports > Improvements		30/10/10 01:00:00 PM
46.176.132.0/24 Old loss 0%, avg rt 60 ms, new loss 0%, avg rt 60 ms, ASN: 51261, Cause: Cost Reports > Improvements		30/10/10 01:00:00 PM
Show all 43223 results in Improvements		
AS-path problems		
46.161.168.0/21 No routing updates, current route is bad. Problem: Unconfirmed prefix. Reports > AS-Path problems		30/10/10 01:00:00 PM
46.262.192.0/19 Pending prefix. Reports > AS-Path problems		30/10/10 01:00:00 PM
46.262.32.0/19 Pending prefix. Reports > AS-Path problems		30/10/10 01:00:00 PM
46.262.192.0/19 Resolving 40.552.162.0/19 from peer PR to peer SINPE. Cost improvement. Problem: Unconfirmed prefix. Reports > AS-Path problems		30/10/10 01:00:00 PM
46.262.32.0/19 Resolving 40.552.32.0/19 from peer SINPE to peer PR. Old loss 100%, avg rt 0 ms, new loss 59%, avg rt 151 ms. Problem: Unconfirmed prefix. Reports > AS-Path problems		30/10/10 01:00:00 PM
Show all 51259 results in AS-Path problems		

Figure 3.2.13: Global search results

The *Show all X results in Y category* buttons can also be used, to display all matching results from the selected category (Improvements or AS Path problems)

3.2.4 Warning bars

The warning bars are shown in case the attention should be drawn to operational/performance issues, some specific IRP configuration settings or license/payment issues.

Attention! All providers have the same precedence value and improve in group feature is not enabled for all of them.

Figure 3.2.14: Precedence and improve in a group warning

Attention! The internal BGP Monitor is not configured, IRP will not be able to detect provider failure and will continue to announce improvements through failed providers too.

Figure 3.2.15: Internal Monitor warning

There are 4 problems in VIP improvements lists. Under current Explorer configuration IRP will not be able to probe all VIP, new and reprobed destinations. It is recommended to either remove some VIP entries, increase VIP re-probing period, or increase the standard re-probing period.

Figure 3.2.16: Explorer performance warning

IRP license will expire in less than one hour. Renew it here

Figure 3.2.17: License expiration warning

3.3 Customizing the dashboard

The Default dashboard is the first screen that can be seen after a successful login to the system. It contains a default set of gadgets, that can be customized by adding additional gadgets, removing unnecessary gadgets, or modifying specific gadget's settings.

3.3.1 Custom dashboard gadgets

3.3.1.1 Platform status

Platform status			
BGPd (1) ✓	Core ✓	Explorer ✓	Irpflowd ✓

Figure 3.3.1: Platform status

3.3.1.2 Interfaces status





Interfaces status	
Active probing eth0	
Management eth0	
Collector eth1	
Collector eth2	

Figure 3.3.2: Interfaces status

3.3.1.3 List of improved prefixes











List of improved prefixes						
#	Time stamp	Prefix	ASN	From	To	Type
1	10/31/13 16:49:17	94.181.20.0/22	41661	SWFE	FR	
2	10/31/13 16:49:17	92.81.224.0/19	9050	SWFE	FR	
3	10/31/13 16:49:17	88.113.64.0/19	719	SWFE	FR	
4	10/31/13 16:49:17	178.127.128.0/19	6697	FR	SWFE	
5	10/31/13 16:49:11	90.189.0.0/19	41440	SWFE	FR	
6	10/31/13 16:49:11	77.34.224.0/20	12332	SWFE	FR	
7	10/31/13 16:49:00	37.79.32.0/19	34875	SWFE	FR	
8	10/31/13 16:49:00	168.62.0.0/19	8075	SWFE	FR	
9	10/31/13 16:48:55	90.157.32.0/19	35154	FR	SWFE	
10	10/31/13 16:48:55	78.30.192.0/19	35816	SWFE	FR	

Figure 3.3.3: List of improved prefixes

3.3.1.4 AS-Path problems











AS-path problems					
#	AS-Pattern	Problem	Detection time	Problem status	
				Confirmation rate	Status
1	50384-34892	Outage	10/31/13 04:52:46 PM	20%	
2	12389-43975	Congestion	10/31/13 04:44:58 PM	1.49%	
3	12956-22085	Congestion	10/31/13 04:21:32 PM	0%	
4	31133-49718	Outage	10/31/13 04:19:22 PM	14.29%	
5	31133-34364	Outage	10/31/13 04:05:23 PM	14.29%	
6	12389-43975	Congestion	10/31/13 04:02:29 PM	1.52%	
7	6453-812	Congestion	10/31/13 03:15:14 PM	0%	
8	1273-10030	Outage	10/31/13 03:05:25 PM	66.67%	
9	8075-8069	Congestion	10/31/13 02:31:27 PM	11.11%	
10	10026-8075	Congestion	10/31/13 02:50:35 PM	22.22%	

Figure 3.3.4: AS-Path problems

3.4 Reports

The IRP Frontend comes with a comprehensive set of reports, reflecting the current state of the network as well as overall statistics on a particular IRP instance performance. Click on the desired report title to open it.

Platform	Improvements	History	Statistics
Platform overview	Current Improvements 📈	Historical records	Country statistics 📈
Platform performance	Improvements To Exchanges	Inbound Improvements history	AS statistics
Loss Improvements	Current Inbound Improvements	Monitor history	Top volume prefixes
Latency Improvements	VIP Improvements	Cost overview	Prefix statistics 📈
Probe conversion	Congestion and outages		Top volume AS
Monitors Status	Providers overall		Top problem AS
	Providers performance		Exchanges statistics
	Providers daily performance		
	Probes today		

Figure 3.4.1: Reports menu

3.4.1 Platform overview

The “Platform Overview” report provides information on the overall rate of improvements based upon latency, loss and cost. This report can be used for a quick overview of the overall system performance and improvements.

One can see here that almost all the improved routes had a loss drop of 20% or more. For more than 80% of the improved routes the loss was fully eliminated. There are also 7% of the improved routes, which were redirected from a complete blackout. The report also provides the amount of the accomplished route improvements made, based upon latency, and cost (if the Cost optimization is enabled, see [Cost optimization](#)).

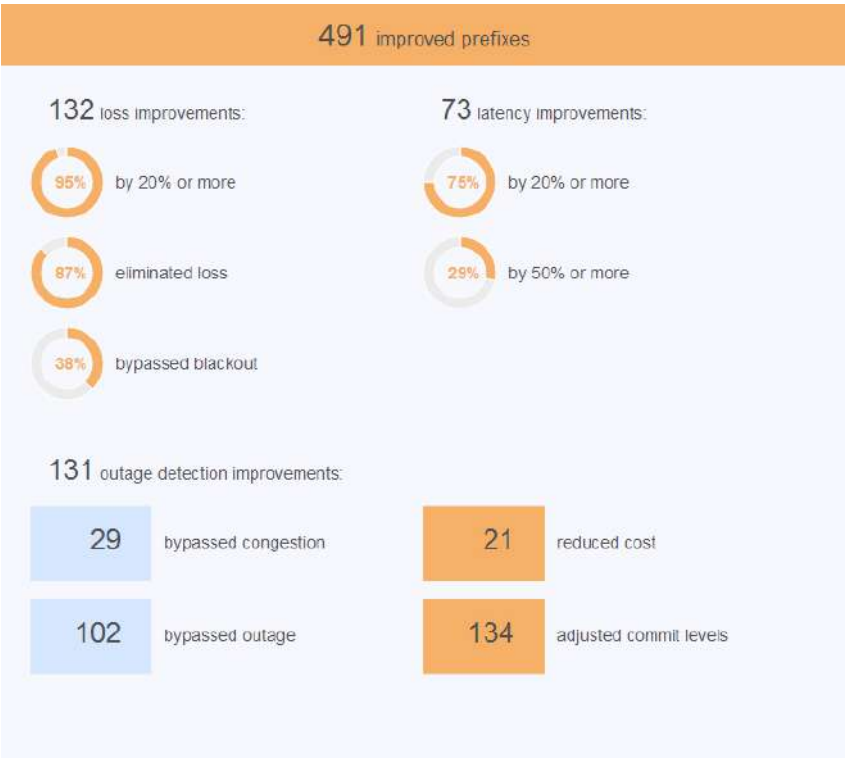


Figure 3.4.2: Platform overview

3.4.2 Platform performance

The “Platform Performance” report offers overall statistics on improvements made, based upon performance reasons, commit control and cost reasons on a daily, weekly and monthly basis.

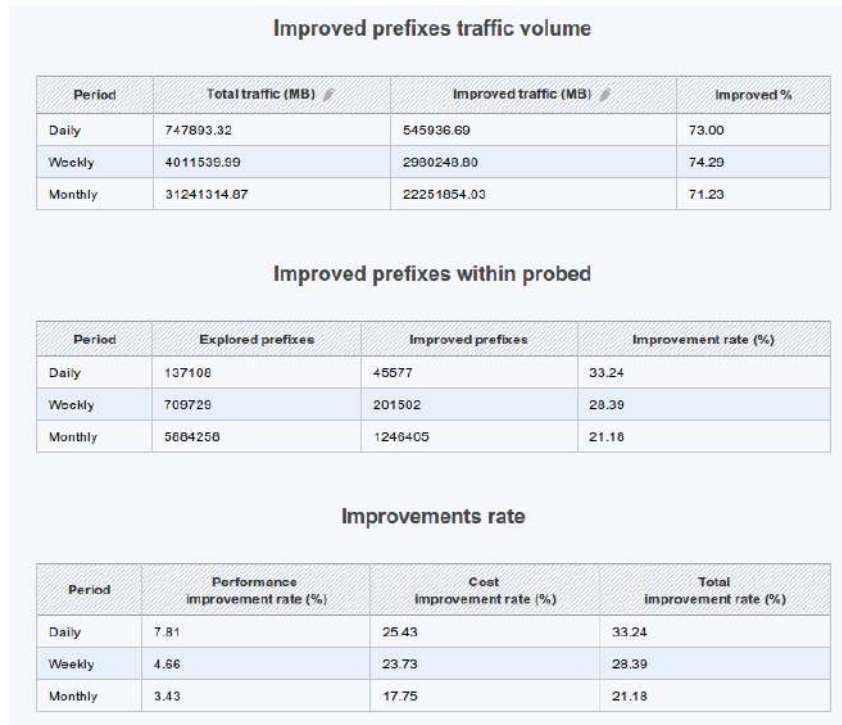


Figure 3.4.3: Platform performance

3.4.3 Loss Improvements

The Loss Improvements report highlights how much of the "Packet loss problem" IRP detected and was able to solve during a specific timeperiod.

The report depicts loss average rates for groups of destinations before and after IRP assessed and re-routed the traffic. The groups proposed and highlighted in number of prefixes depicted by the pie charts below represent:

- Sampled destinations that cover all destination prefixes - good or problematic - that have been probed by IRP;
- Problem destinations include only routes/prefixes where IRP detected packet loss and improved them. The number of such destinations, as expected, is usually significantly lower. At the same time the loss before and after IRP improved them average to a much larger values;
- Loss eliminated represents those problem destinations that IRP was able to improve to a better route with zero loss. By definition this group will have a zero remaining loss but the original average loss highlights how much better those end users perceive their service.;
- Loss reduced represents those problem destinations that IRP improved but was not able to eliminate loss completely.

In absolute values the bar chart highlights the average loss rate values before and after improvement for the groups of destinations listed above. Values of 1, 3, 5,7% are usual for the . Values exceeding 10% are a bit worrisome. While it allows IRP to boast a bigger impact for your network it also means that the routing information you receive is not very accurate and that your providers might not be making the very best effort to service your traffic.




Figure 3.4.4: Loss improvements

Loss relative rates highlighting the percentage magnitude of the improvement in loss rates can be seen by opting for an 100% vertical axis on the reports filters. This view quantifies how much better loss characteristics on the network improved for various groups of end users. Notice how the ~1% loss figure above for all Sampled destinations expands to approximately ~32% of total number of lost packets being eliminated after IRP discovered the better routes and applied the improvements.




Figure 3.4.5: Loss relative rates

Problem destinations - highlights destination counts of affected routes in the groups as above. The charts present the rate of problem destinations identified by IRP and subsequently highlights how well IRP was able to improve on them. The counts allow you to infer whether there are 20 customers suffering minor packet loss (say, 5% each) or 2 customers suffering severe packet loss (for example 50% each). If for the top and middle graphs above IRP was able to reduce packet loss by 80% these pie charts allow you to infer that there were 20 or 2 customers affected and that for some of them IRP solved the problem completely and how many others are still suffering.




Figure 3.4.6: Loss-improved destinations

The report can be switched to table view in order to review actual numbers across the different groups and the corresponding loss average values.

3.4.4 Latency Improvements

The Latency Improvements report presents an overview for average latency values before and after IRP optimization during a specific timeperiod and depicts them across the following populations:

- Sample destinations that cover good and problematic destinations that IRP probed during selected timeperiod;
- Problematic destinations cover those routes/prefixes that IRP assessed to have excessive latency rates and identified a better route via a different upstream provider;
- 20% or more cover those destinations where IRP was able to identify a route with latency improvements of 20% or more;
- 50% or more cover those destinations where IRP was able to identify a route with latency improvements of 50% or more.

i Note that the 20% and 50% groups are depicted only to preserve similarity in the number of groups in Loss Improvement report. At the same time the 20% and 50% thresholds represent significant quality of service improvements for affected end users and should warrant special consideration.

Latency values bar-chart highlights the latency averages before and after improvement for the designated groups of destinations.




Figure 3.4.7: Latency values (ms)

Percentages of latency improvement averages can be seen by opting to use 100% vertical axis instead of the usual miliseconds. The option is available in report filters.



Figure 3.4.8: Latency values (%)

Latency destinations - highlights destination counts in the groups as above.




Figure 3.4.9: Latency destinations

The report can be switched to table view in order to review actual numbers across the different groups and the corresponding averages.

3.4.5 Probe conversion rates

Probe conversion rates - highlights the counts of networks/prefixes probed by IRP and their conversion rate into any of the possible improvements:

- Commit
- Performance
- Cost




Figure 3.4.10: Probe conversion rates

3.4.6 Current improvements

The “Current Improvements” report provides a list of the currently active improvements as per network. A drill-down capability to a specific ASN or prefix is also provided. The report shows the improved prefix and the providers from and to which the traffic was redirected. It also provides the reason, which the improvement was based upon, along with the performance values before and after the traffic was rerouted.

If necessary, specific improvements can be manually deleted by clicking the checkbox next to the needed prefix, afterward using the “Remove selected” button.

Current Improvements											2017-06-07 11:05:52 - 2017-06-08 11:05:52 [switch to map view]	
#	Timestamp	Prefix	ASN	AS Name	From	To	Type	Age	Prob	Details	Rd	
1	17 min 0 seconds ago	115.182.64.0/20	4847	China Networks Inter-Excha	Serverius	Cogent				Old loss 11%, avg rtt 327 ms, new loss 6%, avg rtt 375 ms.	0	
2	17 min 0 seconds ago	122.200.73.0/24	4847	China Networks Inter-Excha	Cogent	Serverius				Old loss 0%, avg rtt 291 ms, new loss 0%, avg rtt 185 ms.	0	
3	17 min 0 seconds ago	125.208.0.0/24	63530	Beijing Guanghuan newslett	Cogent	Serverius				Old loss 15%, avg rtt 364 ms, new loss 0%, avg rtt 363 ms.	0	
4	17 min 0 seconds ago	202.165.181.0/24	4847	China Networks Inter-Excha	Cogent	Serverius				Old loss 0%, avg rtt 363 ms, new loss 0%, avg rtt 260 ms.	0	
5	17 min 0 seconds ago	219.232.56.0/21	4847	China Networks Inter-Excha	Cogent	Serverius				Old loss 0%, avg rtt 361 ms, new loss 0%, avg rtt 251 ms.	0	
6	17 min 0 seconds ago	68.247.0.0/16	3651	Sprint	Cogent	Serverius				Old loss 100%, avg rtt 0 ms, new loss 46%, avg rtt 104 ms.	0	
7	17 min 0 seconds ago	72.60.0.0/16	3651	Sprint	Cogent	Serverius				Old loss 77%, avg rtt 94 ms, new loss 0%, avg rtt 88 ms.	0	
8	17 min 5 seconds ago	117.114.132.0/21	4847	China Networks Inter-Excha	Cogent	Serverius				Old loss 9%, avg rtt 361 ms, new loss 0%, avg rtt 229 ms.	0	
9	17 min 5 seconds ago	117.224.144.0/21	9829	National Internet Backbone I	Serverius	Cogent				Old loss 63%, avg rtt 180 ms, new loss 44%, avg rtt 164 ms.	0	
10	17 min 5 seconds ago	121.52.160.0/24	4847	China Networks Inter-Excha	Cogent	Serverius				Old loss 0%, avg rtt 343 ms, new loss 0%, avg rtt 297 ms.	0	
11	17 min 5 seconds ago	174.157.0.0/17	3651	Sprint	Cogent	Serverius				Old loss 82%, avg rtt 115 ms, new loss 0%, avg rtt 114 ms.	0	
12	17 min 5 seconds ago	174.157.128.0/17	3651	Sprint	Cogent	Serverius				Old loss 41%, avg rtt 111 ms, new loss 0%, avg rtt 116 ms.	0	

Figure 3.4.11: Current improvements

i In some specific cases, the “From” and “To” fields will contain the same provider. This is not a platform issue and is related to the Outage detection algorithm specifics. For more details see the [Outage detection](#) and [VIP Improvements](#) sections.

i The “ASN” field can be empty in case there is no information about the AS number (a prefix belongs to) in the IRP ASN dictionary. As soon as the ASN information is available in the dictionary, the field will be filled accordingly.

Of note is that Current Improvements report also has a map view.

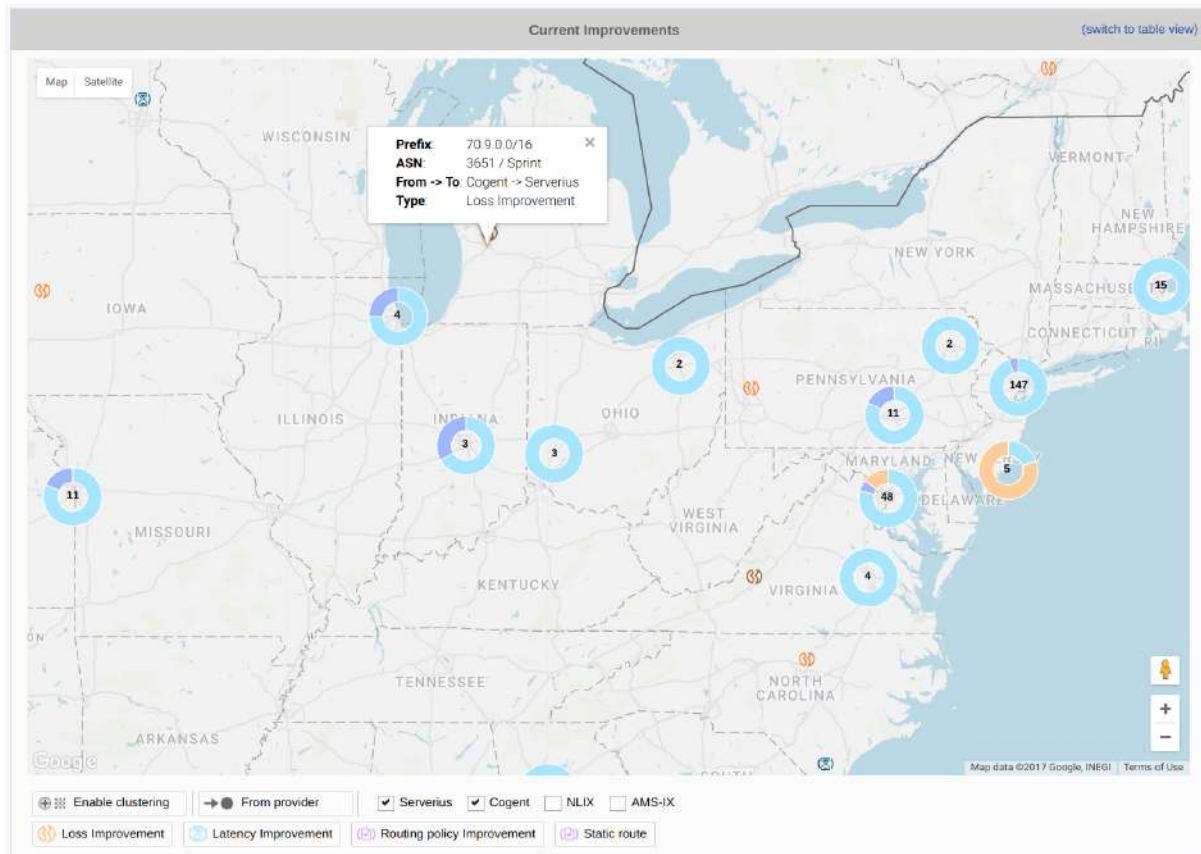


Figure 3.4.12: Current improvements map view

The map view besides offering the usual zooming and panning capabilities also offers clustering and highlighting of improvements by either provider or improvement type.

3.4.7 Improvements to Exchanges

The “Improvements to Exchanges” report contains a complete list of current improvements where the new provider is an Internet Exchange. The report includes data about improved prefixes and their relative traffic by means of last minute and average bandwidth usage. Average bandwidth is estimated based on current hour statistics.

#	Timestamp	Prefix	Asn	From	To	Age	Prob	Details	Cause	RD	Estimated avg BW (Mbps) ↗	BW in last minute (Mbps) ↗
1	03:03:47 PM	106.51.192.0/21	24309	COGENT	NLIX	🕒	🛡️	Old loss 0%, avg rt 158 ms, new loss 0%, avg rt 168 ms	🔥	Global	0.02	9.00
2	03:08:11 PM	113.29.129.0/24	38182	COGENT	NLIX	🕒	🛡️	Old loss 0%, avg rt 271 ms, new loss 0%, avg rt 192 ms	🔥	Global	0.00	9.00
3	03:04:53 PM	117.218.0.0/20	9829	COGENT	NLIX	🕒	🛡️	Old loss 0%, avg rt 213 ms, new loss 0%, avg rt 216 ms	🔥	Global	0.47	9.00
4	03:08:06 PM	122.177.176.0/20	24560	COGENT	NLIX	🕒	🛡️	Old loss 0%, avg rt 201 ms, new loss 0%, avg rt 170 ms	🔥	Global	0.14	9.00
5	03:05:41 PM	175.109.98.0/24	38623	COGENT	NLIX	🕒	🛡️	Old loss 0%, avg rt 371 ms, new loss 0%, avg rt 328 ms	🔥	Global	0.00	9.00
6	03:06:07 PM	180.244.72.0/21	17974	SERVERIUS-PTF1	NLIX	🕒	🛡️	Old loss 0%, avg rt 200 ms, new loss 0%, avg rt 213 ms	🔥	Global	0.00	9.00
7	03:03:59 PM	202.128.124.0/24	38313	COGENT	NLIX	🕒	🛡️	Old loss 0%, avg rt 181 ms, new loss 0%, avg rt 187 ms	🔥	Global	0.06	9.00
8	03:06:07 PM	282.44.232.0/21	4618	COGENT	NLIX	🕒	🛡️	Old loss 0%, avg rt 320 ms, new loss 0%, avg rt 245 ms	🔥	Global	0.02	9.00
9	03:06:27 PM	213.57.132.0/24	12849	COGENT	NLIX	🕒	🛡️	Old loss 100%, avg rt 0 ms, new loss 0%, avg rt 71 ms	🔥	Global	0.00	9.00
10	03:06:12 PM	213.57.84.0/24 🕒	12849	COGENT	NLIX	🕒	🛡️	Old loss 55%, avg rt 87 ms, new loss 0%, avg rt 78 ms	🔥	Global	1.05	9.00
11	03:05:41 PM	218.156.0.0/16	4766	COGENT	NLIX	🕒	🛡️	Old loss 0%, avg rt 317 ms, new loss 0%, avg rt 294 ms	🔥	Global	0.00	9.00
12	03:08:06 PM	38.77.240.0/20	17974	SERVERIUS-PTF1	NLIX	🕒	🛡️	Old loss 0%, avg rt 228 ms, new loss 0%, avg rt 226 ms	🔥	Global	0.89	1.60

Figure 3.4.13: Improvements to Exchanges

3.4.8 Historical records

The “Historical records” report contains a complete list of active and inactive improvements that were made by the system. The report can be sorted using any of the columns by simply clicking on the column

header.

#	Timestamp	Prefix	ASN	AS Name	Old			New			Type	Prob
					Provider	Loss %	Latency (avg, ms)	Provider	Loss %	Latency (avg, ms)		
1	11/11/13 12:01:28 PM	77.52.96.0/19	21497	PrJSC MTS ...	IX2	100	0	IX2	0	1	🔥	🔄
2	11/11/13 11:36:52 AM	77.52.96.0/19	21497	PrJSC MTS ...	A	0	283	IX2	0	1	🔥	🔄
3	11/11/13 09:35:32 AM	213.160.192.0/19	13238	Yandex LLC	A	0	46	IX2	0	1	🔥	🔄
4	11/11/13 04:21:31 AM	77.52.96.0/19	21497	PrJSC MTS ...	B	0	147	IX2	0	1	🔥	🔄
5	11/11/13 03:29:51 AM	209.65.128.0/19	15169	Google Inc...	B	0	94	IX2	0	1	🔥	🔄
6	11/10/13 07:21:06 PM	77.52.96.0/19	21497	PrJSC MTS ...	A	0	161	IX2	0	1	🔥	🔄
7	11/10/13 06:41:20 PM	77.52.96.0/19	21497	PrJSC MTS ...	IX2	100	0	IX2	0	1	🔥	🔄
8	11/10/13 05:41:37 PM	77.52.96.0/19	21497	PrJSC MTS ...	A	0	112	IX2	0	1	🔥	🔄
9	11/10/13 04:21:26 PM	77.52.96.0/19	21497	PrJSC MTS ...	IX2	100	0	IX2	0	1	🔥	🔄
10	11/10/13 04:01:09 PM	77.52.96.0/19	21497	PrJSC MTS ...	A	0	141	IX2	0	1	🔥	🔄
11	11/10/13 02:58:13 PM	83.237.96.0/19	8359	MTS MTS OJ...	B	0	55	IX2	0	1	🔥	🔄
12	11/10/13 02:41:24 PM	77.52.96.0/19	21497	PrJSC MTS ...	IX2	100	0	IX2	0	1	🔥	🔄

Figure 3.4.14: Historical records

📘 In some specific cases, the “From” and “To” fields will contain the same provider. This is not a platform issue, and is related to the Outage detection algorithm specifics. Please see the [Outage detection](#) section for more details

3.4.9 Providers overall

The “Providers Overall” report provides the number of prefixes that were rerouted from and to each of the providers. In the example below, the traffic exchange between the providers is balanced. However, if the difference between the outgoing and incoming traffic is significant, then the quality of that provider’s network should be questioned. The report also shows the number of improved prefixes, based upon performance and cost, for each of the providers.



Figure 3.4.15: Providers overall

3.4.10 Provider performance

During the day IRP makes many probes to determine performance characteristics of alternatives routes.

Provider performance report presents an aggregated view of this measurements for providers. The data on the report highlights Packet loss and Average latency per provider putting the provider with best packet loss data at the top.



Figure 3.4.16: Provider performance

Provider performance report has been extended and now also offers a daily performance view. The extended Provider daily performance report adds the following capabilities:

- a Best (hypothetical) provider that estimates what might be possible to achieve on your network in terms of loss and average latency,
- classification of destinations by distance from your network with the option to review a category of interest,
- filtering by date range allowing review of past data or average performance over longer time intervals,
- a table view with exact details of the data in the report.




Figure 3.4.17: Provider daily performance

Note that while some providers seem better performing (for example, NLIX or AMS-IX in the provided screen capture) the fact that they display a much smaller number of probes clearly indicates they are Internet Exchanges with only a few peers interchanging data with our network.

For a long term distribution of this data refer chart [Provider performance history](#).

3.4.11 Cost overview

 This report is available only if the system is running in Cost optimization mode (see [Cost optimization](#)).

The “Cost Overview” report provides details regarding the cost savings, resulted from running IRP for each provider. It displays the costs that would be normally paid to the providers and the costs incurred after IRP has optimized the routing. The system calculates the total transit cost while running in Cost improvement mode and compares it with the total transit cost that would be incurred without running IRP. This difference is shown on the right side of the table as Cost Savings. The total cost savings include the ones made by the Commit Control as well. If the provider’s billing periods differ, then it is possible to choose a billing period for each provider separately.








Provider	Billing period	Cost (USD)	Default *		Improved by IRP*		Cost savings (USD)
			Billable throughput ** (Mbps) 	Cost(USD)	Billable throughput ** (Mbps) 	Cost (USD)	
A	11/01/13 - 12/01/13 	3.00	121	363	119	357	6
B		3.00	47	141	46	138	3
IX2		0	0	0	0	0	0
			168	504	165	495	9


Figure 3.4.18: Cost overview

3.4.12 Congestion and outages

This report is available only if [Outage detection](#) is enabled.

The “Congestion and outages” report contains all the as-patterns defined as problematic by the system. The problem state is shown in the “Status” column, as follows:

-  - A problem on the specified AS-Pattern is suspected
-  - The problem was not confirmed by additional re-probing
-  - The problem was confirmed, and all related prefixes were rerouted
-  - The problem was confirmed, but no better-performing route has been found

 In the past the report was named “AS pattern problems”.


#	AS-Pattern	Problem	Detection time	Affected prefixes						Problem status	
				Total	Selected for probing	Probed	Pending probing	Problem confirmed	Problem not confirmed	Confirmation rate	Status
1	2914-20377	Congestion	11/11/13 12:21:23 PM	51	51	51	0	0	51	0%	
2	9198-29355	Congestion	11/11/13 12:13:40 PM	57	56	56	0	4	52	7.02%	
3	12956-27699	Congestion	11/11/13 12:05:05 PM	111	111	111	0	0	111	0%	
4	7470-9287	Congestion	11/11/13 11:56:49 AM	2	2	2	0	1	1	50%	
5	12389-43975	Congestion	11/11/13 11:10:18 AM	149	149	149	0	3	146	2.01%	
6	1299-4788	Congestion	11/11/13 11:16:56 AM	14	14	14	0	0	14	0%	
7	9110-12400	Congestion	11/11/13 11:13:06 AM	53	53	53	0	6	45	15.09%	
8	174-12874	Congestion	11/11/13 09:21:45 AM	46	46	46	0	2	44	4.35%	
9	12389-43975	Congestion	11/11/13 09:06:41 AM	148	147	147	0	3	144	2.03%	
10	7473-4804	Congestion	11/11/13 09:01:42 AM	43	43	43	0	3	42	6.67%	
11	21219-28926	Congestion	11/11/13 08:59:31 AM	16	16	16	0	0	16	0%	
12	24708-13194	Outage	11/11/13 08:23:43 AM	37	37	37	0	1	36	2.7%	

Figure 3.4.19: Congestion and outages

3.4.13 Top volume prefixes

The “Top volume prefixes” displays a list of prefixes sorted by the total volume transferred to these remote networks.

#	Prefix	Volume(MB) 
1	146.120.112.0/23	307639.45
2	82.145.220.0/22	110744.63
3	66.249.78.0/24	96863.82
4	192.187.96.0/19	93099.26
5	217.69.128.0/20	73946.92
6	188.134.32.0/19	64603.56
7	213.87.136.0/21	54826.76
8	5.10.64.0/19	49511.55
9	188.123.224.0/19	48647.03
10	91.221.66.0/23	47962.48
11	178.124.172.0/24	45565.62
12	109.207.0.0/20	39653.37

Figure 3.4.20: Top volume prefixes

3.4.14 Top volume AS

The “Top Volume AS” reports show the Autonomous Systems with the highest volume of traffic coming from the monitored network. It can help you understand your traffic flow directions. If according to this report a specific AS has significantly more traffic than the others, then the network operator may consider getting a direct connection to this AS, to reduce traffic costs.

Top volume AS				
#	ASN	AS Name	95th (Mbps)	Volume (GB)
141	14420	CORPORACION NACIONAL DE TELECOMUNICACIONES - CNT EP	390.71	0.33
142	37611	Afrihost	1271.44	0.32
143	50266	T-Mobile Thuis BV	0.66	0.32
144	26615	Tim Celular S.A.	0.00	0.32
145	37963	Hangzhou Alibaba Advertising Co.,Ltd. East Software Park, 99 Huaxing Rd. Hangzhou, Ch	190.29	0.31
146	39309	Edutel BV	2.04	0.30
147	3265	Xs4all Internet BV	225.35	0.30
148	59050	Beijing Cloud-Ark Technology Co.,Ltd.	182.05	0.29
149	23650	AS Number for CHINANET jiangsu province backbone	299.18	0.26
150	7600	Escape.net	144.84	0.26
151	45090	Shenzhen Tencent Computer Systems Company Limited Tencent Building, Kejizhongyi Av	152.31	0.26
152	1136	KPN B.V.	396.24	0.25

Figure 3.4.21: Top volume AS

3.4.15 Top problem AS

The “Top Problem AS” report reveals the most problematic Autonomous Systems that the monitored networks are sending traffic to. It does that by showing how many times the traffic going to that AS was rerouted by the IRP.

#	ASN	AS Name	Problems
1	25577	Connexions4London Ltd	3
2	6849	JSC UKRTELECOM,	2
3	48503	Tele2 Sverige AB	1
4	44640	Cactus Ltd.	1
5	21151	Ukrcom Ltd	1
6	35377	A.B.N. JSC	1
7	25592	NETIS TELECOM Inc. Yaroslavl region ISP provider Russia	1
8	34829	LLC Nauka-Svyaz	1
9	15557	Societe Francaise du Radiotelephone S.A	1
10	50780	EAST.NET Ltd.	1
11	42945	Limited liability company Insit	1
12	51004	Sakhalin Cable Telesystems Ltd	1

Figure 3.4.22: Top problem AS

3.4.16 Prefix statistics

This report lists specific prefixes with relevant monthly values such as traffic volume and number of improvements or current unique and top hosts values. To see top hosts expand the + button to expand to latest individual hosts and their corresponding traffic volume. Note that the upper limit for the number of top hosts is constrained by `collector.export.top_volume_ips`.








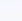






Prefix Statistics										
August, 2017  (switch to map view)										
#	Month	Prefix	ASN	AS Name	Country	Volume (GB) 	Improvements	Policy	Unique IPs	Top Hosts
1	Aug. 2017	89.45.8.0/24	54103	MOD Mission Critical		5033.97	0		2	
2	Aug. 2017	42.112.10.0/24	18403	The Corporation for Finan	VN	70.85	1		5	
3	Aug. 2017	118.107.176.0/24	64050	BGPNET Global ASN	SG	67.28	2		7	
4	Aug. 2017	88.198.0.0/16	24940	Hetzner Online GmbH		65.82	1		4	
#	Timestamp	IP Address			Volume (KB) 			Rank		
1	20 seconds ago	88.198.68.8			561.24			1		
2	20 seconds ago	88.198.68.23			19.73			2		
3	20 seconds ago	88.198.25.3			1.91			3		
5	Aug. 2017	112.121.190.0/24	45753	NETSEC NOC	HK	64.57	15		1	
6	Aug. 2017	93.93.96.0/23	44054	Media Network Services		61.08	4		2	
7	Aug. 2017	116.251.224.0/19	133771	Rapid Shield Company Li		60.34	7		1	
8	Aug. 2017	179.0.12.0/23	52472	FOCUS EL SALVADOR S	SV	60.01	0		2	
9	Aug. 2017	93.93.98.0/23	44054	Media Network Services		57.74	1		2	
10	Aug. 2017	96.44.188.0/22	8100	QuadraNet, Inc	US	55.84	16		6	
11	Aug. 2017	104.223.76.0/23	8100	QuadraNet, Inc	US	55.54	4		6	
12	Aug. 2017	103.9.105.0/24	58599	Cybergate Limited	BD	54.89	1		1	

Figure 3.4.23: Prefix statistics

Note that Prefix statistics report introduced a map view highlighting with heat-maps those regions of the world where either most traffic or highest number of improvements are made.

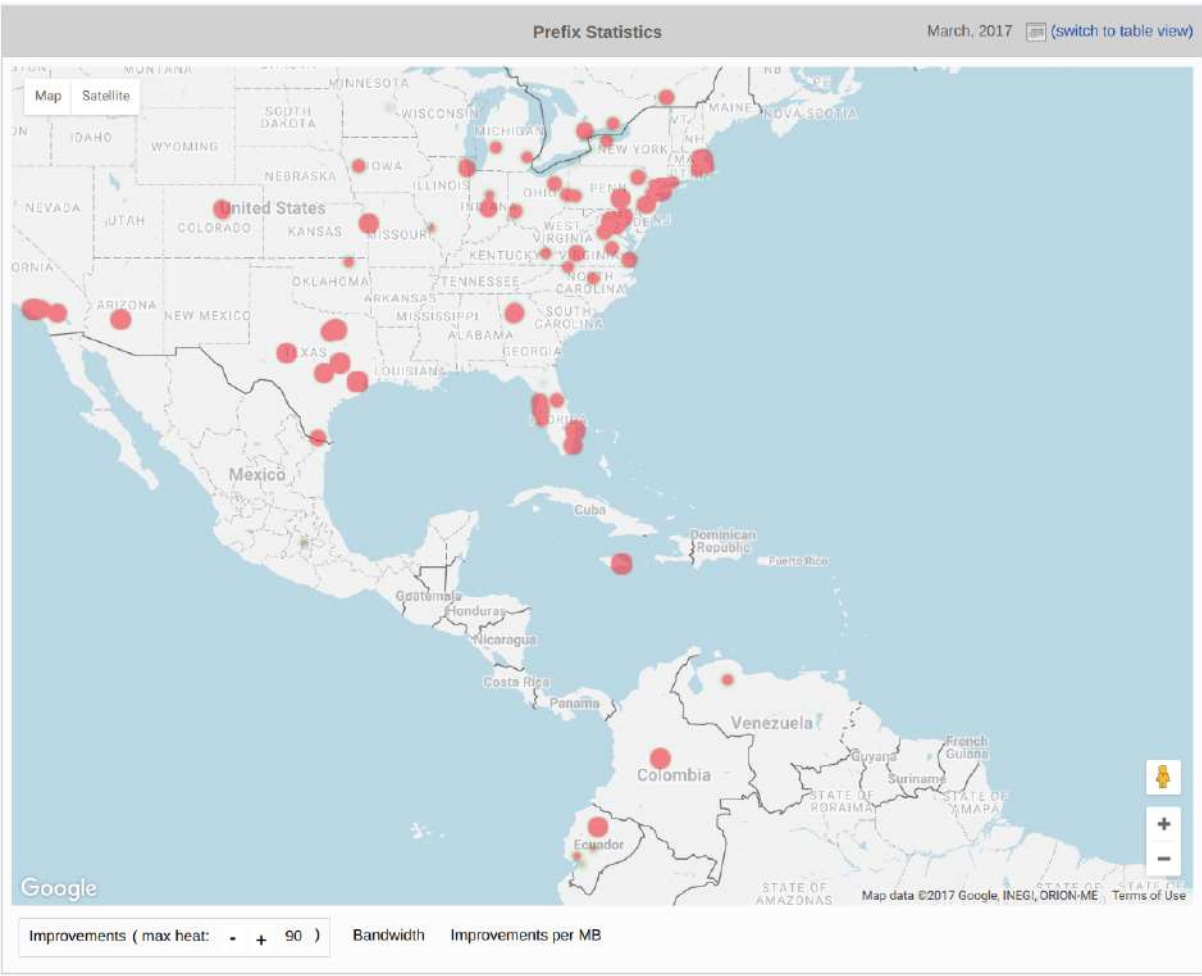


Figure 3.4.24: Prefix statistics map view

3.4.17 AS statistics

The “AS statistics” report aggregates by Autonomous System the total number of improvements, the volume of traffic and the 95th for a selected past month. Note that the choice of month for the 95th and the time period of the report are independent and they can set non-overlapping ranges.

AS statistics						
#	Date	ASN	AS Name	95th July, 2017	Volume (GB)	Improvements
81	2017-11-01	24444	Shandong Mobile Communication Company	0.26	1.69	10800
82	2017-10-26	4808	China Unicom Beijing Province Network	0.14	0.29	10797
83	2017-03-20	4847	China Networks Inter-Exchange	0.51	0.00	10772
84	2017-11-15	23724	IDC. China Telecommunications Corporation	0.31	0.10	10756
85	2017-10-21	3651	Sprint	0.41	0.15	10744
86	2017-05-31	4847	China Networks Inter-Exchange	0.51	0.00	10693
87	2017-10-21	4837	CHINA UNICOM China169 Backbone	0.21	0.06	10672
88	2017-11-12	4847	China Networks Inter-Exchange	0.51	4.85	10662
89	2017-10-29	24444	Shandong Mobile Communication Company	0.26	0.54	10608
90	2017-11-13	23724	IDC. China Telecommunications Corporation	0.31	0.21	10576
91	2017-11-11	701	MCI Communications Services, Inc. d/b/a Ver	0.34	3.90	10568
92	2017-10-31	4837	CHINA UNICOM China169 Backbone	0.21	0.17	10566
93	2017-04-06	17638	ASN for TIANJIN Provincial Net of CT	0.23	0.00	10545

Figure 3.4.25: AS statistics

3.4.18 Country statistics

The “Country statistics” report distributes traffic volume and number of improvements by country. It helps see to what regions most of a network’s traffic is addressed and also the relative number of suboptimal routes towards them that IRP identified and addressed by injecting improvements.

Country Statistics					(switch to map view)
#	Country Name	Country	Volume (GB)	Improvements	
1	China	CN	646.05	452027	
2	United States	US	4567.72	171039	
3	India	IN	171.92	98173	
4	South Africa	ZA	36.26	44059	
5	Australia	AU	63.19	39997	
6	Rest of the world	--	5295.77	9296	
7	Colombia	CO	10.18	8749	
8	Ecuador	EC	2.77	5894	
9	Jamaica	JM	3.49	4786	
10	Singapore	SG	268.79	3227	
11	Japan	JP	1.42	1144	

Figure 3.4.26: Country statistics

Country statistics report also displays a map view highlighting countries with most/least number of improvements.

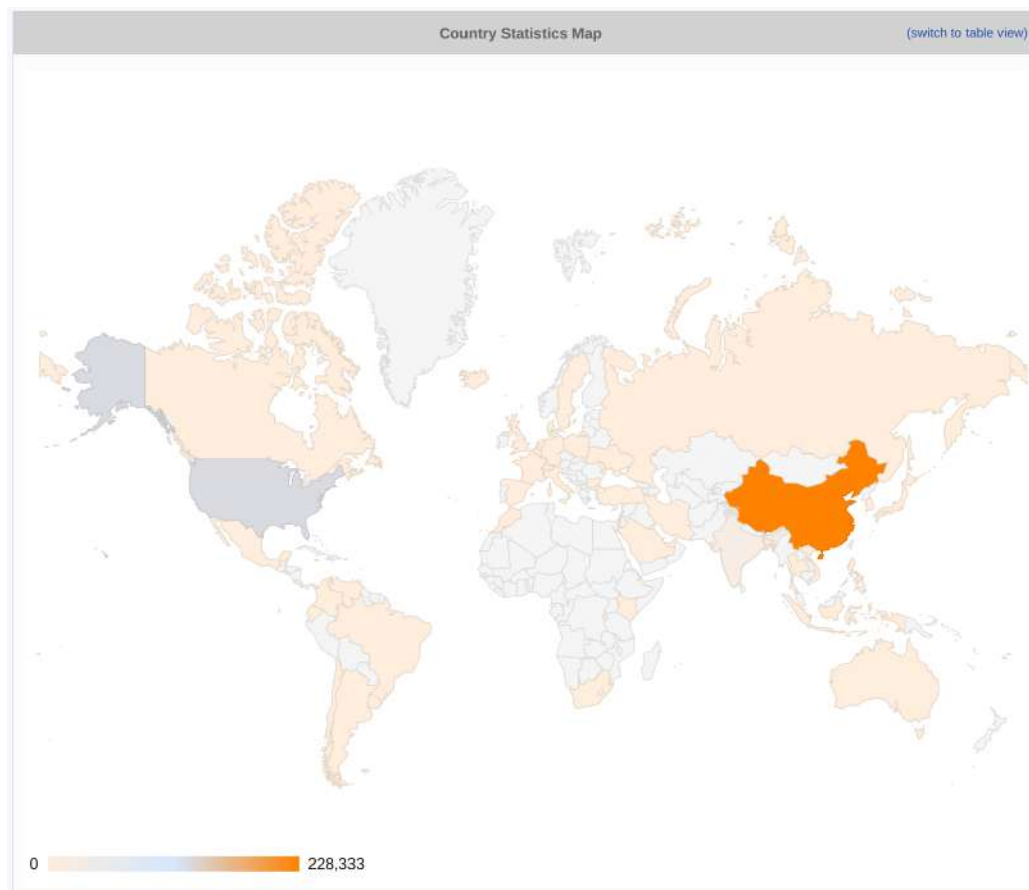


Figure 3.4.27: Country statistics map view

3.4.19 Exchanges statistics

The “Exchanges statistics” report lists statistics regarding communications with peers on Internet Exchanges. Filters and sorting help identify the required details.

i In the past this report was named “Bandwidth by Peer”

#	Timestamp	Peering partner	Provider	IP Version	Volume Outbound (MB) ↗	Volume Inbound (MB) ↗	Packets Outbound	Packets Inbound
1	Jun. 2017	80.249.208.90	AMS-IX	4	368.70	0.00	156450	0
2	Jun. 2017	80.249.212.146	AMS-IX	4	104.11	0.00	8648	0
3	Jun. 2017	80.249.209.150	AMS-IX	4	62.79	0.00	52525	0
4	Jun. 2017	193.239.117.165	NLIX	4	16.31	0.00	8810	0
5	Jun. 2017	193.239.117.234	NLIX	4	12.71	0.00	7098	0
6	Jun. 2017	193.239.117.110	NLIX	4	11.92	0.00	15339	0
7	Jun. 2017	80.249.208.247	AMS-IX	4	10.15	0.00	10213	0
8	Jun. 2017	80.249.210.125	AMS-IX	4	8.86	0.00	5599	0
9	Jun. 2017	193.239.116.25	NLIX	4	7.76	0.00	4305	0
10	Jun. 2017	80.249.208.200	AMS-IX	4	7.72	0.00	5835	0
11	Jun. 2017	80.249.209.20	AMS-IX	4	7.02	0.00	5217	0
12	Jun. 2017	80.249.208.50	AMS-IX	4	3.89	0.00	2177022	0

Figure 3.4.28: Exchanges statistics

3.4.20 Current inbound improvements

The “Current inbound improvements” report provides a list of the currently active improvements. A drill-down capability to a specific ASN or prefix is also provided. The report shows the improved prefix and the providers from and to which the traffic was redirected. It also provides the reason, which the improvement was based upon, along with the performance values before and after the traffic was rerouted. If necessary, specific improvements can be manually deleted by clicking the check-box next to the needed prefix, afterward using the “Remove selected” button.

Current Inbound Improvements						
#	Timestamp	Link1	Link2	Prefix	Type	<input type="checkbox"/>
1	03/04/16 05:49:09 AM	-	-	10.1.7.0/24		<input type="checkbox"/>
2	03/04/16 04:14:10 AM	-	-	10.1.3.0/24		<input type="checkbox"/>
3	03/04/16 03:58:34 AM	2	-	10.1.8.0/24		<input type="checkbox"/>
4	03/04/16 03:23:08 AM	-	2	10.1.6.0/24		<input type="checkbox"/>
5	03/04/16 02:01:42 AM	-	2	10.1.9.0/24		<input type="checkbox"/>
6	03/04/16 02:01:34 AM	1	-	10.1.4.0/24		<input type="checkbox"/>
7	03/04/16 01:34:11 AM	2	-	10.1.2.0/24		<input type="checkbox"/>
8	03/04/16 01:30:56 AM	1	-	10.1.1.0/24		<input type="checkbox"/>

1 - 14 of 14 results found

Figure 3.4.29: Current inbound improvements

Refer [Inbound bandwidth optimization](#) for details.

3.4.21 Inbound Improvements history

The “Inbound Improvements history” report provides a detailed list of past inbound improvement actions. Filtering for example by specific prefix or improvement type is provided. The report shows the improved prefix and the new prepends count towards a provider.

Inbound Improvements History							
#	Timestamp	Prefix	IP Version	Provider	Prepends	Type	Action
1	52 min 38 seconds ago	149.11.0.0/16	4	Cogent	-		Deleted
2	2017-06-22 16:13:06	185.34.201.0/24	4	Cogent	3		Updated
3	2017-06-22 15:28:30	185.34.200.0/24	4	Serverius	2		Updated
4	2017-06-22 15:24:58	185.34.202.0/24	4	Serverius	3		Updated
5	2017-06-22 14:05:10	185.34.202.0/24	4	Serverius	4		Updated
6	2017-06-22 13:21:37	185.34.200.0/24	4	Serverius	3		Updated
7	2017-06-22 11:00:21	185.34.200.0/24	4	Serverius	2		Updated
8	2017-06-22 11:00:21	185.34.202.0/24	4	Serverius	1		Updated

Figure 3.4.30: Inbound Improvements history

Refer [Inbound bandwidth optimization](#) for details.

Historical monitors										
#	Timestamp	Provider	Next Hop	Type	IP Version	PBR State	PBR Observed Hop	PBR Diag Hop	BGP Internal Monitor State	BGP External Monitor State
61	2017-05-16 05:01:01	Serverius		pbrTransit	6	Fail	2a00:1ca8:ad::2	2a00:1ca8:ad::1/128		
62	2017-05-16 04:59:44	Serverius		pbrTransit	6	Fail	2a00:1ca8::2:20	2a00:1ca8:ad::1/128		
63	2017-05-16 04:42:23	AMS-IX	80.249.208.173	pbrPartial		Fail				
64	2017-05-16 04:32:46	AMS-IX	80.249.209.150	bgpIntPartial					Ok	
65	2017-05-16 04:30:46	AMS-IX	80.249.211.239	bgpIntPartial					Ok	
66	2017-05-16 04:07:52	AMS-IX	80.249.209.20	pbrPartial		Ok				
67	2017-05-16 03:14:25	AMS-IX	80.249.211.95	pbrPartial		Ok				
68	2017-05-16 01:23:52	AMS-IX	80.249.211.95	pbrPartial		Fail				
69	2017-05-16 01:10:53	Serverius		pbrTransit	6	Fail	2a00:1ca8:ad::2	2a00:1ca8:ad::1/128		
70	2017-05-16 01:09:36	Serverius		pbrTransit	6	Fail	2a00:1ca8::2:20	2a00:1ca8:ad::1/128		
71	2017-05-16 01:06:51	Serverius		pbrTransit	6	Fail	2a00:1ca8:ad::2	2a00:1ca8:ad::1/128		
72	2017-05-16 01:05:34	Serverius		pbrTransit	6	Fail	2a00:1ca8::2:20	2a00:1ca8:ad::1/128		
73	2017-05-16 00:18:32	Serverius		pbrTransit	6	Fail	2a00:1ca8:ad::2	2a00:1ca8:ad::1/128		
74	2017-05-16 00:17:16	Serverius		pbrTransit	6	Fail	2a00:1ca8::2:20	2a00:1ca8:ad::1/128		
75	2017-05-16 00:12:21	AMS-IX	80.249.211.95	pbrPartial		Ok				
76	2017-05-15 23:52:26	AMS-IX	80.249.209.150	bgpIntPartial					Fail	

Figure 3.4.33: Monitor history

3.5 Graphs

The IRP graphs are visual representations of data found in the reports. They provide an opportunity to quickly identify problems in the network, as well as forecast periodic changes in the traffic patterns and network behavior. For each graph you can adjust the time period which you want to get the results for, and export them as PDF (see the [Sidebar](#) section).

Unless otherwise noted, all graphs are plotted using 5 minute intervals.



Figure 3.5.1: Graphs menu

3.5.1 Improvements by type

The “Improvements by type” graph shows the improvements made by IRP based upon: performance, cost, commit, and outage detection. Various problem patterns can be detected in compliance with the IRP activity shown in this graph. The graph provides information on the types of improvements that are more frequent for the network in the current state. Different patterns can be noticed depending on the improvement mode that is currently running.



Figure 3.5.2: Improvements by type

3.5.2 Performance improvements

The “Performance improvements” graph displays the prefixes improved based upon a performance reason. By following this graph, the solved loss and latency issues, occurring in the network, can be monitored.

The graph also provides the average, maximum and minimum number of improvements for each of the improvement reasons during the selected time frame. The maximum peaks are indicators of a loss or latency problem in the network.




Figure 3.5.3: Performance improvements

3.5.3 Probed prefixes and Improvements

The "Probed prefixes and Improvements" graph shows the number of improvements (Loss/Latency) per time unit as reported to the total probed prefixes.



Figure 3.5.4: Probed prefixes and Improvements

3.5.4 New and retry improvements

After a particular improvement was made, IRP analyzes the path periodically. If the improvement is still valid it leaves it as it is. The "New and retry improvements" graph provides the number of improvements made for the first time as related to those, which were simply confirmed after re-probing.




Figure 3.5.5: New and retry improvements

3.5.5 Improvements by probing source

The current graph displays the percentage of improvements that are differentiated by the probing source including Commit Control, Outage Detection, VIP Probing, Regular and Retry Probing.




Figure 3.5.6: Improvements by probing source

3.5.6 Prefixes rerouted from provider

The current graph displays the number of prefixes that were rerouted from a specific provider in order to resolve performance or cost issues.




Figure 3.5.7: Prefixes rerouted from provider

3.5.7 Prefixes rerouted to provider

The graph displays the number of prefixes that were rerouted to a specific provider in order to resolve performance or cost issues.




Figure 3.5.8: Prefixes rerouted to provider

3.5.8 Improvements by provider

The “Improvements by provider” graph shows all the improvements made by IRP for each of the providers. You can monitor the issues occurring in the network, which triggered IRP to reroute the traffic. You can compare the providers based on how they perform in the context of performance and cost. It also provides the average, maximum and minimum number of improvements per provider for each of the improvement reasons.

i Improvements by provider graph categorizes improvements by performance and cost. If cost mode is disabled this graph shows the same data as Prefixes rerouted to provider graph.




Figure 3.5.9: Improvements by provider

3.5.9 Performance improvements by provider

Similar to the previous graph, the “Performance improvements by provider” displays a graphical representation of historical values for all latency and loss-based improvements for each provider. The charts are displayed next to each other for easier comparison and identification of regularities. Keep in mind that specific provider lines can be removed from the graph. This way we are able to focus only on the interesting data.




Figure 3.5.10: Performance improvements by provider

3.5.10 Commit improvements by provider

If the Commit Control algorithm is enabled (see [Commit Control](#)), this graph will display an overview of all the commit improvements that were enforced by IRP for a specific time period.




Figure 3.5.11: Commit improvements by provider

3.5.11 Total and Improved traffic

The Total and Improved Traffic graph allows you to see the positive impact of IRP on the network. It provides the amount of improved traffic, shown in blue color, as related to the total traffic, shown in green. Thus, you can monitor how much of the traffic is affected by IRP.



Figure 3.5.12: Total and improved traffic

3.5.12 Bandwidth usage by provider

The Bandwidth Usage graph shows the total usage for each of the providers. It allows you to compare the current traffic volume to the current 95th percentile and the commit 95th. The graph provides average, maximum and minimum traffic usage values for each of the providers during the selected time period.

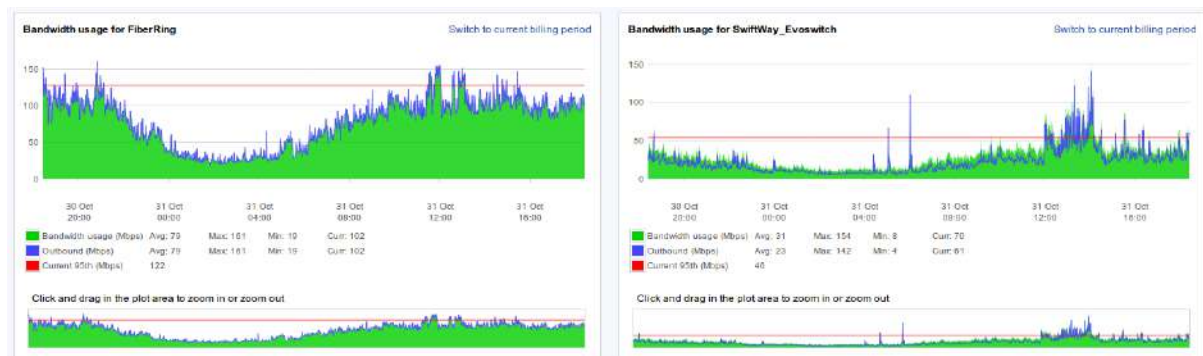


Figure 3.5.13: Bandwidth usage by provider

3.5.13 Providers bandwidth usage

The Providers bandwidth usage graph combines all usage lines in a single chart. It allows cross-comparison between providers themselves. The chart includes options to hide/show providers of interest during selected time period. The top and bottom charts display outbound and inbound traffic accordingly.




Figure 3.5.14: Providers bandwidth usage

3.5.14 Bandwidth usage and improvements

Bandwidth usage and improvements chart superimposes two charts for easy analysis and comparison of how improvement counts correlate with bandwidth.



Figure 3.5.15: Bandwidth usage and improvements

3.5.15 Total bandwidth usage




Figure 3.5.16: Total bandwidth usage

3.5.16 Inbound traffic distribution

The graph highlights distribution of inbound traffic across inbound prefixes and providers.

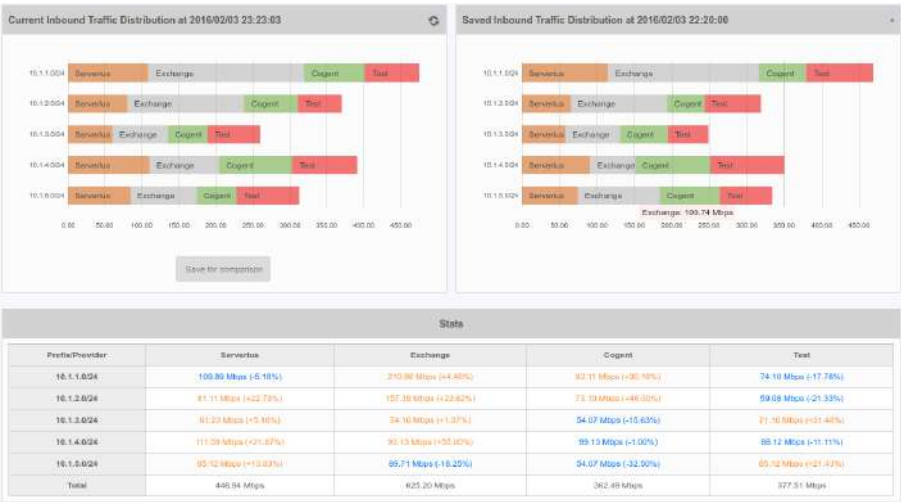


Figure 3.5.17: Inbound traffic distribution

Inbound traffic distribution report includes a feature to cross-compare current values with a saved inbound traffic distribution from recent past. To cross-compare saved inbound traffic distribution with freshly retrieved statistics, save and compare traffic shape changes and if they conform to expectations. Refer [Inbound bandwidth optimization](#) for details.

3.5.17 Provider performance history

Provider performance history uses a scatter chart to plot past provider performance data. Refer report [Provider performance](#) for a different way to display this data. While Provider daily performance report for a longer time period averages the individual metrics this chart allows spotting consistent or hectic performance for providers.

i Individual data points represent past daily averages. IRP needs to run for a few days for this chart to have any meaningful data to plot.

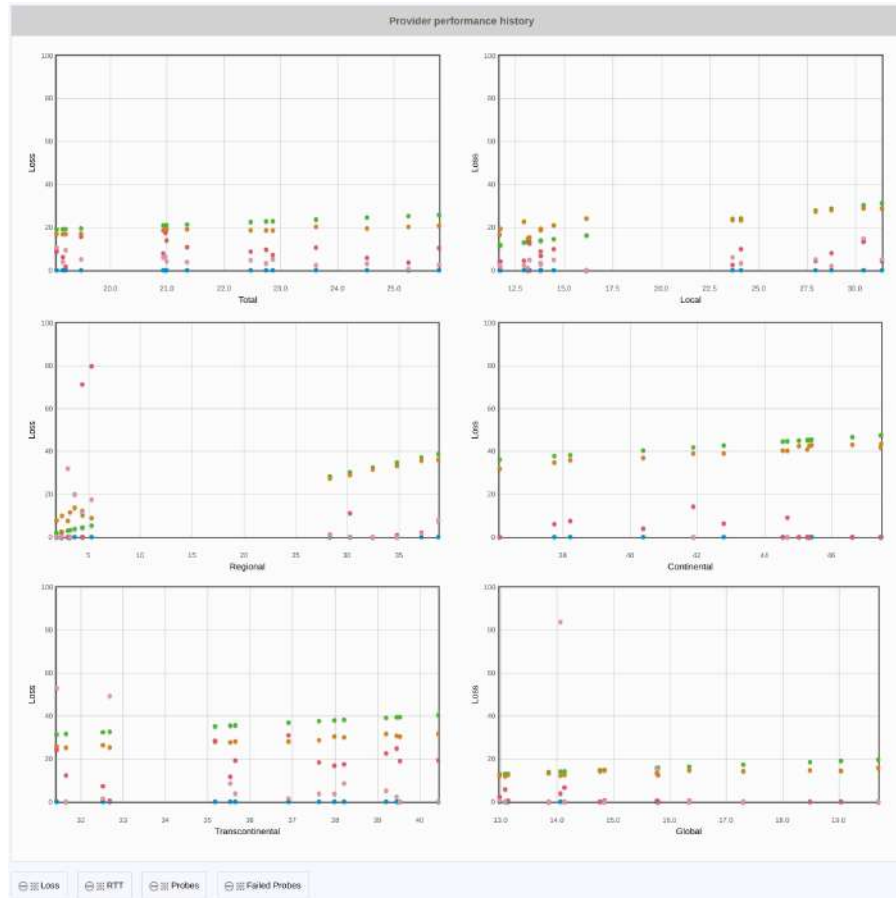


Figure 3.5.18: Provider performance history

Refer legend on the left-side sidebar for color points assigned to different providers. The same 'Best' category as in report [Provider performance](#) is distinguished to plot the absolute best loss or RTT observed during IRP probing. Designated charts plot local, regional, continental, transcontinental and worldwide data. The totals including all data points is available too.

3.6 Routing Policies

To see the full list of routing policies configured in the system, select the Routing Policies option from the main menu and click on the Routing Policies option in the submenu. The list displays all the enabled and disabled routing policies configured in IRP.

Priority	Prefix/ASN	Note	ASN	AS Name	VIP	Policy Type	Providers	Last probed	Next probe	From	To	Details	Actions
1	66.87.70.0/24 (1 prefixes)		10607	Sprint Personal Communications Systems		Allow through	NLIX, Serverius	15 minutes ago	-				
2	5.101.41.0/24 (1 prefixes)		44050	Holandburg Internet Network Ltd		Allow through	Cogent, Serverius	a few seconds ago	In next 4 minutes				
2	2.2.2.0/24 (1 prefixes)		3215	Orange		Static route	Cogent	-	-	Cogent	Cogent	Global Static route through Cogent[6]	
7	17623 (107 prefixes)		17623	China Unicom Shenzhen network		Allow through	AMS-IX, Cogent, NLIX, Serverius	17 minutes ago	In next 4 hours				
7	27.38.0.0/24		17623	China Unicom Shenzhen network		Allow through	AMS-IX, Cogent, NLIX, Serverius	17 minutes ago	In next 4 hours				
7	27.38.9.0/24		17623	China Unicom Shenzhen network		Allow through	AMS-IX, Cogent, NLIX, Serverius	17 minutes ago	In next 4 hours				

Figure 3.6.1: Routing Policies

As shown in the screenshot above, the list provides the following details:

- Priority of a policy.

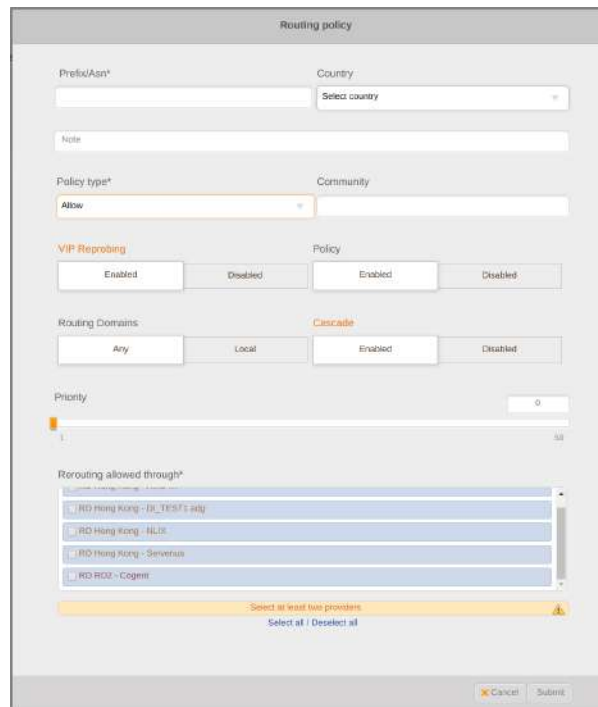
i Priority specifies what policy to use in case of overlapping prefixes as might be the case of a large aggregate prefix policy extending over an ASN policy

- The country, prefix or ASN to which the policy applies.
- A note to describe the policy.
- The specific ASN and AS Name of the prefix to which the policy was applied.
- The status of the VIP reprobing (enabled/disabled).
- The type of the policy (allow, deny, static, static_exact).
- Providers which are affected by the policy.
- Time passed after the last probe was performed towards this prefix.
- The time left before the next scheduled probe towards this prefix.
- From and To details of an existing improvement relating to this policy.
- The available actions for the policy. These include:

- On-demand probing
- Editing the routing policy
- Enabling the routing policy
- Disabling the routing policy

You can collapse all the policies by clicking the "Collapse All" button and then expand them back by clicking the "Expand all" button. You can use the "Remove all" button to delete all the policies. You can also select a set of policies and use the "Remove selected" button to delete them.

A new routing policy can be added by pressing the "Add routing policy" button. The following window will open up:



The image shows a 'Routing policy' configuration window. It contains several input fields and checkboxes. At the top, there are fields for 'Prefix/ASN*' and 'Country' (with a 'Select country' dropdown). Below these is a 'Note' text area. The 'Policy type*' is set to 'Allow' with a dropdown arrow. To its right is a 'Community' text field. Under 'VIP Reprising', there are 'Enabled' and 'Disabled' buttons. To the right, under 'Policy', are also 'Enabled' and 'Disabled' buttons. The 'Routing Domains' section has 'Any' and 'Local' buttons, followed by a 'Cascade' label and 'Enabled'/'Disabled' buttons. A 'Priority' slider is shown with a value of 0. At the bottom, there is a list titled 'Rerouting allowed through*' with four entries: 'RD Hong Kong - HKT-TRANS-AS', 'RD Hong Kong - HKT', 'RD Hong Kong - Serinet', and 'RD RIX - Cogent'. Below the list are buttons for 'Select all / Deselect all'. At the very bottom right are 'Cancel' and 'Submit' buttons.

Figure 3.6.2: Routing policy window

The following parameters should be configured to add a routing policy:

- Prefix/ASN: The prefix/ASN to which the policy is applied
- Country: The country to which the policy is applied. Such a policy applies to all prefixes known to be located in the country.
- Note: A note to describe the policy
- Policy Type: The type of the policy (allow, deny, static, static_exact)
- A community to mark improvements
- A flag to set if policy is VIP
- The policy status (enabled/disabled)
- A flag to indicate if an ASN policy should propagate to downstream ASN.
- A flag indicating if a global improvement allowed.
- Priority slider or text-box will set a policy's priority.
- The Providers for which the policy should be respected
- Cascade flag for ASN policies that indicate if the policy should be enforced only for this AS or also for downstream AS. If cascading is enabled it ensures that the policy is enforced for all traffic that will eventually be routed through this AS.

Refer policy configuration parameters [Routing Policy](#) for details.

To see a list of Static Route policies, select Static Routes in the Routing Policies submenu. Alternatively the filter on the left side-bar can be used. A list like the one below will be displayed.

Add static route													Collapse all		Expand all	Remove all	Remove selected
Priority	Prefix/ASN	Note	ASN	AS Name	VIP	Policy Type	Providers	Last prebid	Next probe	From	To	Details	Actions				
2	222.027 (1 prefixes)		3215	Orange		Static route	Cogent	-	-	Cogent	Cogent	Global Static route through Cogent(4)					

Figure 3.6.3: Static Routes

The list structure and the actions available are similar to those found in the Routing Policies list. Here you can also add a new Static Route by using the "Add static route" button.

You can select the VIP option in the Routing Policies submenu to access a list of all the policies with VIP reprobing enabled.

Priority	Prefix/ASN	Note	ASN	AS Name	VIP	Policy Type	Providers	Last probed	Next probe	From	To	Details	Actions
1	66.67.70.0/24 (1 prefixes)		10607	Sprint Personal Communications Systems		Allow through	NLIX, Serverius	3 minutes ago	in next 2 minutes				
2	5.301.41.0/24 (1 prefixes)		44050	Petersburg Internet Network Ltd.		Allow through	Cogent, Serverius	4 minutes ago	in next a minute				

Figure 3.6.4: VIP policies

Here you can also enable VIP reprobing for a new prefix/ASN by using the "Add VIP" button.

3.7 Flowspec policies

To review Flowspec policies, select the Routing Policies option from the main menu and click on the Flowspec Policies option in the submenu.

Flowspec must be enabled for Bgpd to announce them.

The list displays all the enabled and disabled Flowspec policies configured in IRP.

Note	Source prefix	Source port	Destination prefix	Destination port	Rule action	Redirect IP	Protocol	Provider	Rate (Mbps)	Actions
Throttle torrent		6881-6889		6881-6889	Throttle				100	
		4450			Throttle		TCP		100	
	185.34.222.0/24				Redirect IP	1.2.3.4				
	185.34.200.0/24				Redirect			Cogent		
Drop IPv6 from Oz	2a14:5ec0::431:96				Drop					

Figure 3.7.1: Flowspec policies

As shown in the screenshot above, the list highlights:

- A note to describe the policy.
- A source ASN/prefix and port(s) for matching packets.
- A destination prefix and port(s) for matching packets.
- DSCP traffic classification value.
- The policy action (Throttle, Drop, Redirect or Redirect IP).
- Redirect IP or Provider for redirect policies.
- Protocols of matching packets, for example TCP, UDP or ICMP.
- Rate limits in Mbps for Throttle policies.

Flowspec policies like routing policies can be edited, enabled, disabled or deleted altogether.

A Flowspec policy is added by clicking on the designated button and specifying the following:

Figure 3.7.2: Flowspec policy popup

The following parameters should be configured to add a routing policy:

- Notes to describe the policy
- Status to set if policy is Enabled or Disabled
- Source ASN/Prefix/Port(s): The source ASN/prefix and port(s) of the IP packets that match. A prefix in CIDR notation or a single IP address should be provided. Multiple valid TCP/UDP ports can be provided as well as port ranges.

i One Flowspec rule is created for each of the prefixes in the designated AS. ASN Flowspec policies can be setup only on the source of IP packets.

- Ensure that routers are capable of processing a large number of Flowspec rules before setting policies for an AS consisting of a very large number of network segments (prefixes).

- Destination Prefix/Port: The destination prefix/port attribute of the IP packets that match. Same rules as for Source Prefix/Port(s) apply.

- Prefixes **MUST** not equal or include probing IP addresses. Refer to for example `global.master_probing_interface`, `peer.X.ipv4.master_probing`, `peer.X.ipv6.master_probing`.

- Protocols: Packet protocols that match the policy. Can be any combination of TCP, UDP and ICMP.
- Policy Type: The type of the policy (Throttle, Drop or Redirect)

- Provider specifies one of the provider identifiers where traffic will be redirected. The provider is set only for Redirect policies.

i In order for Redirect policies to be implemented on routers a community value must be assigned to them and VRF must be configured on the router itself. Consult Noction support for details.

- Rate limits the allowed bandwidth usage for matching traffic. The value is set only for Throttling policies. The rate specifies a number in the range 1-4200 Mbps.

i In case the provided Flowspec policy attributes are incomplete or invalid on attempting to submit it a warning will be raised similarly as depicted in the popup.

For further details refer Flowspec configuration parameters for example: `global.flowspec`, `core.flowspec.max`, `core.flowspec.max_ipv6`, `bgpd.peer.X.flowspec`, `peer.X.flowspec.ipv4.redirect_community`, `peer.X.flowspec.ipv6.redirect_community`.

3.8 Throttling excessive bandwidth use with Flowspec

Throttling excessive bandwidth use with Flowspec policies use the same Flowspec policies reports. Select the Routing Policies option from the main menu and click on the Flowspec Policies option in the submenu.

i Flowspec must be enabled for Bgpd to announce them.

The list displays besides all the other Flowspec policies the automatic throttling rules marked with Throttle at. These rules cannot be edited.

Note	Source	Source port(s)	Destination	Destination port(s)	Rule action	Redirect IP	Protocol	Provider	Rate (Mbps)	Announced prefixes	Actions	
	60238	1.2.4.5.6.7	185.34.200.0/24		Throttle		ICMP, TCP, UDP		100	0	 	
	18002				Drop					0	 	
Overusage			185.34.201.0/24		Throttle at				101	0	 	
Overusage			185.34.200.0/24		Throttle at				45	0	 	
Overusage			185.34.202.0/24		Throttle at				300	0	 	
Overusage			185.34.212.0/24		Throttle at				15	0	 	

Figure 3.8.1: Throttling excessive bandwidth use with Flowspec policies

Each of the rules sets a specific destination prefix to control outbound bandwidth use and sets an Mbps rate based on past average bandwidth use.

Refer [Core configuration](#) for configuration details.

3.9 Maintenance windows

To review Maintenance windows, navigate to Policies option in main menu and click on Maintenance Window option. The list displays all the current and future maintenance windows configured in IRP. Refer [Maintenance windows](#) for details about maintenance windows.



	Begin	End	Providers	Withdraw Improvements	Seconds to Unload	Seconds to Prepend	Actions
●	2016-09-14 16:08	2016-09-14 16:30	AMS-IX	Yes	300	0	[Edit] [Delete] [Add]
●	2016-09-15 00:00	2016-09-28 00:00	Cogent	Yes	300	150	[Edit] [Delete] [Add]
●	2016-09-23 10:48	2016-09-30 10:48	NL-IX	Yes	0	0	[Edit] [Delete] [Add]

Figure 3.9.1: Maintenance windows

As shown in the screen-shot above, the list highlights:

- Current color-coded status of the maintenance window where Red and Orange depict unloading and actual maintenance phases while Blue and Green depict future planned windows and past finished windows correspondingly.
- Begin and End date-times of maintenance windows.
- Provider or router under maintenance.
- Withdraw Improvement option enables or disables withdrawal of existing Outbound Improvements to that provider.
- Seconds to Unload defines time interval in seconds prior to beginning of the maintenance window when IRP starts to unload outbound traffic from provider.

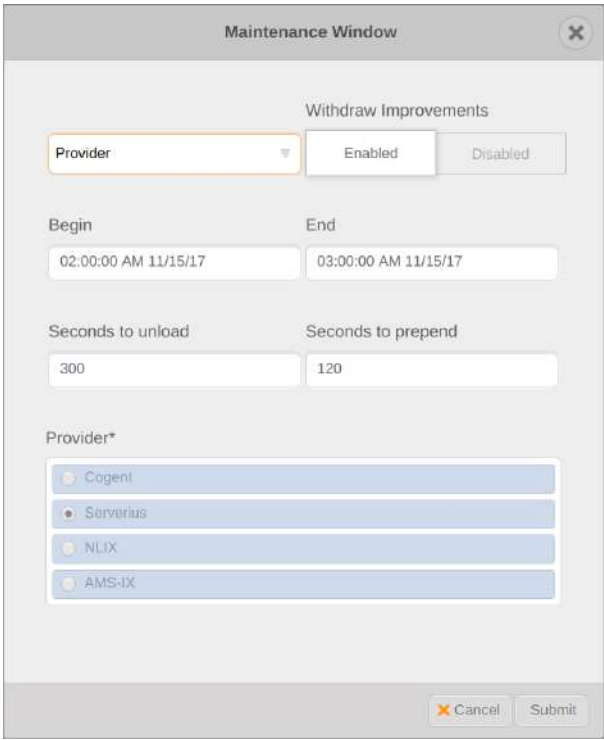
⚠ If Seconds to Unload/Seconds to Prepend time intervals set to zero IRP does not perform corresponding action.

ℹ When seconds to unload is specified it should be at least 5 minutes (300 seconds) long with larger time intervals allowing IRP more time to assess and reroute traffic flowing through provider with scheduled maintenance window. Very short time intervals might not have the desired effect since usually an IRP optimization cycle takes 2-3 minutes to finish.

- Seconds to Prepend defines time interval in seconds prior to beginning of the maintenance window when IRP announces all inbound prefixes with maximum prepends to provider in order to deflect inbound traffic towards other providers.

Maintenance windows can be added, edited or deleted using the designated action buttons.

A maintenance window is added by specifying the following:

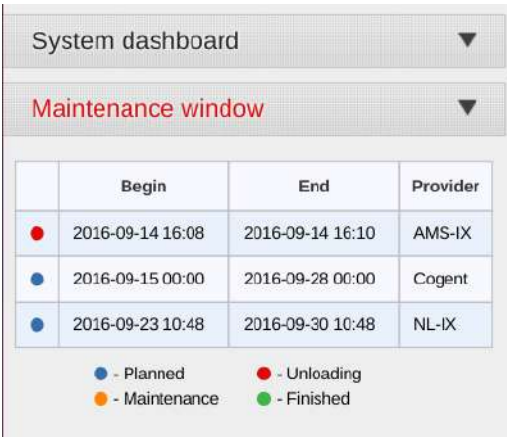


The image shows a 'Maintenance Window' configuration popup. It has a title bar with a close button. The main content area includes a 'Withdraw Improvements' section with a 'Provider' dropdown menu and 'Enabled'/'Disabled' toggle buttons. Below this are 'Begin' and 'End' time pickers. Further down are 'Seconds to unload' and 'Seconds to prepend' input fields. A 'Provider*' section contains a list of providers: Cogent, Serverius, NLIX, and AMS-IX, each with a radio button. At the bottom right are 'Cancel' and 'Submit' buttons.

Figure 3.9.2: Maintenance window popup

The parameters mentioned above including the option to choose a single provider or all providers on a router are provided for Maintenance window configuration.

Maintenance windows are also highlighted on IRP Frontend sidebar with the title of the sidebar box highlighting the status of the next maintenance window so that it is visible even if the contents of the sidebar is collapsed.



The image shows a sidebar with two expandable sections: 'System dashboard' and 'Maintenance window'. The 'Maintenance window' section is expanded and contains a table of maintenance events. Below the table is a legend for the event status icons.

	Begin	End	Provider
●	2016-09-14 16:08	2016-09-14 16:10	AMS-IX
●	2016-09-15 00:00	2016-09-28 00:00	Cogent
●	2016-09-23 10:48	2016-09-30 10:48	NL-IX

● - Planned ● - Unloading
● - Maintenance ● - Finished

Figure 3.9.3: Maintenance window sidebar

3.10 Events

The system events are diagnostic messages that were sent by the IRP Components (See also: [IRP Components](#)), as well as notifications caused by the BGP monitoring algorithm (See also: [BGP Monitoring](#)). The events are displayed in the status bar located at the bottom of the system frontend (See [Frontend structure](#)). Several or all the events can be marked as “Read”, so they will no longer be displayed in the status bar.

Mark all as read

<input type="checkbox"/>	#	Timestamp	Module	Event type	Notification message
<input type="checkbox"/>	1	10/30/13 18:40:44	BGPd	Error	BGP session s019/IPv4 (178.236.177.1 AS 48232) outgoing connection disconnected by peer
<input type="checkbox"/>	2	10/30/13 18:40:43	BGPd	Error	BGP session s019/IPv4 (178.236.177.1 AS 48232) disconnected by Network layer request (TCP reset...
<input type="checkbox"/>	3	10/30/13 18:22:50	BGPd	Error	BGP session s019/IPv4 (178.236.177.1 AS 48232) outgoing connection disconnected by peer
<input type="checkbox"/>	4	10/30/13 18:13:00	BGPd	Error	ExternalMon: Peer FR[2] entering FAIL state
<input type="checkbox"/>	5	10/30/13 18:13:00	BGPd	Error	ExternalMon: Peer SWFE[3] entering FAIL state

1

1 - 5 of 5 results found

Figure 3.10.1: System events

3.11 Troubleshooting

The IRP Frontend provides several troubleshooting tools, that can offer you quick information on the specific remote networks.

3.11.1 Traceroute

A traceroute to a specific IP address or hostname can be run from the Frontend. Results will be displayed in the graphical and detailed formats, as in the examples below. The “Allow automatic improvement” checkbox can also be used for automatic improvement of the provided IP address.



Figure 3.11.1: Traceroute

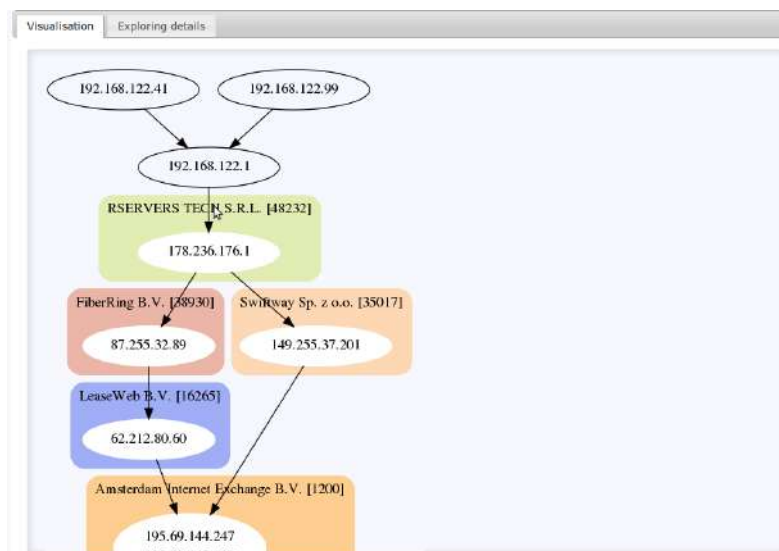


Figure 3.11.2: Traceroute results

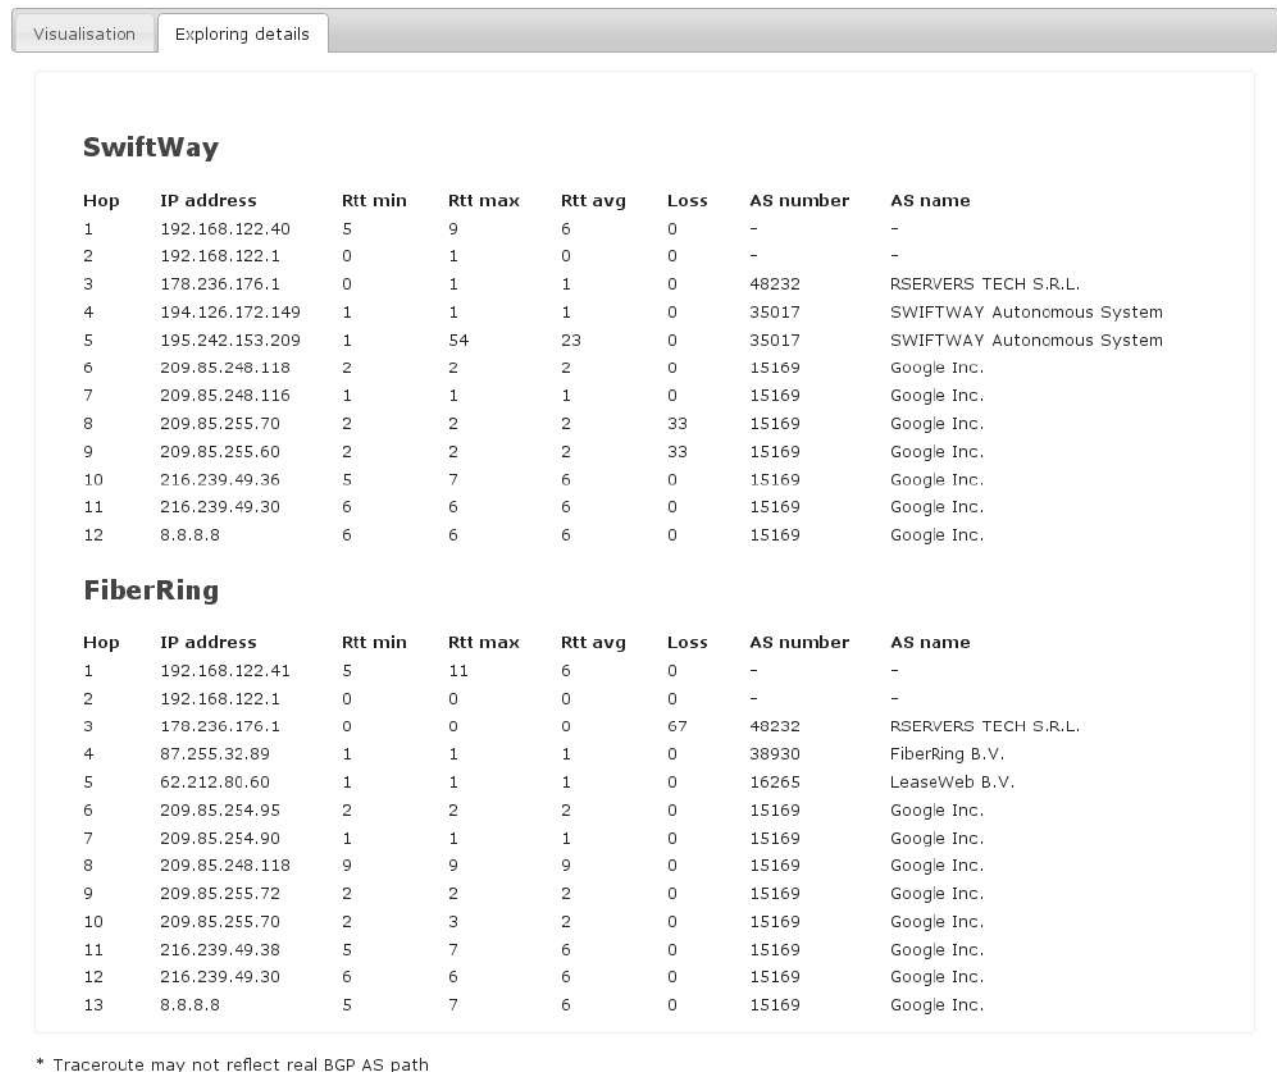


Figure 3.11.3: Traceroute exploring details results

3.11.2 Looking glass

A ping, traceroute or MTR can be run from the Frontend via the default route or via a specific provider.

The screenshot shows the 'Looking glass' web interface. At the top, there's a header 'Looking glass'. Below it, the 'Query type:' section has three radio buttons: 'Ping' (selected), 'Trace', and 'MTR'. The 'Number of packets:' is set to '3'. The 'Send through:' dropdown is set to 'Default route'. The 'Host:' field contains 'google.com'. There is a checkbox for 'Resolve host names' which is checked. A 'Submit' button is at the bottom right. Below the form, the results of the ping test are displayed in a monospaced font.

```

PING google.com (74.125.132.100) 56(64) bytes of data:
64 bytes from wb-in-f100.1e100.net (74.125.132.100): icmp_seq=1 ttl=50 time=6.34 ms
64 bytes from wb-in-f100.1e100.net (74.125.132.100): icmp_seq=2 ttl=50 time=5.89 ms
64 bytes from wb-in-f100.1e100.net (74.125.132.100): icmp_seq=3 ttl=50 time=6.72 ms

--- google.com ping statistics ---
 3 packets transmitted, 3 received, 0% packet loss, time 206ms
rtt min/avg/max/mdev = 5.894/6.328/6.727/0.345 ms

```

Figure 3.11.4: Looking glass

3.11.3 Whois

The “Whois” tool can query IP addresses, hostnames, network prefixes, or AS Numbers. The following valid values can be entered in the search box:

- Hostname: fully qualified domain name, e.g: “domain.com”
- IP: valid IP address, e.g: “10.20.30.40”
- IP Block: any valid prefix, in CIDR format, e.g: “10.20.30.40/24”
- ASN: valid ASN, using the following format: “AS1234”

The screenshot shows the 'Whois' web interface. At the top, there's a header 'Whois'. Below it, the 'Search:' field contains '8.8.8.8'. A 'Whois' button is to the right. Below the search field, the results of the whois query are displayed in a monospaced font.

```

[Querying whois.arin.net]
[whois.arin.net]

#
# ARIN WHOIS data and services are subject to the Terms of Use
# available at: https://www.arin.net/whois_tou.html
#

#
# Query terms are ambiguous. The query is assumed to be:
# "n 8.8.8.8"
# Use "?" to get help.
#

# The following results may also be obtained via:
# http://whois.arin.net/rest/nets;q=8.8.8?showDetails=true&showARIN=false&ext=netref2
#

Level 3 Communications, Inc. LVL3-ORG-8-8 (NET-8-0-0-0-1) 8.0.0.0 - 8.255.255.255
Google Incorporated LVL3-ORG-8-8 (NET-8-8-8-0-1) 8.8.8.0 - 8.8.8.255

#
# ARIN WHOIS data and services are subject to the Terms of Use
# available at: https://www.arin.net/whois_tou.html
#

```

Figure 3.11.5: Whois

3.11.4 Prefix probing

IRP features a manual prefix probing function. One can manually submit a specific IP address for probing and optimization by the system. Valid hostnames and IP addresses can be entered. The probing

results will be displayed on the same page. Please note that IRP probes the exact entered IP address. In case the IP address doesn't respond, IRP runs the indirect probing process and optimizes the whole prefix based on the indirect probing results.

Prefix probing allows to automatically or manually choose the best path for a particular IP address or prefix.

In order to check what are the possible paths for a particular prefix, enter the IP address or prefix into the field. If you prefer the prefix to be improved automatically, tick the check box "Improve automatically" and press the "Submit for probing" button. Otherwise, just press the "Submit for probing" button and wait until the system returns the probing results.



Figure 3.11.6: Prefix Probing

The system returns the probing results by displaying the loss, latency and cost values for each of the providers.

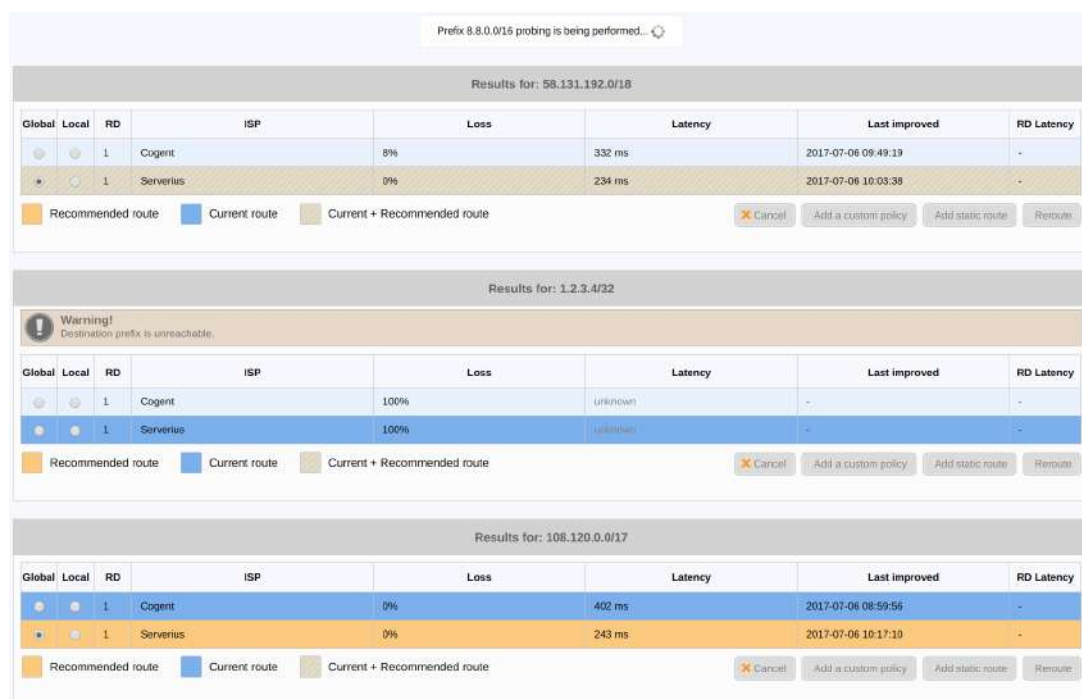


Figure 3.11.7: Prefix probing results

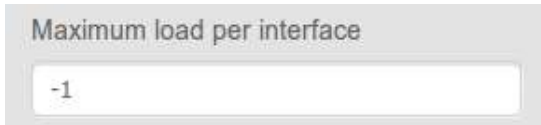
After the probing results are received, choose the path you consider the best one. Note, that IRP highlights the current and the best path. Use the "Reroute" button to redirect the traffic through the selected provider. You can also set the selected provider as static for the probed prefix by pressing the "Add static route" button or apply a specific routing policy to the probed prefix by clicking the "Add a custom policy" button.

3.12 Configuration editor

Starting with IRP version 1.7, all system parameters can be modified from the Frontend, using the “Configuration” menu.

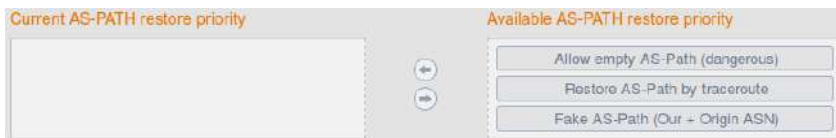
Depending on the parameter type, several control types are available:

- Text boxes, which can hold IP addresses, provider names and numerical values




A text input field with the label "Maximum load per interface" and the value "-1".

- Auto-complete boxes, used for parameters that can take a limited set of values



An auto-complete box with the label "Available AS-PATH restore priority". It contains three options: "Allow empty AS-Path (dangerous)", "Restore AS-Path by traceroute", and "Fake AS-Path (Our + Origin ASN)".

- Auto-complete list boxes, usually containing a list of IPs / prefixes (in CIDR format) / ASNs



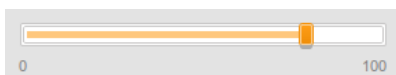
An auto-complete list box with the label "Infrastructure IPs*". It contains two entries: "192.168.122.0/24" and "178.236.176.0/20".

- Buttons



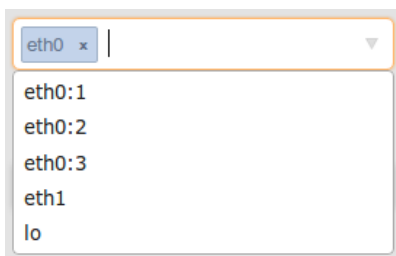
Two buttons: "SUBMIT" with a checkmark icon and "RESET" with a circular arrow icon.

- Sliders, used to select a value from a predefined range



A slider control with a range from 0 to 100. The slider is currently positioned at approximately 50.

- Drop-down lists



A drop-down list with the label "eth0 x". The list contains the following items: "eth0:1", "eth0:2", "eth0:3", "eth1", and "lo".

- Toggle buttons (used for specific algorithms and BGP peers)



Four toggle buttons: "OFF", "ON", "IPv4", and "IPv6". Below them are six buttons: "DISABLED", "ENABLED", "DISABLED", "ENABLED", "INTRUSIVE", and "NONINTRUSIVE".

- Provider control buttons, can be used to suspend, resume, stop, or remove a provider



Three provider control buttons: a pause button (two vertical bars), a stop button (a square), and a remove button (an 'X').

3.12.1 Global configuration

Global settings can be adjusted from the Frontend, by using the “Configuration” menu.

Please refer to [Global and Core Configuration](#) for a detailed parameters description.

Several global parameters can be configured by selecting their desired values and clicking on the “Submit” button.

Available parameters:

- **Improvement mode** ([global.improve_mode](#))
- **Management interface** ([global.master_management_interface](#))
- **BGP mode** (Intrusive / Non-intrusive, see [global.nonintrusive_bgp](#))
- **Probing interface(s)**. See [global.master_probing_interface](#)
- **Prefix aggregation**. If this parameter is enabled, the system operates at aggregate level, up to the **Maximum aggregate mask** (see: [global.aggregate](#), [global.agg_ipv4_max](#), [global.agg_ipv6_max](#))

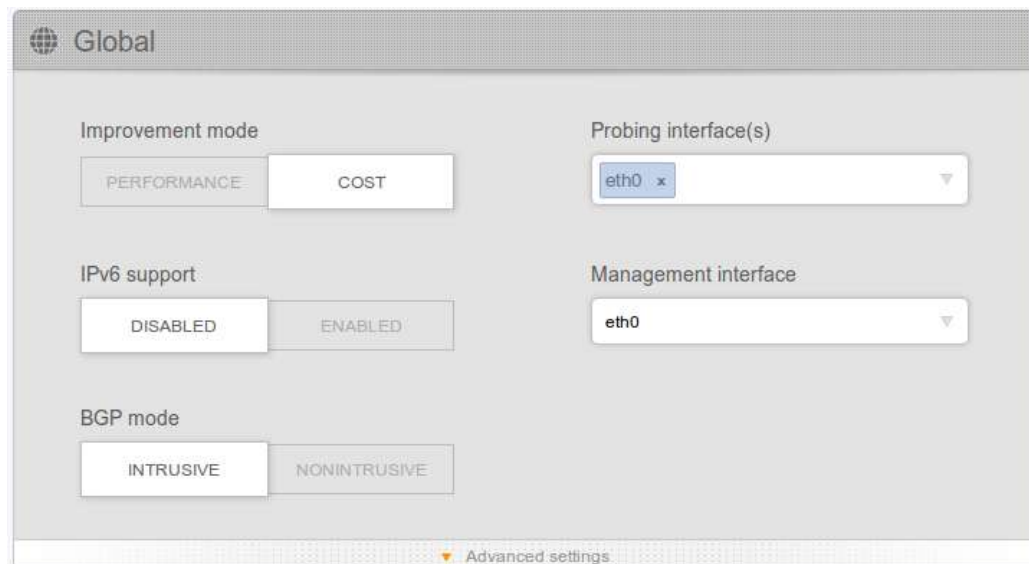


Figure 3.12.1: Configuration editor: Global settings

The list of ignored networks can be modified on the same page. The options allow listing as appropriate:

- prefixes([global.ignorednets](#)),
- ASNs ([global.ignored.asn](#)),
- BGP Community attributes that mark prefixes/routes to ignore ([global.ignored_communities](#)).

Setting prefixes and ASNs is possible either manually, or by importing a text file containing one IP/network in CIDR format or ASN per line.

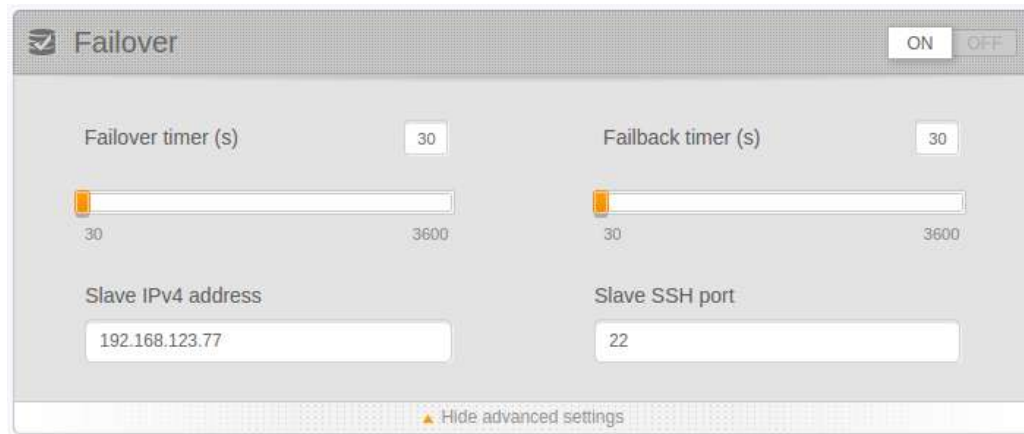
Figure 3.12.2: Configuration editor: Ignored prefixes

Failover functionality is switched on or off ([global.failover](#)) as part of Global configuration too.

i Failover license is required for Failover tab to be available.

Failover configuration covers the following parameters:

- Failover timer - time period in seconds during which master is considered alive and before slave becomes active if there are no announcements from master ([global.failover_timer_fail](#)).
- Failback timer - time period before slave node will release control back to master ([global.failover_timer_failback](#)). This is to protect from failover flapping between master and slave in case of master instability.
- Slave IPv4 address - indicates to master what is the slave node ([global.failover_slave.ip](#)). Master uses this address to setup SSH connections and sync its own configuration files.
- Slave SSH port - indicates the TCP port number for SSH on the slave node ([global.failover_slave.port](#)). Master uses this port to setup SSH connections and sync its own configuration files.



The screenshot shows a configuration window titled "Failover" with a checked checkbox icon on the left and "ON" / "OFF" toggle buttons on the right. The window contains four main settings:

- Failover timer (s):** A numeric input field set to "30" and a slider below it ranging from 30 to 3600.
- Failback timer (s):** A numeric input field set to "30" and a slider below it ranging from 30 to 3600.
- Slave IPv4 address:** A text input field containing "192.168.123.77".
- Slave SSH port:** A text input field containing "22".

At the bottom of the window, there is a button with a yellow triangle icon and the text "Hide advanced settings".

Figure 3.12.3: Configuration editor: Global failover

3.12.2 BGP and Routers configuration

BGP daemon-related parameters can be configured in the “Configuration→BGP” section.

- Bgpd parameters (See also: [Bgpd Configuration](#), [Bgpd settings](#)):
 - Split announced improvements to preserve original route attributes (`bgpd.updates.split`)
 - Remove prefixes (improvements) on next-hop update (`bgpd.improvements.remove.next_hop_eq`)
 - Strip non-IRP communities for improvements (`bgpd.improvements.strip_non_irp_communities`)
 - Remove prefixes (improvements) on aggregate withdrawal (`bgpd.improvements.remove.withdrawn`)
 - Re-probing on new (improved) or old route changes if BMP is available (`bgpd.retry_probing.new.bmp_path_change` `bgpd.retry_probing.old.bmp_path_change`)
- BGP monitoring settings can be found under advanced settings and include but is not limited to (See also: [BGP Monitoring](#)):

i BGP monitor can be enabled/disabled. This option is useful during DoS/DDoS attacks.

- Guard time (`bgpd.mon.guardtime`)
- Holdtime (`bgpd.mon.holdtime`)
- Keepalive (`bgpd.mon.keepalive`)
- Flapping protection (`bgpd.mon.internal.flap_guardtime`)

The screenshot displays two configuration panels. The top panel, titled 'BGPd', contains several toggle switches and dropdown menus. The bottom panel, titled 'Router R1', contains text input fields for various network parameters.

BGPd Configuration:

- Split announcements to preserve original route attributes: ☐ DISABLED ☒ ENABLED
- Remove next hop update: ☐ DISABLED ☒ ENABLED
- Strip non-IRP communities: ☐ DISABLED ☒ ENABLED
- Remove on aggregate withdrawal: ☐ DISABLED ☒ ENABLED
- Re-probe on new path change:
- Re-probe on old path change:

Router R1 Configuration:

- Autonomous System*:
- Improvement localpref:
- IRP's IPv4 address*:
- Router IPv4 address*:

Figure 3.12.4: Configuration editor: BGP settings

BGP routers can be added, removed or modified on the same page.

A new BGP neighbor can be added using the “Add Router” button. Several parameters must be configured:

- General settings:
 - Router name
 - ASN to be used for the iBGP session (`bgpd.peer.X.as`)
 - Improvement local-pref (`bgpd.peer.X.master_localpref`)
 - Local IP address (`bgpd.peer.X.master_our_ip`, `bgpd.peer.X.master_our_ipv6`)
 - Neighbor IP address, usually the router's IP (`bgpd.peer.X.master_peer_ip`, `bgpd.peer.X.master_peer_ipv6`, `bgpd.peer.X.slave_peer_ip`, `bgpd.peer.X.slave_peer_ipv6`)

The screenshot shows the 'Add Router' window with the 'General' tab selected. It contains the following fields:

- Autonomous System* (text input)
- Improvement localpref (text input)
- IRP's IPv4 address* (text input)
- Router IPv4 address* (text input)
- Router name (text input)

A 'Save' button with a checkmark icon is located at the bottom right.

Figure 3.12.5: Configuration editor: New Router, general settings

- Inbound settings:
 - Local IPv4/IPv6 inbound next hop (`bgpd.peer.X.inbound.ipv4.next_hop`, `bgpd.peer.X.inbound.ipv6.next_hop`)
 - Announced inbound LocalPref value (`bgpd.peer.X.inbound.master_localpref`, `bgpd.peer.X.inbound.slave_localpref`)
 - Transit SNMP host (`bgpd.peer.X.transit.snmp`)
 - Transiting traffic (`global.inbound_transit`)

The screenshot shows the 'Add Router' window with the 'Inbound' tab selected. It contains the following fields and controls:

- Local IPv4 inbound next_hop* (text input)
- Announced inbound localpref value (text input)
- Transit SNMP host (dropdown menu showing 'None')
- Transiting traffic (toggle buttons for 'DISABLED' and 'ENABLED')

A 'Save' button with a checkmark icon is located at the bottom right.

Figure 3.12.6: Configuration editor: New Router, inbound settings

- Advanced settings:
 - Keepalive (`bgpd.peer.X.keepalive`)

- Improvement communities to be appended to the Community attribute (`bgpd.peer.X.master_communities`)
- Improvement local-pref (`bgpd.peer.X.master_localpref`)
- Router ID, mandatory for IPv6 only (`bgpd.peer.X.slave_router_id`)
- BGP session password (`bgpd.peer.X.master_password`, `bgpd.peer.X.slave_password`)
- Improvement MED value (`bgpd.peer.X.med`)

The screenshot shows a window titled "Add Router" with a close button in the top right. It has three tabs: "General", "Advanced", and "Inbound". The "Advanced" tab is selected. The configuration area contains several fields and a slider:

- Keepalive interval (sec):** A slider ranging from 1 to 100, with a value of 60 displayed.
- Improvement communities:** A text input field.
- Improvement localpref:** A text input field.
- BGP Router ID:** A text input field.
- BGP session password:** A text input field.
- Improvement MED value:** A text input field.

A "Save" button with a checkmark icon is located at the bottom right of the window.

Figure 3.12.7: Configuration editor: New Router, advanced settings

See also: [Bgp Configuration](#), [Bgp settings](#), [BGP sessions settings](#)

- Failover settings - the same set of parameters shall be setup for both master and slave nodes. Use the toggle to switch between them.
 - Improvement communities to be appended to Community attribute (`bgpd.peer.X.slave_communities`, `bgpd.peer.X.slave_communities`)
 - Improvement LocalPref value (`bgpd.peer.X.master_localpref`, `bgpd.peer.X.slave_localpref`)
 - Local IPv4 address (`bgpd.peer.X.master_our_ip`, `bgpd.peer.X.slave_our_ip`)
 - BGP session password (`bgpd.peer.X.master_password`, `bgpd.peer.X.slave_password`)



Figure 3.12.8: Configuration editor: New Router, advanced settings

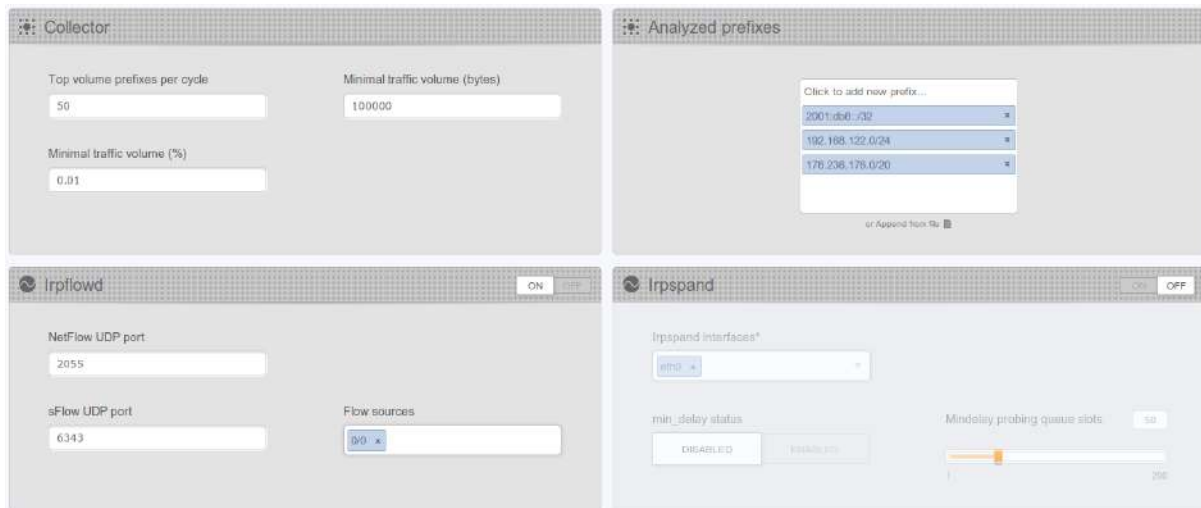
See also: [Bgp Configuration](#), [Bgp settings](#), [BGP sessions settings](#)

3.12.3 Collector configuration

Global collector parameters, as well as Span and Flow-specific settings can be adjusted by using the “Configuration→Collector” menu item.

See also: [Collector Configuration](#), [Collector settings](#).

- Global collector settings:
 - Top volume prefixes per cycle ([collector.export.volume.high.top_n](#))
 - Minimal traffic volume ([collector.export.volume.min](#))
 - Minimal traffic volume (%) ([collector.export.volume.min_pct](#))
- Flow collector settings:
 - UDP port to listen for NetFlow or sFlow packets ([collector.flow.listen.nf](#), [collector.flow.listen.sf](#))
 - Flow sources ([collector.flow.sources](#))
- SPAN collector settings:
 - Capture interfaces ([collector.span.interfaces](#))
 - Mindelay algorithm enable/disable ([collector.span.min_delay](#))
 - Mindelay probing queue slots ([collector.span.min_delay.probing_queue_size](#))



The figure displays four panels of the configuration editor for the Collector settings:

- Collector**: Contains input fields for 'Top volume prefixes per cycle' (50), 'Minimal traffic volume (bytes)' (100000), and 'Minimal traffic volume (%)' (0.01).
- Analyzed prefixes**: Features a list of prefixes: '2001:db0::32', '192.168.122.0/24', and '176.236.176.0/20'. A button 'Click to add new prefix...' is at the top, and a link 'or Append from file' is at the bottom.
- Irpflowd**: Includes a toggle switch (ON) and input fields for 'NetFlow UDP port' (2055), 'sFlow UDP port' (6343), and 'Flow sources' (9/0).
- Irpspand**: Includes a toggle switch (OFF), a section for 'Irpspand interfaces*' with a dropdown menu showing 'eth0', a 'min_delay status' section with 'DISABLED' and 'ENABLED' buttons, and a 'Mindelay probing queue slots' slider set to 50.

Figure 3.12.9: Configuration editor: Collector settings

The list of prefixes announced by the monitored network can be edited in the same form. Its contents can be also imported from a text file.

See also: collector.ournets.net.

3.12.4 Core configuration

All Core parameters affecting the decision-making algorithms can be modified on the next page (“Configuration→Core”). See also: [Global and Core Configuration](#), [Core settings](#).

The following configuration parameters can be adjusted:

- **Max IPv4 improvements** (`core.improvements.max`, `core.improvements.max_ipv6`)
- **Standard reprobing period** (`core.improvements.ttl.retry_probe`) - the time period after which a specific improvement is being scheduled for re-probing.
- **VIP reprobing period** (`core.vip.interval.probe`) - defines the VIP prefixes probing interval, in seconds.
- **Minimal probe lifetime** (`core.probes.ttl.min`) - Ordinary probing will not be performed for a specific prefix if its probe age is lower than this value.
- **Allowed latency worsening** (`core.cost.worst_ms`)
- **Exploring queue slots** (`core.eventqueueulimit`)
- **Relevant loss** for loss-based improvements (`core.performance.loss_pct`) - a prefix will be improved based on a loss cause only if the packet loss can be decreased by this value (in %)
- **Relevant RTT difference** (`core.performance.rtt.diff_ms`)
- **Relevant RTT difference (%)** (`core.performance.rtt.diff_pct`)
- **Maximum probe lifetime** (`core.probes.ttl.max`)



Figure 3.12.10: Configuration editor: Core configuration

[Outage detection](#) algorithm can be enabled and configured on the same page. Outage detection algorithm can be enabled or disabled using the toggle On/Off buttons at the top of each form.

Figure 3.12.11: Configuration editor: Outage detection

Throttling excessive bandwidth use with Flowspec can be enabled and configured on the same page. The feature can be enabled or disabled using the toggle On/Off buttons at the top of each form.

Figure 3.12.12: Configuration editor: Overusage

The following configuration parameters can be adjusted:

- **Overusage interval** (`core.overusage.check_interval`) sets the frequency in seconds of checking for excessive bandwidth use.
- **Overusage rule retention** (`core.overusage.hold_timer`) sets how much time a rule is kept after bandwidth use returns to normal.
- **Prefix BW average time** (`core.overusage.out.average.period`) sets the number of hours used to determine the average bandwidth use of a prefix.
- **Prefix relevant BW** (`core.overusage.out.average.relevant_min`) sets the relevant bandwidth use in Mbps by a prefix before considering any rules for it.

- **Overusage throttle multiplier** (`core.overusage.out.threshold.throttle`) sets the multiplier to apply to average prefix bandwidth use when setting a rule.
- **Overusage threshold multiplier** (`core.overusage.out.threshold.trigger`) sets the threshold multiplier used to determine excessive bandwidth use.

➖ Prefix relevant BW and Overusage threshold multiplier parameters control the number of Flowspec rules that will be generated automatically. Adjust these values to control the number of Throttling Flowspec policies.

Circuit issues detection parameters are listed and can be adjusted:

The screenshot shows a configuration window titled "Circuit issues detection". It contains the following controls:

- Delta loss to shutdown:** A slider from 2 to 50 with a value of 15.
- Delta loss to warn:** A slider from 1 to 49 with a value of 10.
- Delta loss to restore:** A slider from 0 to 48 with a value of 5.
- Restore interval (min):** A slider from 1 to 30 with a value of 5.
- Issues time horizon (min):** A slider from 1 to 240 with a value of 5.
- Restore after (sec):** A slider from 60 to 3600 with a value of 600.
- Withdraw improvements on warn:** A toggle switch currently set to "DISABLED".

Figure 3.12.13: Configuration editor: Circuit issues detection

The following configuration parameters can be adjusted:


- **Delta loss to shutdown** (`core.circuit.high_loss_diff`) sets threshold to initiate shutting down a provider with circuit issues as a difference between examined provider's average loss and overall average loss during the given time horizon. Shutdown is attempted only for providers marked accordingly.
- **Delta loss to warn** (`core.circuit.warn_loss_diff`) sets threshold to raise a warning regarding issues with a provider as difference between examined provider's average loss and overall average loss during the given time horizon. Usually this threshold is significantly lower than the shutdown threshold.
- **Delta loss to restore** (`core.circuit.recover_loss_diff`) sets low loss level threshold when IRP will restore full functionality over provider that had circuit issues.
- **Issues time horizon** (`core.circuit.hist_interval`) sets how many minutes in the past IRP looks for probes with loss over both analyzed provider and all the other providers on the network.
- **Withdraw improvements on warn** (`core.circuit.withdraw_on_warn`) instructs IRP to withdraw outbound improvements over provider with circuit issues when warning threshold is reached. By default improvements are withdrawn only when shutdown threshold is exceeded.

- **Restore after** (`core.circuit.recover_hold_time`) sets the interval in seconds after which IRP should re-evaluate a provider's circuit loss and attempt restoring it to normal function. This will be performed only for providers that are configured to attempt to restore after shutdown.
- **Restore interval** (`core.circuit.recover_monitored_intervals`) sets the interval in minutes during which IRP will periodically re-evaluate a provider's circuit loss and attempt restoring it to normal function. This will be performed only for providers that are configured to attempt to restore after shutdown.

➡ The configuration parameters above are applied for all providers even though for individual providers different level of control can be set starting from disabling and ending with attempting both to automatically shutdown and later restore connectivity with it. See also [Providers configuration](#)

3.12.5 Commit Control configuration

Commit Control algorithm can be enabled and configured on the Commit Control group of pages. It can be enabled or disabled using the toggle On/Off buttons at the top of the form.

 If NetFlow is used to collect statistics Routers MUST be configured to export statistics every minute (or as often as possible). Some router models have default export intervals for either inactive or active flows of up to 1800 seconds. Big delays cause IRP to react very slowly to increased load and reduce the effectiveness of Commit Control feature.

If NetFlow is used to collect statistics Routers MUST be configured to export statistics every minute (or as often as possible). Some router models have default export intervals for either inactive or active flows of up to 1800 seconds. Big delays cause IRP to react very slowly to increased load and reduce the effectiveness of Commit Control feature.

The most important parameters for Commit Control are, as follows:

- **Commit Control probing queue slots** (`core.commit_control.probing_queue_size`)
- **Minimal prefix bandwidth** in Mbps (`core.commit_control.agg_bw_min`)

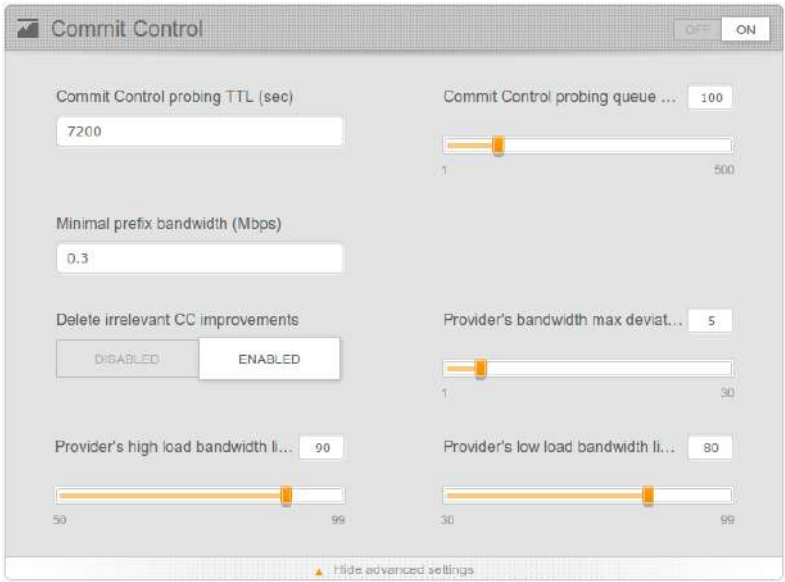


Figure 3.12.14: Configuration editor: Commit Control

Providers Overall block displays Commit Control configuration for all providers including their precedence. See details in the figure below.

Individual provider parameters can be adjusted as well as detail regarding current Commit Control changes applied by IRP can be accessed by following the provided links.


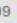















Providers Overall - Commit Control							
Provider name	Commit Control	Configured 95th (Mbps)	Current bw usage (Mbps)	Monthly 95th (Mbps)	Cost (USD)	Load balancing	Precedence
B	On	80 	29.09  	51 (64%) 	3 	Disabled	80 
IX2	On	5 	0  	-	2 	Disabled	73 
A	On	180 	101.86  	107 (59%) 	3 	Disabled	70 

Figure 3.12.15: Providers Overall

3.12.6 Explorer configuration

To adjust the Explorer-related parameters, the “Configuration→Explorer” menu item must be accessed. The [Explorer Configuration](#) and [Explorer settings](#) sections should be consulted for detailed Explorer configuration and parameters description.

Explorer parameters:

- Infrastructure IPs ([explorer.infra_ips](#))
- Explorer worker threads ([explorer.maxthreads](#))
- Probing algorithms and Traceroute algorithms ([explorer.probe.algorithm](#), [explorer.trace.algorithms](#))
- High volume task precedence ([explorer.high_vol_precedence](#))
- Process max collected IPs ([explorer.max_collector_ips](#))
- First probing packets amount ([explorer.probing.sendpkts.min](#))
- Adaptive probing packets count ([explorer.probing.sendpkts.adaptive_max](#))
- ICMP timeout ([explorer.timeout](#))
- Traceroute retry packets and packets per hop ([explorer.traceroute.sendpkts](#), [explorer.traceroute.retrypkts](#))
- Traceroute minimum and maximum ttl ([explorer.traceroute.ttl.min](#), [explorer.traceroute.ttl.max](#))




Figure 3.12.16: Configuration editor: Explorer settings

3.12.7 Providers and Peers configuration

Providers can be added and modified via the “Configuration→Providers and Peers” menu.

See also: [Provider](#)

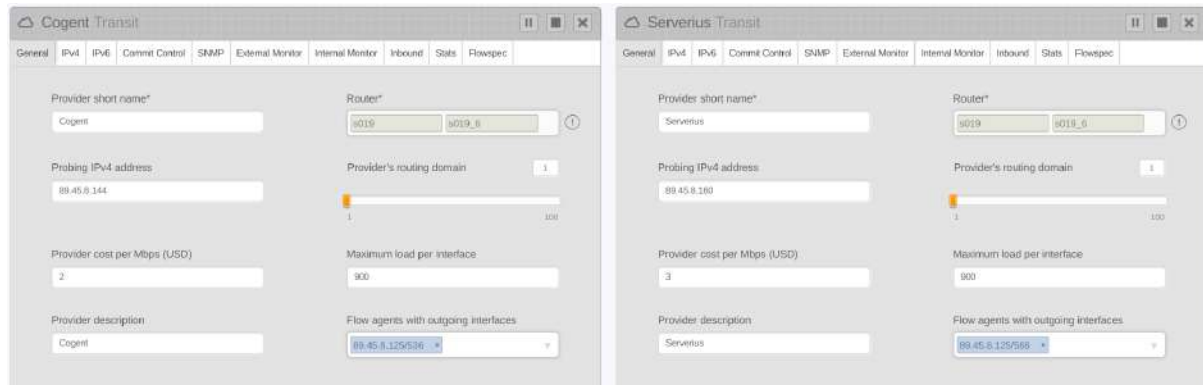


Figure 3.12.17: Configuration editor: Providers configuration

Using the provider control buttons, a specific provider can be temporarily suspended (e.g for a short maintenance in the monitored network), stopped (for long maintenance), or completely deleted. See also: [peer.X.shutdown](#). Commit Control can also be disabled or enabled for each provider specifically. ([peer.X.cc_disable](#)).

i The sections below describe the most common groups of provider configuration parameters. Some parameters are repeated on more than one of these groups for convenience.

3.12.7.1 Providers configuration

To add a new provider, the “Add provider” button must be used. Several parameters should be configured for the new provider:

- General settings:
 - Provider short name ([peer.X.shortname](#))
 - Router ([peer.X.bgp_peer](#))
 - Probing IPv4 address ([peer.X.ipv4.master_probing](#)) or Probing IPv6 address ([peer.X.ipv6.master_probing](#))
 - Provider’s routing domain ([peer.X.rd](#))
 - Provider cost per Mbps ([peer.X.cost](#))
 - Maximum load per interface ([peer.X.limit_load](#))
 - Provider description ([peer.X.description](#))
 - Flow agents ([peer.X.flow_agents](#))
 - Circuit issues detection ([peer.X.circuit.control](#)) indicating how IRP should react to excessive persistent loss over a particular provider.

The screenshot shows a web-based configuration interface titled "Add Transit Provider". It has a tabbed interface with the following tabs: General, IPv4, IPv6, Commit Control, SNMP, External Monitor, Internal Monitor, Inbound, and an unlabeled tab. The "General" tab is currently selected. The form contains the following fields and controls:

- Provider short name***: A text input field.
- Router***: A dropdown menu with a warning icon.
- Probing IPv4 address**: A text input field.
- Provider's routing domain**: A numeric input field with the value "1" and a slider below it ranging from 1 to 100.
- Provider cost per Mbps (USD)**: A text input field.
- Maximum load per interface**: A dropdown menu with the value "No limit".
- Provider description**: A text input field.
- Flow agents with outgoing interfaces**: A dropdown menu.

A "Save" button with a checkmark icon is located at the bottom right of the form.

Figure 3.12.18: Configuration editor: Adding a new provider, General settings

- Commit Control configuration:
 - Provider cost per Mbps (`peer.X.cost`)
 - Maximum load per interface (`peer.X.limit_load`)
 - Provider 95th percentile (`peer.X.95th`)
 - Provider inbound 95th percentile (`peer.X.95th.in`)
 - Provider 95th calculation mode for inbound traffic (95th calculation modes)
 - Commit Control status for this provider (`peer.X.cc_disable`)
 - CC provider precedence - used for Commit Control and grouping (`peer.X.precedence`)
 - Performance/Cost improvements within provider group flag in case Provider load balancing is configured, (`peer.X.improve_in_group`, `peer.X.precedence`)
 - Provider billing day (`peer.X.95th.bill_day`)
 - Centile value (`peer.X.95th.centile`)

The screenshot shows the 'Add Transit Provider' configuration window with the 'Commit Control' tab selected. The window has tabs for General, IPv4, IPv6, Commit Control, SNMP, External Monitor, Internal Monitor, and Inbound. The Commit Control tab contains the following fields and controls:

- Provider cost per Mbps (USD): Text input field.
- Maximum load per interface: Text input field with 'No limit' selected.
- Provider 95th percentile*: Text input field.
- Provider 95th percentile (inbound): Text input field.
- Provider 95th Calculation Mode (inbound...): Dropdown menu with '95th from greater of in, out' selected.
- Commit Control for provider: Two buttons, 'ENABLED' (active) and 'DISABLED'.
- CC provider precedence*: A slider control ranging from 0 to 100, currently set near 0.
- Performance/Cost improvements with...: Two buttons, 'DISABLED' and 'ENABLED'.
- Provider billing day: Text input field with '1' entered.
- Centile value: Text input field with '95' entered.

A 'Save' button with a checkmark icon is located at the bottom right of the window.

Figure 3.12.19: Configuration editor: Adding a new provider, Commit Control

- SNMP-related settings:

The screenshot shows the 'Add Transit Provider' configuration window with the 'SNMP' tab selected. The window has tabs for General, IPv4, IPv6, Commit Control, SNMP, External Monitor, Internal Monitor, and Inbound. The SNMP tab contains the following fields and controls:

- IPv4 / IPv6: Two buttons, with 'IPv6' selected.
- Provider SNMP IPv4 address: Text input field.
- Provider SNMP community: Text input field.
- Provider SNMP interface: Text input field with 'ifIndex.ifName' and a refresh icon.
- SNMP version (stats): Dropdown menu with 'v2c' selected.

A 'Save' button with a checkmark icon is located at the bottom right of the window.

Figure 3.12.20: Configuration editor: Adding a new provider, SNMP settings

- SNMP v3 uses additional parameters depending on security services used for statistics collection:
 - SNMP security services selects if Authentication and/or Privacy services are used (`snmp.X.secllevel`)
 - SNMP authentication password if authentication is used (`snmp.X.auth_password`)
 - SNMP authentication protocol if authentication is used (`snmp.X.auth_protocol`)
 - SNMP encryption password if privacy is used (`snmp.X.priv_password`)
 - SNMP encryption protocol if privacy is used (`snmp.X.priv_protocol`)
 - SNMP Username if authentication is used (`snmp.X.auth_username`)
 - SNMP version (`snmp.X.version`)
- External Monitor settings:
 - External monitor status flags for IPv4 / IPv6 (`peer.X.mon.ipv4.external.state`, `peer.X.mon.ipv6.external.state`)

- ICMP/UDP ping monitored IPv4 / IPv6 addresses (`peer.X.ipv4.mon`, `peer.X.ipv6.mon`)

The screenshot shows the 'Add Transit Provider' configuration window with the 'External Monitor' tab selected. The window has tabs for General, IPv4, IPv6, Commit Control, SNMP, External Monitor, Internal Monitor, and Inbound. The 'External Monitor' section contains two columns of settings. The left column is for IPv4 and the right for IPv6. Each column has an 'External Monitor' status toggle (currently 'DISABLED'), an 'ICMP/UDP ping monitored' address field (IPv4: 208.67.222.222, IPv6: 2620:0:ccc::2), and a 'Save' button at the bottom right.

Figure 3.12.21: Configuration editor: Adding a new provider, External Monitor

- Internal Monitor settings:
 - Internal monitor status flags for IPv4 / IPv6 (`peer.X.mon.ipv4.internal.state`, `peer.X.mon.ipv6.internal.state`)
 - BGP session monitoring IPv4 / IPv6 addresses (`peer.X.mon.ipv4.bgp_peer`, `peer.X.mon.ipv6.bgp_peer`)
 - BGP MIBs for IPv4 / IPv6 (`peer.X.mon.ipv4.internal.mode`, `peer.X.mon.ipv6.internal.mode`)
 - SNMP host for BGP Internal Monitor (`peer.X.mon.snmp`)

The screenshot shows the 'Add Transit Provider' configuration window with the 'Internal Monitor' tab selected. The window has tabs for General, IPv4, IPv6, Commit Control, SNMP, External Monitor, Internal Monitor, and Inbound. The 'Internal Monitor' section contains two columns of settings. The left column is for IPv4 and the right for IPv6. Each column has an 'Internal Monitor' status toggle (currently 'DISABLED'), a 'BGP session monitoring' address field, a 'BGP MIB (IPv4/IPv6)' dropdown (IPv4: Generic (BGP4-MIB), IPv6: Cisco (CISCO-BGP4-MIB)), a 'BGP Internal Monitor SNMP' address field, a 'BGP Internal Monitor SNMP Community' field, and an 'SNMP version (monitor)' dropdown (currently 'v2c'). A 'Save' button is at the bottom right.

Figure 3.12.22: Configuration editor: Adding a new provider, Monitoring

- SNMP v3 uses additional parameters depending on security services used for monitoring:
- IPv4 / IPv6 settings:
 - IPv4 / IPv6 diagnostic hops (`peer.X.ipv4.diag_hop`, `peer.X.ipv6.diag_hop`)

- Probing IPv4 / IPv6 address (`peer.X.ipv4.master_probing`, `peer.X.ipv6.master_probing`)
- Remote provider ASN (`peer.X.ipv4.next_hop_as`), used by the AS-Path restoration algorithm (`bgpd.as_path`)
- Router next-hop address (`peer.X.ipv4.next_hop`), that sets the next-hop value for injected routes related to this provider

Figure 3.12.23: Configuration editor: Adding a new provider, Commit Control

The same parameters can be adjusted for the already configured providers.

3.12.7.2 Internet Exchanges configuration

An Exchange is configured as a special type of provider. The important things to consider when setting up an Exchange are listed below.

i It is recommended that Noction systems engineers guide you through Exchange configuration process.

- IRP needs one ACL and one PBR rule for each Peering Partner on an Exchange. Internet Exchanges can have hundreds of Peering Partners. Ensure that the number of Access Lists, Access List entries and PBR entries will not exceed router hardware limitations.

- Once the Exchange is setup IRP will retrieve the routing table on the router in order to setup Exchange peering partners. IRP will require access to the router in order to do so.

- When configuring an Exchange its Diag Hop parameter must be provided. Diag Hop for an Exchange represents the list of direct-connected networks configured on the router's interface towards the Exchange network. This parameter is labeled "IPv4 diagnostic hop" on IRP Frontend.

- Probing IPs for Exchanges no longer require one IP address per peering partner since this number might be impractically big. Instead a combination of probing IP and DSCP value is used to configure the PBRs. A single probing IP in conjunction with DSCP can cover up to 64 peering partners. A sufficient number of probing IPs will be required in order to cover all peering partners on your exchange.

⚠ As a rule of thumb it is better to list the PBR rules for transit providers before the (large number of) rules for Internet Exchange peers. If the many PBR rules assigned to peers on a large Internet Exchange are placed at the beginning of the PBR list some routers evaluate all of them before finding the terms for transit providers and this consumes valuable router CPU resources.

The screenshot shows the 'Ix2 - Exchange' configuration window. The 'IPv4' tab is selected. The configuration fields are as follows:

- Provider short name*:** IX2
- Router*:** s020
- IPv4 diagnostic hop*:** 10.10.20.20
- Probing IPv4 address:** 10.10.20.20, 10.10.20.21, 10.10.20.22
- Provider cost per Mbps (USD):** 2
- Maximum load per interface:** -1
- Provider description:** IX2

A note below the probing IP addresses states: '3 IP addresses can be used to setup up to 192 peering partners'.

Figure 3.12.24: An Internet Exchange is similar to a provider and it requires configuration of the Router, the Probing IPs and the other usual attributes

- Once the Exchange is configured, IRP will need to reset its BGP session towards the router interconnecting with the Exchange in order to retrieve the required routing table information about all peering partners. Once the BGP session is restarted IRP will start fetching Exchange routing tables.

i Give IRP sufficient time (~10 minutes) to fetch Exchange routing tables before continuing configuration.

- IRP offers peering partner autoconfiguration. It retrieves the list of Next Hops (peering partners) on the Exchange with the corresponding list of prefixes individual peers accept.

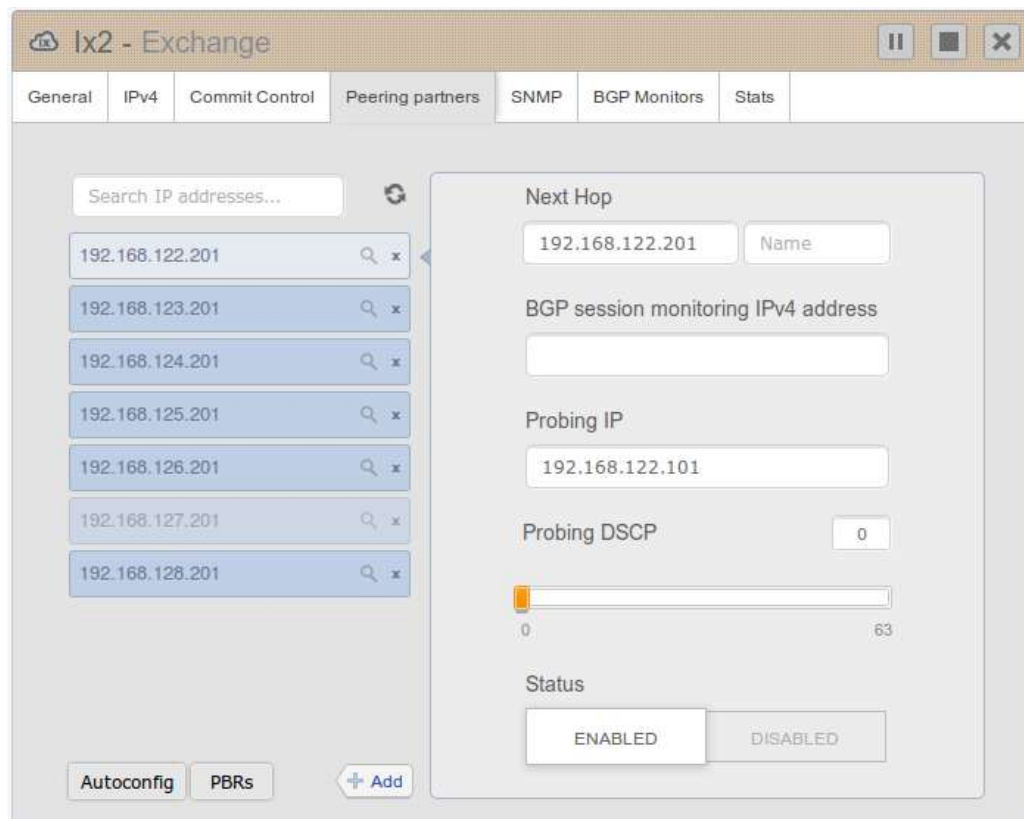


Figure 3.12.25: The list of peering partners for the Exchange shows details about each of them and offers access to IRP Autoconfiguration feature and PBR ruleset generator for the router

- Use the Autoconfiguration feature to create an initial configuration for IRP. Review Exchange peering partners before starting the Exchange. Consider enabling periodic auto re-configurations for selected IX in order to update periodically the value of BGP session monitoring IPv4 address from the Exchange and to add new peering partners. Refer [global.exchanges.auto_config_interval](#) and [peer.X.auto_config](#).

i Keep in mind that Autoconfiguration might tamper with all the changes you've made directly within IRP config files for peering partners. So, use Autoconfiguration sparingly after you've applied manual changes or try avoiding direct configuration file changes altogether.

- Once the Exchange is configured correctly within IRP you will need to apply the PBR rules on the router(s). IRP includes the functionality to generate PBR rules for different router vendors. You will need to review the generated ruleset and apply it on the router manually.

Figure 3.12.26: PBR Generator

- It is possible that the Autoconfiguration feature on the Exchange has been run with incomplete routing tables or that new peering partners have been connected. This is especially true for very big Exchanges. In this case it is possible that some peering partners are neither in IRP nor router configurations. When IRP detects such a case it raises a warning about newly discovered peering partners like in the image below. At this stage you will need to both re-autoconfigure IRP and extract the PBRs for the router.

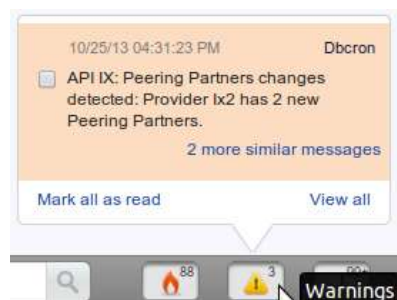


Figure 3.12.27: Warning about new peering partners that have been added to the Exchange

- After its creation and up to its complete configuration the Exchange is kept in a Stopped state. When both IRP and the Router PBRs are configured IRP is ready to start optimizing Exchange traffic too. Keep in mind that starting the Exchange will require a BGP session restart.

i We must mention that before applying changes to IRP configuration the changes are validated. If errors are detected these will be flagged and the previous good configuration will be kept intact so you will have the option to review the erroneous data and correct.

3.12.7.3 Switch a provider from one router to another

Switching a provider from one router to another is possible via manual IRP reconfiguration only.

Use the following steps to switch the provider from one router to another:

1. Suspend the provider using IRP Frontend “Providers and Peers” configuration section. IRP withdraws all existing and stops new improvements towards shutdown provider(s).

2. You can remove traffic from the provider (by denying incoming/outgoing announces in the BGP configuration) and then physically re-cable the provider from one router to another. Then configure BGP session towards the provider on new router.
3. The PBR rule for the provider should be configured on new router (move provider's PBR settings from the old router to the new one).
4. Change IRP box local PBR rules if any.
5. Check if PBR works properly using traceroute tool.
6. Modify (/etc/noction/irp.conf) assigned router for the provider (refer to `peer.X.bgp_peer`).
7. Modify (/etc/noction/irp.conf) SNMP interface/IP/community for the provider (refer to `SNMP Host`, `peer.X.snmp.interfaces`, `peer.X.mon.snmp`).
8. Reload Bgpd configuration (affected BGP sessions could be reset).
9. Re-activate the provider using IRP Frontend “Providers and Peers” configuration section
10. Check IRP BGP Internal and External monitor states. They must be UP.

3.12.8 SNMP hosts configuration

⚠ Starting with version IRP 3.7 SNMP is configured as SNMP hosts and provider specific SNMP settings are being deprecated.

SNMP hosts are nodes on the network that provide or read SNMP data. Multiple SNMP hosts are configured in this section of configuration and references to them are used from elsewhere in the configuration. SNMP hosts can be configured by navigating via the “Configuration→SNMP hosts” menu.

See also: [SNMP Host](#)

The screenshot shows a configuration window titled "host1" with a tab labeled "SNMPv3". The window contains the following fields and options:

- SNMP host short name***: Text input field with value "snmphostname".
- SNMP host address***: Text input field with value "185.34.200.17".
- SNMP community**: Text input field with value "n0c-irp".
- SNMP v3 security level**: Dropdown menu with value "Authentication and privacy".
- SNMP v3 authentication password**: Text input field with a "Show password" checkbox.
- SNMP v3 authentication protocol**: Dropdown menu with value "MD5".
- SNMP encryption password**: Text input field with a "Show password" checkbox.
- SNMP encryption**: Dropdown menu with value "DES".
- SNMP v3 Username**: Text input field.

Figure 3.12.28: Configuration editor: SNMP hosts configuration

- SNMP-related settings:
 - Flag indicating SNMP supported version (`snmp.X.version`)

- SNMP host short name that assigns a recognizable name of the SNMP host in IRP (`snmp.X.name`)
- SNMP host IPv4/IPv6 address (`snmp.X.ip`)
- SNMP v2 uses additional parameter:
 - SNMP host community (`snmp.X.community`)
- SNMP v3 uses additional parameters depending on security services used:
 - SNMP security level selects if Authentication and/or Privacy services are used (`snmp.X.seclevel`)
 - SNMP authentication password if authentication is used (`snmp.X.auth_password`)
 - SNMP authentication protocol if authentication is used (`snmp.X.auth_protocol`)
 - SNMP encryption password if privacy is used (`snmp.X.priv_password`)
 - SNMP encryption protocol if privacy is used (`snmp.X.priv_protocol`)
 - SNMP Username if authentication is used (`snmp.X.auth_username`)

3.12.9 Notifications and Events

Notifications are messages sent to alert about some events registered by IRP. There are two prerequisites in order to start using notifications:

1. Configure event parameters and channels
2. Subscribe for event notifications

IRP can send notifications as SMS, email and SNMP Trap.

3.12.9.1 Configure notification and events

Events configuration changes default threshold values when an event is raised. The most relevant events are presented in the following figure. For example Overloaded by (%) indicates that this event will fire when the aggregated outbound bandwidth usage for all providers exceeds by 10% the configured limits. Refer section 5.12 for details.

The screenshot shows the 'Events configuration' window with the following settings:

Configuration Item	Value
Provider outbound overload by (%)	10
Provider outbound overload by (Mbps)	100
Overload by (%)	10
Overload by (Mbps)	100
Inbound overload by (Mbps)	100
Inbound overload by (%)	10
Provider inbound overload by (Mbps)	100
Provider inbound overload by (%)	10

At the bottom, there is a link to 'Show advanced settings'.

Figure 3.12.29: Configuration editor: Events configuration

IRP uses a local email server to send emails. Still, email sent by this server might trigger filtering policies enforced by some third parties that might block some important event notifications. It is possible to provide the details of another email server on your network that is protected for this. The following figure highlights the available email channel configuration parameters. Besides specifying the server email configuration sets the sender of email messages that will show in receiver's inbox.

i Only SMTP servers without authentication are currently supported.

Figure 3.12.30: Configuration editor: Email configuration

SMS configuration is mandatory for sending SMS notifications.

IRP supports the following SMS gateways:

- [Twilio](#)
- [Plivo](#)

A valid account with a supported SMS gateway is required in order to finish configuration. Check the following figure for details:

Figure 3.12.31: Configuration editor: SMS configuration

The settings are:

- SMS gateway - choose one of the supported gateways.
- Account ID - the public identifier assigned to your account by the SMS gateway. The following figure highlights this value for Plivo.
- Max message size - the maximum length of the SMS text. The SMS will be trimmed by IRP before sending. Subsequently the SMS gateway will split the SMS into multiple parts if required.
- Secret - the secret identifier assigned to your account by the SMS gateway. The following figure highlights this value for Plivo.
- From phone number - the phone number that shows as sender on the receiver's mobile device. Use a phone number supported by the SMS gateway.

Refer section [5.12](#) for complete details about configuration parameters.

⚠ Some SMS gateways enforce the From phone number and will reject numbers that do not comply with their policy.




Figure 3.12.32: Plivo account ID and secret

SNMP Traps configuration requires a list of destination IP addresses and an SNMP community parameter to deliver SNMP traps. The following figure highlights the basic configuration. There is a series of advanced configuration parameters like the SNMP Trap port and rate of packets. Refer section 5.12 for details.



Figure 3.12.33: Configuration editor: SNMP Traps configuration

Web Hooks configuration requires only a Webhook URL to be provided. There is a series of advanced configuration parameters in case their default settings are not satisfactory. Refer section 5.12 for details.



Figure 3.12.34: Configuration editor: Web Hooks configuration

The settings are:

- Webhook URL - the unique URL the team is assigned by Web hook provider. An example for Slack.com is shown.
- Avatar icon URL - optional parameter that designates an icon (usually PNG or JPEG image are supported) assigned to the Webhook bot in the channel.
- Avatar emoji - an alternative avatar specified in the form of an emoji in “:robot-face:” format.
- Bot name - optional bot name assigned to Webhook bot.

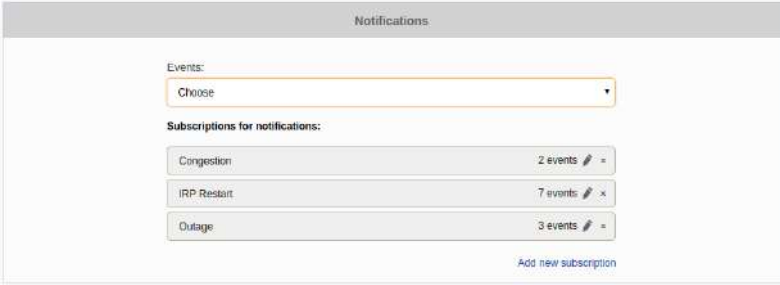
Refer section 5.12 for complete details about configuration parameters.

⚠ Web hooks support has been tested and verified to work for Slack.com API.

3.12.9.2 Subscriptions

Notifications are sent only if a valid subscription has been created. Subscriptions, similar to regular emailing of reports can be found under <Username> → Notifications menu.

Notifications page provides an overview of existing subscriptions with features to create, edit, delete them. Check the following figure for a preview.



The screenshot shows a web interface titled "Notifications". It features an "Events:" section with a dropdown menu currently set to "Choose". Below this is a section titled "Subscriptions for notifications:" which contains a list of three subscriptions: "Congestion" (2 events), "IRP Resort" (7 events), and "Outage" (3 events). Each subscription entry has an edit icon (pencil) and a delete icon (X). At the bottom right of the list is a link that says "Add new subscription".

Figure 3.12.35: Notifications

Use Add new subscription or Events drop-down to create a new subscription or edit, delete existing subscriptions by clicking corresponding icons. Creating or editing a subscription will bring up a pop-up like the one in the following figure.

Figure 3.12.36: Subscribe notifications

Subscription details include:

- Topic - the title given to subscription in Notification page and also a textual part of email and SMS notifications.

i If SMS notifications are used it is advised to keep the topic short so that there is enough space left to include details about the event.

- Interval between notifications - sets a rate limit of how often an email and/or SMS notification is sent. Value 'No limit' for the interval depicts no rate limits and all events will trigger a notification.

i SNMP Traps notification are not constrained by the interval between notifications. SNMP Traps are sent immediately as they are raised.

a The rate limit is enforced per subscription so it is still possible to receive multiple notifications if the subscribed events are part of different subscriptions.

- Destinations - IP address of a trap receiver for SNMP Trap notifications, email addresses, phone numbers and webhook channels. Multiple destinations of same type can be provided. At least one valid destination must be provided per subscription.

i IRP parses and recognizes whether a provided destination is a valid IP address, email address, phone number or web hook channel. Web hook channels for example are recognized by detecting that the first character of a channel name is the “#” mark.

- Event filters - allows filtering of events by textual description or by IRP component. Remove the filters to see all the subscribed events.
- Subscribed events - the full list of events supported by IRP and tick marks for the events in this subscription. At least one event is required for a valid subscription.

3.12.9.3 Notification text

When a subscribed event is fired IRP will send notifications. A sample notification text (email) looks like the following figure:



Figure 3.12.37: Subscribe notifications

The email sample above identifies the different parts of the message:

1. IRP server name
2. Subscription topic as specified in the subscription
3. From email address - as configured in email channel
4. Time - date-time of the last event that caused the notification. In case of rate limitation this might be older than the time of the email.
5. Textual description of the event and any relating details.
6. Older events in this subscription that have been retained due to rate limitations. Keep in mind that all past events are aggregated into a single email or SMS message.

3.12.10 Security Configuration

Security configuration covers restrictions on the ranges of IP addresses that can access IRP's Frontend and also user's allowed to access IRP as well as their roles.

3.12.10.1 Allowed ranges of IP addresses

IRP Frontend includes a feature to restrict access by IP address and this feature can be turned on or off as needed. When this feature is used allowed IP addresses are maintained as a list of prefixes in CIDR notation.



Figure 3.12.38: Configuration editor: Security

3.12.10.2 Users and user directories

Users who can access IRP are managed in User Directories. User directory is a generic name for a external user management provider for example LDAP (either POSIX or Active Directory). IRP can interface with as many user directories as needed in a specific environment.

IRP recognizes Admin and User privileges that can be granted to users. Once a user logs into the system IRP creates a profile that maintains his personal configurations like dashboards or notifications.

Users					
#	Username	Email	Privileges	User Directory	+ Add
1	admin	admin@noction.com	Admin	Internal	
2	noction	support@noction.com	User	Internal	
9	windowsUser	winUser@NOCTION.local	Admin	AD Test	
11	ldapUser	ldapUser@noction.com	Admin	LDAP TEST	

User Directories			
Name	Type	Hostname	+ Add
LDAP TEST	LDAP POSIX	ldap.noction.com	
AD Test	Active Directory	23.45.54.32	
Default internal user directory (MySQL)	Internal	-	

Figure 3.12.39: Configuration editor: Users and Directories

3.12.10.3 LDAP and Active Directory

LDAP or AD user directories can be added, updated and removed from IRP by accessing “Configuration→Security” menu item in the User Directories tab. Each user directory takes a series of parameters specific for the protocol.

All operations with DN's (initial bind DN, group DN's, user names) are case insensitive and also strip redundant whitespace.

Refer individual protocol documentation for how to correctly configure one or another user directory. The example below offer a generic set of parameters required to configure IRP to use Active Directory for access management.

The screenshot shows the 'Edit user directory LDAP TEST' window with the 'General' tab selected. The window has three tabs: 'General', 'Advanced', and 'Bindings'. The 'General' tab contains the following fields and controls:

- User directory name:** A text field containing 'LDAP TEST'.
- Type:** A dropdown menu showing 'Active Directory'.
- State:** Two buttons, 'DISABLED' and 'ENABLED', with 'ENABLED' being the active state.
- Order:** A slider control ranging from 0 to 255, currently set at 0.
- Hostname:** A text field containing 'ldap.noction.com'.
- Port:** A text field containing '389'.
- Save button:** A button with a checkmark icon and the text 'Save'.

Figure 3.12.40: Configuration editor: LDAP User Directory configuration

The general tab covers:

- User directory name - the name assigned to the directory within IRP,
- User directory type - one of the supported User Directories,
- Enabling or disabling a user directory,
- Order specifies when this user directory will be examined by IRP compared to other user directories,
- User directory hostname in the form of either IP address or domain name
- User directory port

The screenshot shows the 'Edit user directory LDAP TEST' window with the 'Advanced' tab selected. The window has three tabs: 'General', 'Advanced', and 'Bindings'. The 'Advanced' tab contains the following fields and controls:

- Initial bind DN:** A text field containing 'cn=LDAP TEST,ou=app-accounts,dc:'.
- Initial bind Password:** A password field with masked characters '*****' and a 'Show password' checkbox.
- Timeout:** A slider control ranging from 0 to 5, currently set at 5.
- TLS:** Two buttons, 'DISABLED' and 'ENABLED', with 'ENABLED' being the active state.
- Certificate verification:** Two buttons, 'DISABLED' and 'ENABLED', with 'ENABLED' being the active state.
- TLS CA Certificate file:** A text field.
- Save button:** A button with a checkmark icon and the text 'Save'.

Figure 3.12.41: Configuration editor: Users and Directories

Advanced configuration of a user directory covers:

- Initial binding user name that IRP uses to authenticate itself,
- Initial bind password assigned to IRP,

i Usually bind passwords are not required to access directories. If a password is required and configured via IRP Frontend, it will be scrambled in IRP configuration files. If the password is specified directly in IRP configuration it is provided in clear-text with the condition that it does not start/end with scrambled password encapsulation characters.

- Timeout before failing a connection to this user directory,
- TLS use,
- Certificate verification and
- CA certificate used to verify server's certificate.

Figure 3.12.42: Configuration editor: Users and Directories

Bindings maps User Directory attributes to IRP specific attributes, for example:

- Base DN and User DN specifies the root distinguished name and user subtree

i IRP recognizes BOTH short and full user identifiers. Examples below are both valid directory entries that will match user “chris” with long name “Mr. Chris Smith”:

```
cn: ops
uniqueMember: chris
```

- Admin role groups are user directory objects that list users with Administrator privileges within IRP as compared to all users. Both short and full paths are accepted for Admin role groups.
- User role, Username, Email and Common name fields map User Directory attributes to IRP user attributes,
- User filters can specify any valid LDAP search criteria (without outside parenthesis).

3.12.10.4 TACACS+ directory

TACACS+ user directories can be added, updated and removed from IRP by accessing “Configuration→Security” menu item in the User Directories tab. Each user directory takes a series of parameters specific for the protocol.

The example below highlights the parameters required to configure IRP to authenticate against a TACACS+ directory/host.

The screenshot shows a configuration window titled "New user directory TAC". It has a "General" tab. The fields are as follows:

- User directory name:** TAC
- Type:** TACACS+
- State:** ENABLED (with a DISABLED button also visible)
- Order:** 42 (with a slider from 0 to 255)
- Hostname:** 1.2.3.4
- Port:** 49
- Secret:** asecret

At the bottom right, there are "Check" and "Save" buttons.

Figure 3.12.43: Configuration editor: TACACS+ configuration

The parameters include:

- User directory name - the name assigned to the directory within IRP
- User directory type - one of the supported User Directories in our case TACACS+
- Enabling or disabling a user directory
- Order specifies when this user directory will be examined by IRP compared to other user directories

- User directory hostname in the form of either IP address or domain name
- User directory port
- A shared secret used by IRP and TACACS+ host to secure communications.

Configuring a user account in the TACACS+ server IRP uses “service=irp” in the authorization request and checks the response for “priv-lvl=15” or “admin=true” to verify if the account has administrator privileges.

See also: [User Directory](#).

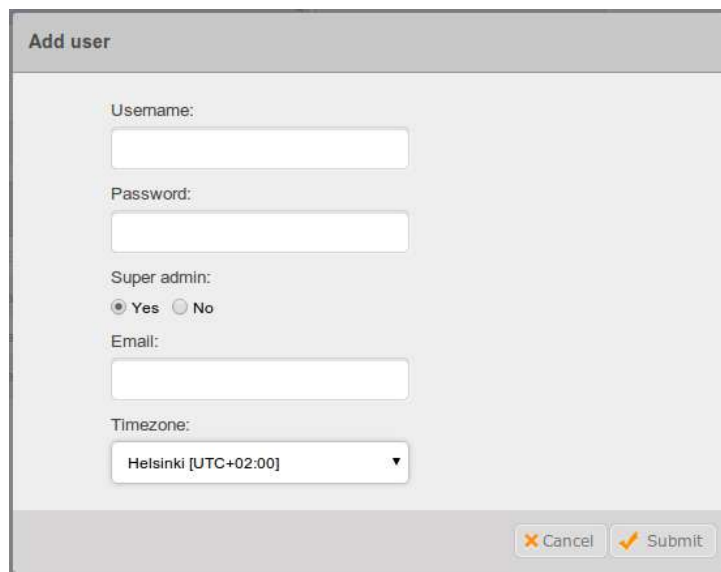
3.12.10.5 Internal user directory

IRP includes an Internal user directory that allows multiple user accounts to be setup. User accounts management is done by accessing “Configuration→Security” menu item.

i User management for external user directories is performed outside IRP using designated tools. The Internal directory is used when external user directories are not available or their use is not desirable.

To edit a specific user’s profile, press the Edit button next to it. To add a new user, click on the “Add” button. Specific users may be given Admin privileges. In such case they can also manage all system users.

The timezone specified in the user profile will be considered while displaying system Reports and Graphs.



The screenshot shows a web form titled "Add user". It contains the following fields and controls:

- Username:** A text input field.
- Password:** A text input field.
- Super admin:** Radio buttons for "Yes" (selected) and "No".
- Email:** A text input field.
- Timezone:** A dropdown menu currently showing "Helsinki [UTC+02:00]".
- Buttons:** "Cancel" and "Submit" buttons at the bottom right.

Figure 3.12.44: Configuration: Adding a new user account

3.12.11 Frontend configuration

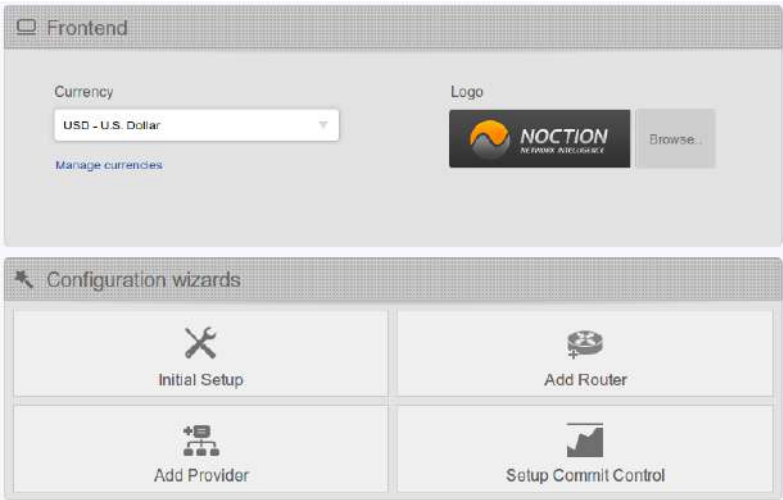


Figure 3.12.45: Configuration editor: Frontend

The currency used in various cost-related reports as well as the system logo (shown in the login form and on the menu bar) can be changed on the same page.
Currencies can be added/deleted manually by using the “Manage currencies” feature.

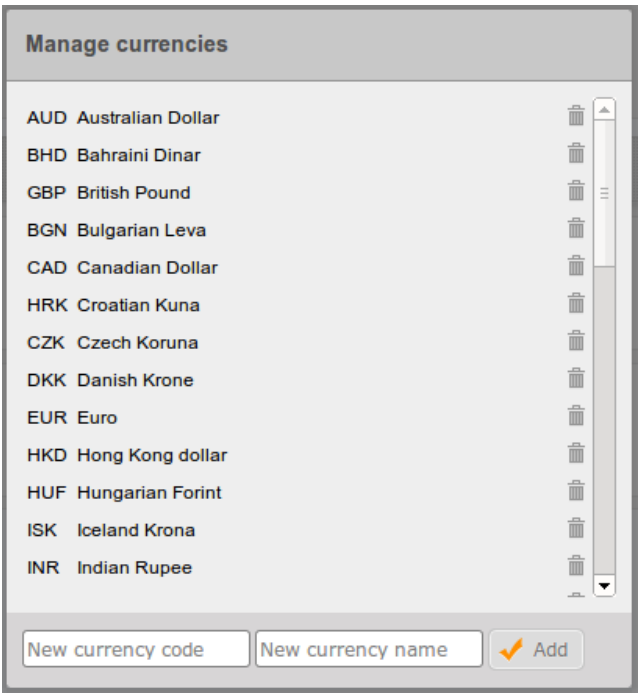


Figure 3.12.46: Configuration Editor: Manage currencies

⊖ Only a Super admin user can change the currency setting. The currency is a global system value, and cannot be adjusted per-user.
The system logo is a global setting and can be changed by the Super administrator only.

i The default currency is “USD - U.S. Dollar”.

A custom logo image can be uploaded by using the Upload button. Afterward, the image can be cropped to fit the standard logo size. The new settings are applied after clicking the “Submit” button.

3.12.12 Inbound configuration

Refer [Inbound bandwidth optimization](#) for an overview of Inbound feature.

- Flow (sFlow, NetFlow) is a prerequisite for Inbound optimization. IRP collects and uses statistics about upstream provider incoming traffic targeting inbound prefixes. This statistics should be made available by properly setting up Flow and provider specific Flow agents (refer [peer.X.flow_agents](#)).

i SPAN does not carry details about the specific interface packets enter your network and as such it is impossible to distribute inbound traffic targeting Inbound prefixes by upstream provider.

All Inbound related features, including configuration, moderation, and reports are listed under main menu Inbound:

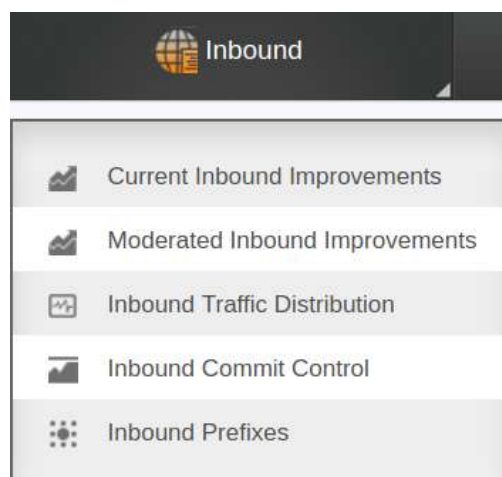


Figure 3.12.47: Main menu: Inbound

3.12.12.1 Inbound prefixes

i Note that optimization of transiting traffic feature does not rely on manually configured transit prefixes. Instead filters by ASN and/or prefix are used to constrain the specific prefixes that can be improved.

Inbound optimization manipulates how much traffic targets YOUR prefixes through different providers. In order to do this you need to configure IRP and let it know what are the Inbound prefixes. You can specify a subset of your network prefixes and you can also split larger prefixes into smaller ones in order to have a more fine-grained control over the Inbound traffic. These prefixes will be announced back to your originating routers marked with the community designated to prepends count to use when announcing to individual upstream providers.

➡ Refer for details regarding IRP announcement of inbound prefixes.

❗ Splitting larger prefixes into subprefixes offers IRP the opportunity to better control the granularity of traffic shaping changes enabling improvement of smaller chunks of overall traffic. At the same time original prefixes can be announced as before and they will guarantee continuity of the prefix for cases when improvements announced by IRP become unavailable for some reason (for example if smaller prefixes are filtered in some networks).

➡ Verify during Inbound configuration that the split inbound prefixes are not filtered by your providers and propagate as expected on the Internet.

IRP Frontend includes facilities that let you define the list of YOUR prefixes covered by Inbound traffic optimization as highlighted below and includes:

- prefix belonging to the network (`inbound.rule.X.prefix`),
- router(s) where announcements are sent (`inbound.rule.X.bgp_peer`),

➡ Inbound improvements must be advertised by IRP to core routers that originate the prefix. This way the improvements will propagate consistently through your network to edge routers and beyond.

- `prefixstatus` as each prefix can be disabled in order to exclude from further inbound optimization (`inbound.rule.X.enabled`),
- next hop used by improvements for this prefix (`inbound.rule.X.next_hop`),
- an option to instruct IRP whether to fully control the prefix or to only announce improvements when there are any. Refer the section for further details.
- providers for which this inbound prefix can be optimized (`inbound.rule.X.providers`). It must be noted that if no provider is selected then the prefix is optimized for all providers. Otherwise, the prefix will be optimized only through selected providers.

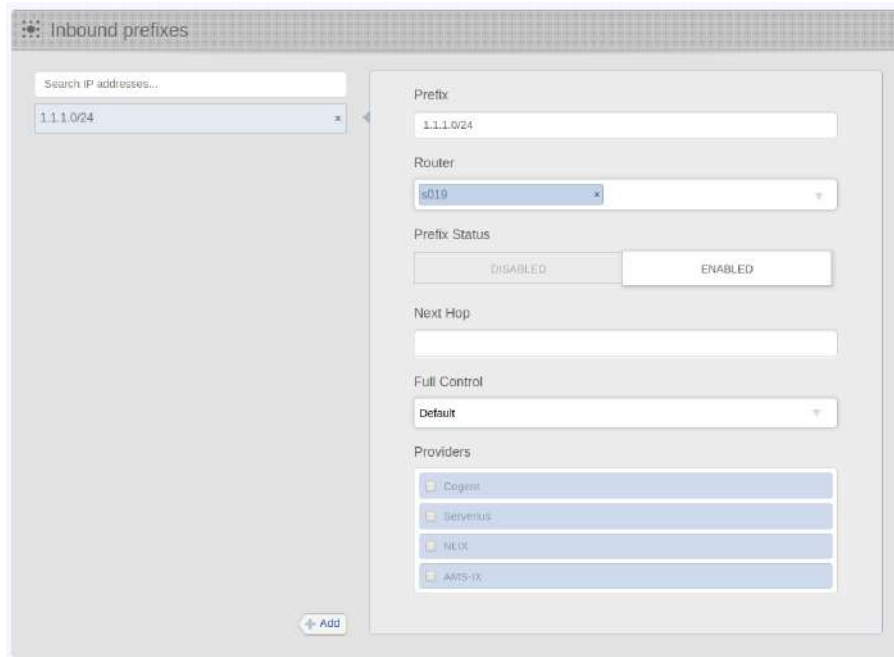


Figure 3.12.48: Configuration editor: Inbound prefixes

i Inbound prefixes can be found both via Inbound main menu or Configuration →Collectors

Inbound prefixes full control IRP operates under the assumption that the network already originates all network’s prefixes and only a subset are configured in IRP for inbound optimization purposes. This subset includes the intentionally split prefixes that allow better bandwidth control granularity by IRP. Under these assumptions IRP announces to the network only the improvements it makes. This avoids the overhead of announcing too often without real benefit for the network.

On some networks there might be a preference to guarantee that an inbound prefix is announced to all (allowed) upstreams at the same time thus making sure that they all have identical knowledge about the route with possibly different number of prepends.

Improvements Announce only improved inbound prefixes with communities for improved providers

If improved Announce only improved inbound prefixes with communities for all allowed providers

All Always announce all inbound prefixes with communities for all allowed providers

Options “If improved” and “All” are useful when the network itself does not originate a prefix, then the inbound prefix announced by IRP fully describes the route and provider preference to all upstream providers thus ensuring that all of them can route our traffic and the corresponding prepends count will influence what providers are preferred.

To enable inbound prefix full control either individual inbound prefix explicit value should be set in Configuration →Collectors →Inbound Prefixes

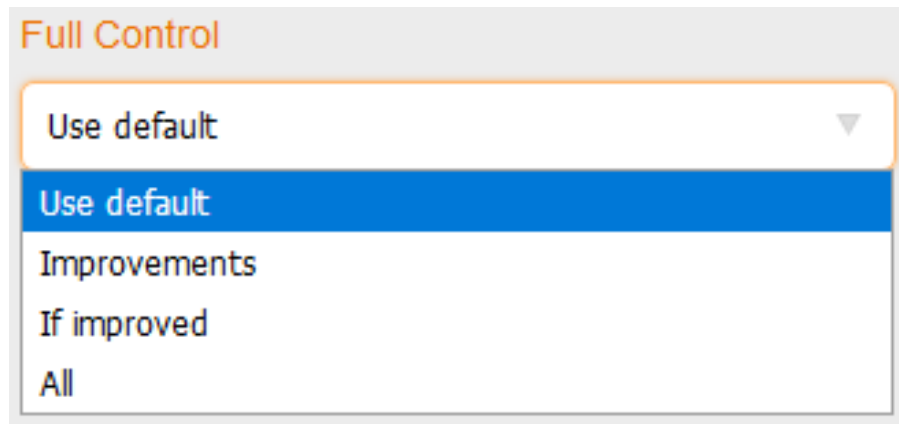


Figure 3.12.49: Configuration editor: Control of individual inbound prefix

or default IRP behavior should be specified in Configuration →BGP and Routers




Figure 3.12.50: Configuration editor: Inbound prefixes control

3.12.12.2 Inbound prepends

Inbound optimization in IRP relies on BGP priority rules that generally elect best routes based on AS PATH length. By instructing your network to add or remove prepends for specific prefixes over different providers, IRP is able to influence how your inbound traffic flows. Inbound optimization in IRP is less effective for those regions of the Internet where your providers or other networks on the Internet use other policies for electing best routes. For example, Internet Exchange routes always take precedence for peers. This makes impossible prepend based inbound traffic optimization for Internet Exchanges. Often partial providers employ similar policies and inbound traffic flowing through them cannot be optimized. Prepends are not added directly by IRP. Instead IRP inbound improvements are marked with designated communities that are recognized and acted upon by network routers.

Base communities Inbound optimization improvements advertise larger or smaller counts of prepends to be announced to different providers. IRP uses BGP session to communicate to routers and relies on BGP communities to communicate the improvements. Each provider is assigned a base community which represents zero prepends. Additional prepends are represented by incrementing accordingly the right-side part of the provider's community.

⚠ IRP uses 8 prepend levels. This means that base community should be chosen to allow additional 7 values without intersecting with other community values.



Figure 3.12.51: Configuration editor: Base communities

i Inbound base communities are assigned for each individual provider Configuration → Providers and Peers

Additionally routemaps are applied so that routers depict the designated communities in IRP's improvements and change the announcement of the inbound prefix towards that provider appropriately. Different routemaps are produced and applied on each router in your network. See a few examples below. Request assistance from Noction's support team to prepare exact routemaps for your network configuration and router models.

3.12.12.3 Routers configuration for Inbound

w Complete sample configurations for Cisco, Juniper, Brocade and Vyatta will be provided by Noction Support on request.

3.12.12.4 Inbound Commit Control

Inbound bandwidth control is part of the Commit Control feature of IRP. You need to enable Commit Control and further to enable Inbound Commit control to make use of this feature. Besides the above mentioned settings for Inbound Commit Control separate low and high limits are configured to instruct IRP when to start and when to stop improving and also how to estimate the effects of inbound improvements.

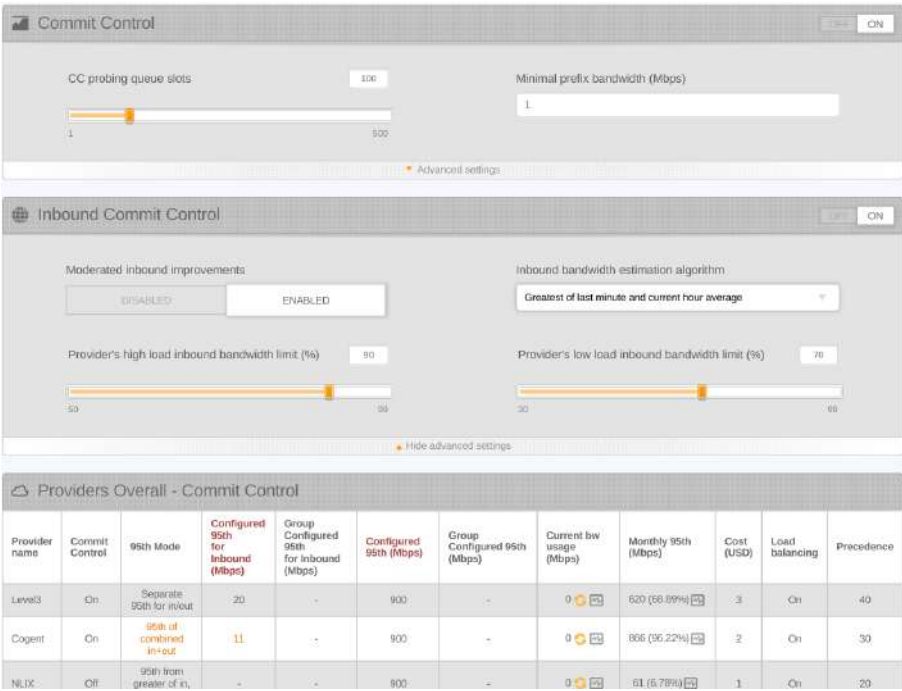


Figure 3.12.52: Configuration editor: Inbound commit control

i Inbound commit control can be found both via Inbound main menu or Configuration →Commit Control

Commit control for inbound as can be seen above screencapture highlights the relevant configuration parameters:

- whether inbound commit control is on or off
- whether review and moderation feature is enabled or disabled
- what is the bandwidth estimation algorithm for inbound and
- high and low limits for inbound in case the 95th mode warrants independent inbound vs outbound optimization.

Refer also to `peer.X.95th.mode`, `core.commit_control.inbound.enabled`, `core.commit_control.inbound.moderated`.

Estimating inbound traffic Internet traffic has high variability and adjacent measurements can differ significantly between them. In these conditions a leveling function helps mitigate the risk of generating excessive numbers of Inbound improvements due to higher traffic variability of some networks. This estimation can be fine tuned by the Inbound bandwidth estimation algorithm parameter that tells IRP what value to use as the basis of the calculation. See the possible base values below:




Figure 3.12.53: Configuration editor: Inbound BW estimation algorithms

Moderated Inbound Improvements IRP includes a moderation feature that allows review and approval of Inbound improvements. If this feature is enabled Inbound improvements are not automatically announced but instead they are queued for review. If an inbound improvement has not been approved in time for a subsequent cycle of improvements IRP will review the proposal itself and will adjust as needed thus making another recommendation.

See the capture below.

#	Timestamp	Link1	Link2	Prefix	Type			
1	03/03/16 04:30:56 PM	2 (+1)	4 (+2)	10.1.1.0/24		Commit	Reject	
2	03/03/16 04:34:11 PM	3 (+1)	3 (+1)	10.1.2.0/24		Commit	Reject	
3	03/03/16 07:14:10 PM	-	2 (+2)	10.1.3.0/24	-	Commit	Reject	
4	03/03/16 05:01:34 PM	1	2	10.1.4.0/24		Commit	Reject	
5	03/03/16 06:23:08 PM	2	2	10.1.6.0/24		Commit	Reject	
6	03/03/16 08:49:09 PM	-	2 (+2)	10.1.7.0/24	-	Commit	Reject	
7	03/03/16 06:58:34 PM	2	4	10.1.8.0/24		Commit	Reject	
8	03/03/16 05:01:42 PM	2	2	10.1.9.0/24		Commit	Reject	

Commit selected Reject selected

Figure 3.12.54: Configuration editor: Inbound BW estimation algorithms

Inbound Improvements moderation highlights the number of additional prepends proposed by IRP and allows submitting or rejecting individual improvements as well as a batch of selected improvements.

3.12.12.5 Optimization of transiting traffic

Optimization of transiting traffic is setup as part of [Inbound Commit Control](#).

Inbound Commit Control [OFF] [ON]

Moderated inbound improvements

Transit optimization

Transit ASNs
 Click to add new asn...
 3 x
 2 x

Transit Improvements TTL min
 14401

Top transit prefixes per cycle
 101

Inbound bandwidth estimation algorithm
 Greatest of last minute and current hour average

Max transit Improvements
 101

Transit prefixes
 Click to add new prefix...
 1.2.0.0/16 x
 60.240.242.217/32 x
 2001.db8./32 x
 192.168.0.0/16 x
 172.16.0.0/12 x

Transit Improvements TTL max
 86401

Match transit at egress

Figure 3.12.55: Configuration editor: Inbound optimization of transiting traffic parameters

i Inbound commit control can be found both via Inbound main menu or Configuration → Commit Control

Optimization of transiting traffic configuration parameters are highlighted:

- Transiting traffic toggle that enables or disables the feature (`global.inbound_transit`)
- Transit Improvements Max sets the maximum number of transit improvements (`core.improvements.inbound_transit.max`)
- Transit ASNs and Transit prefixes specify those segments of the Internet that IRP monitors and optimizes (`bgpd.prefixlist.asn`, `bgpd.prefixlist.prefixes`)
- Transit Improvements TTL min and max set the lower and upper bounds in seconds to keep a transit improvement (`core.improvements.inbound_transit.ttl.min`, `core.improvements.inbound_transit.ttl.max`)

i Note that traffic belonging to a specific prefix that fragments a large transit prefix is collected as belonging only to the specific prefix and excluded from the larger prefix traffic statistics. If the specific prefix is also filtered from IRP configuration as for example is the case of a /26 prefix, then the specific prefix will be excluded from optimization of transiting traffic and will not show up in any of the subsequent decisions or reports.

- Transit Top N prefixes sets the number of transit prefixes that are collected each cycle and considered for optimization (`collector.flow.export.inbound_transit.topn`)
- Match transit at egress enables or disable statistics collection for transit prefixes when packets exist the network (`collector.flow.process_transit_in_outbound`)

3.12.13 Wizards

Wizards can be accessed from Configuration → Setup Wizards or from Configuration → Frontend → Configuration Wizards.

3.12.13.1 Initial Setup

IRP initial setup requires to:

1. Specify Infrastructure IP addresses and list of analyzed networks
2. Enable and configure Flow or Span Collector to gather network statistics
3. Set IRP Improvement mode
4. Setup the management interface
5. Indicate interfaces that IRP uses for active probing

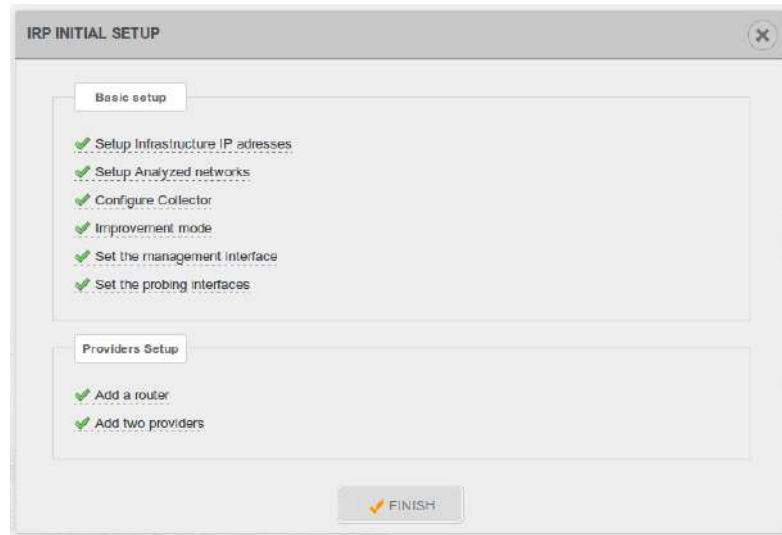


Figure 3.12.56: Configuration editor: Initial Setup

- Basic Setup
 - Setup Infrastructure IP addresses (`explorer.infra_ips`)
 - Setup Analyzed networks (`collector.ournets`)

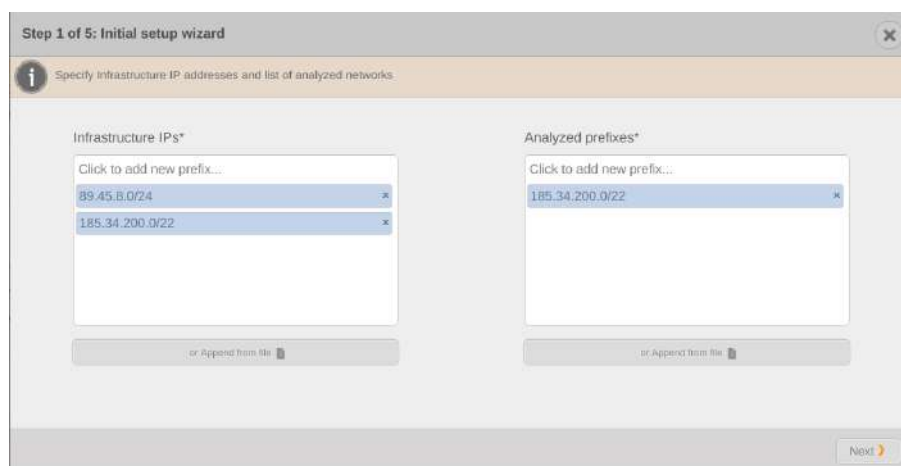


Figure 3.12.57: Configuration editor: Setup Infrastructure IP addresses and Analyzed networks

- Configure Collector:
 - * Irpflowd
 - Irpflowd enable/disable(`collector.flow.enabled`)
 - NetFlow UDP port(`collector.flow.listen.nf`)
 - sFlow UDP port(`collector.flow.listen.sf`)
 - Flow Sources(`collector.flow.sources`)
 - * Irpspand
 - Irpspand enable/disable (`collector.span.enabled`)
 - Irpspand interfaces (`collector.span.interfaces`)
 - Mindelay status (`collector.span.min_delay`)
 - Mindelay probing queue slots (`collector.span.min_delay.probing_queue_size`)

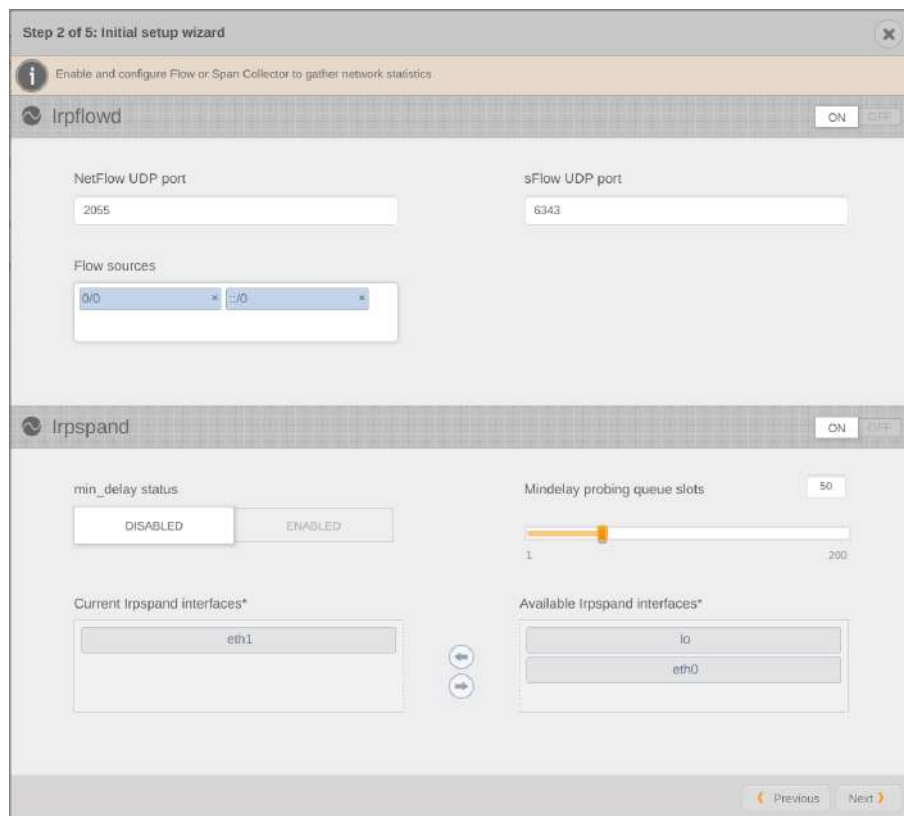


Figure 3.12.58: Configuration editor: Configure Collector

- Improvement mode (`global.improve_mode`)



Figure 3.12.59: Configuration editor: Improvement mode

- Set the managements interface (`global.master_management_interface`)

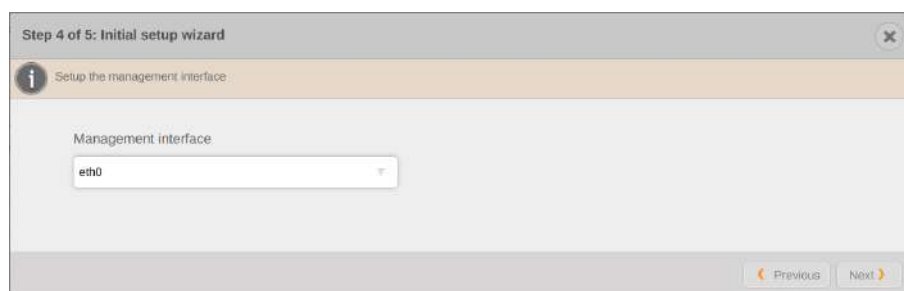


Figure 3.12.60: Configuration editor: Management Interface

- Set the probing interface (`global.master_probing_interface`)

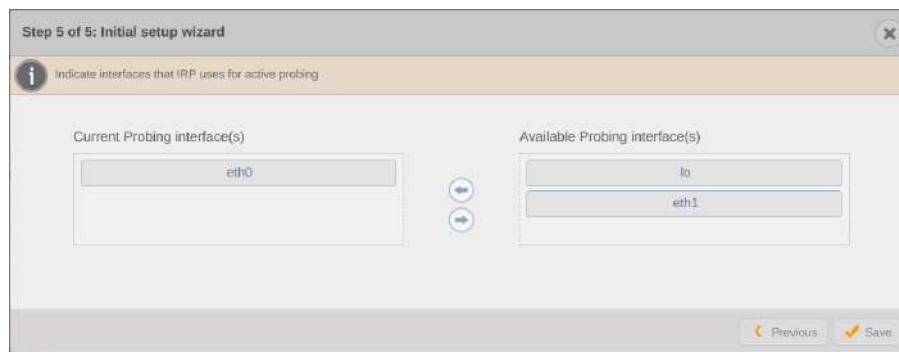


Figure 3.12.61: Configuration editor: Probing Interface

- Providers Setup
 - Add a router, see [Add Router](#)
 - Add two providers, see [Add Provider](#)

3.12.13.2 Add Router

IRP communicates improvements to your edge routers. Add Router wizard guides router configuration through the following steps:

1. Identify the router and its AS
2. Router IP address to setup BGP session
3. BGP announcement attributes to distinguish and prioritize IRP improvements on this router

Check below the corresponding wizard steps:

- Router name assigned within IRP for easy identification and
- Autonomous System number of the network (`bgpd.peer.X.as`)

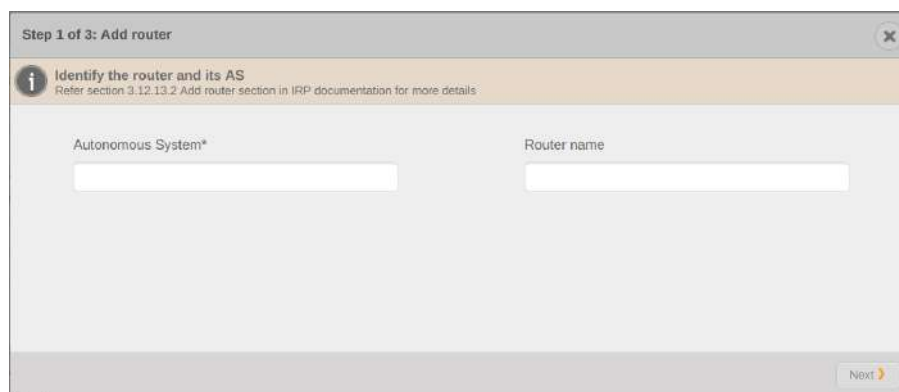


Figure 3.12.62: Configuration editor: Router name and AS

- Router IPv4 addresses
 - Local router IPv4 address (`bgpd.peer.X.master_our_ip`)
 - Router IPv4 address (`bgpd.peer.X.master_peer_ip`)

Figure 3.12.63: Configuration editor: Router IPv4 addresses

- Router announcement attributes for IRP improvements
 - Improvement LocalPref value (`bgpd.peer.X.master_localpref`)
 - Improvement communities(`bgpd.peer.X.master_communities`)
 - Improvement MED value (`bgpd.peer.X.med`)

Figure 3.12.64: Configuration editor: Router announcing

3.12.13.3 Add Provider

The wizard Add Provider covers the following:

1. Router - Identify the router and routing domain where the provider interconnects with your network
2. Provider - Specify what is a provider ASN, short name and description
3. IP addresses - Specify and assign provider network addresses that IRP will use
4. Commit Control - Optionally set provider usage thresholds to be used for Commit Control
5. SNMP host - Specify the SNMP host for this provider
6. External monitor - Indicate if an External monitor is used and designated external IP addresses used to verify this provider link status
7. Internal monitor - Indicate if an Internal monitor based on BGP session with provider will be used and the corresponding SNMP resource
8. Pre-checks - Validate given provider parameters before submitting them

- Choose a Router ([peer.X.bgp_peer](#))

Figure 3.12.65: Configuration editor: Choose a Router

- Provider name
 - Provider short name ([peer.X.shortname](#))
 - Provider description ([peer.X.description](#))

Figure 3.12.66: Configuration editor: Provider name

- Provider IPv4 addresses
 - Probing IPv4 address ([peer.X.ipv4.master_probing](#))
 - IPv4 diagnostic hop ([peer.X.ipv4.diag_hop](#))
 - Router next-hop address ([peer.X.ipv4.next_hop](#))
 - Router ASN ([peer.X.ipv4.next_hop_as](#))

Figure 3.12.67: Configuration editor: Provider IPv4 addresses

- Provider IPv6 addresses

- Provider probing IPv6 address (`peer.X.ipv6.master_probing`)
- IPv6 diagnostic hop (`peer.X.ipv6.diag_hop`)
- Router nex-hop IPv6 address (`peer.X.ipv6.next_hop`)
- Router ASN (`peer.X.ipv6.next_hop_as`)

Figure 3.12.68: Configuration editor: Provider IPv6 addresses

- Provider Commit Control

- Provider 95th percentile (`peer.X.95th`)
- Provider cost per Mbps (`peer.X.cost`)
- Commit Control provider precedence (`peer.X.precedence`)
- Maximum allowed load per interface (`peer.X.limit_load`)
- Commit Control status for this provider (`peer.X.cc_disable`)
- Allow Performance/Cost Improvements within provider group members (`peer.X.improve_in_group`)

Figure 3.12.69: Configuration editor: Provider Commit Control

- Provider Monitoring setup IPv4
 - Provider SNMP interfaces ([peer.X.snmp.interfaces](#))
 - SNMP host settings ([SNMP Host](#))

Figure 3.12.70: Configuration editor: Provider Monitoring setup IPv4

- Provider Monitoring setup IPv6
 - Provider SNMP interfaces ([peer.X.snmp.interfaces](#))
 - SNMP host settings ([SNMP Host](#))

Figure 3.12.71: Configuration editor: Provider Monitoring setup IPv6

- External Monitor setup IPv4
 - External Monitor enable/disable (`peer.X.mon.ipv4.external.state`)
 - Router monitoring IPv4 address(`peer.X.ipv4.mon`)




Figure 3.12.72: Configuration editor: External Monitor setup IPv4

- External Monitor setup IPv6
 - External Monitor enabled/disabled (`peer.X.mon.ipv6.external.state`)
 - Router monitoring IPv6 address`peer.X.ipv6.mon`)



Figure 3.12.73: Configuration editor: External Monitor setup IPv6

- Internal Monitor setup
 - Internal Monitor enable/disable (`peer.X.mon.ipv4.internal.state`)
 - Monitoring IPv4 address (`peer.X.mon.ipv4.bgp_peer`)

Figure 3.12.74: Configuration editor: Internal Monitor setup

- Provider pre-checks

Figure 3.12.75: Configuration editor: Provider pre-checks

3.12.13.4 Setup Commit Control

Step 1: Enable Commit Control Providers that should maintain a target 95th limit are identified and the exact limits are set. Link upper limit represents 80-90% of interface capacity and will be pre-populated for you.

- Commit control enables/disables the feature for a specific provider([peer.X.cc_disable](#))
- 95th target specifies the contracted bandwidth usage target in Mbps ([peer.X.95th](#))
- Link upper limit specifies the maximum allowed bandwidth on the link in Mbps ([peer.X.limit_load](#)). This represents ~80-90% of interface capacity and will be pre-populated for you. Change it as appropriate.




Figure 3.12.76: Configuration editor: Commit Control basic setup

Step 2: Provider precedence Precedence indicates how important it is to keep the commit levels for different providers. Traffic will be unloaded to the provider with least precedence in situations when all other providers are overloaded. If two or more providers have equal precedence they will form groups. Move the slider or use drag and drop to specify the order of providers from most important to least important.



Figure 3.12.77: Configuration editor: Commit Control provider precedence setup

Step 3: Groups configuration Groups are formed by a set of providers (as configured in IRP) that share some common characteristics, for example a redundant link or a virtual link consisting of multiple physical links towards the same upstream provider. Depending on customer's needs, groups can be configured to permit or forbid performance rerouting from one provider in the group towards another. If providers are grouped they will be assigned the same precedence. Groups are optional and can be skipped.

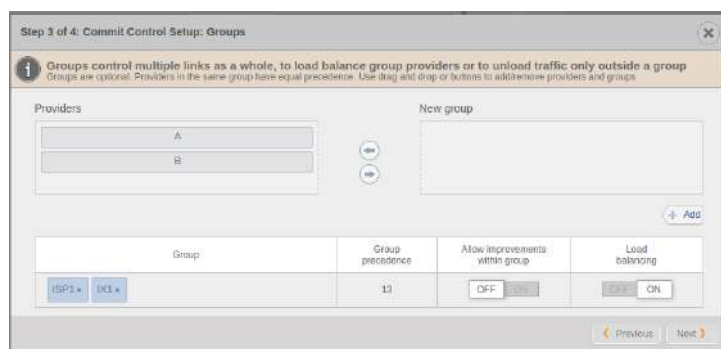
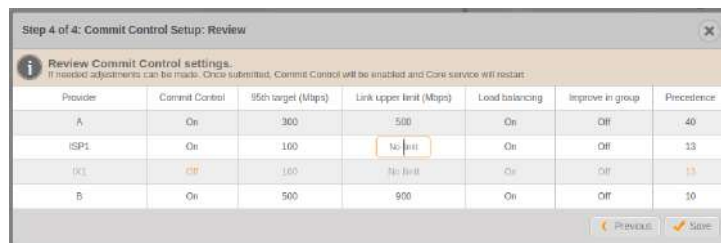


Figure 3.12.78: Configuration editor: Commit Control groups setup

Step 4: Commit Control Overview The last step of the wizard is there to review the configuration. Last minute changes can be made to some parameters. When submitted the configuration changes will be validated and applied if correct.

- Provider name (`peer.X.shortname`)

- Commit control status for a specific provider (`peer.X.cc_disable`)
- Configured 95th (`peer.X.95th`)
- Upper limit (`peer.X.limit_load`)
- Load balancing for groups
- Improve in group (`peer.X.improve_in_group`)
- Precedence (`peer.X.precedence`)



Provider	Commit Control	95th target (Mbps)	Link upper limit (Mbps)	Load balancing	Improve in group	Precedence
A	On	300	500	On	Off	40
ISP1	On	100	No limit	On	Off	13
DC1	Off	100	No limit	On	Off	13
B	On	500	900	On	Off	10

Figure 3.12.79: Configuration editor: Commit Control Provider overall

3.12.13.5 Setup Failover wizard

Step 1: Enable failover Failover configuration has a series of mandatory global settings before enabling it.

- Failover status (`global.failover_role`)
- Slave IP address (`global.failover_slave.ip`)
- Slave SSH port (`global.failover_slave.port`)
- Failover timer in seconds (`global.failover_timer_fail`)
- Failback timer in seconds (`global.failover_timer_failback`)

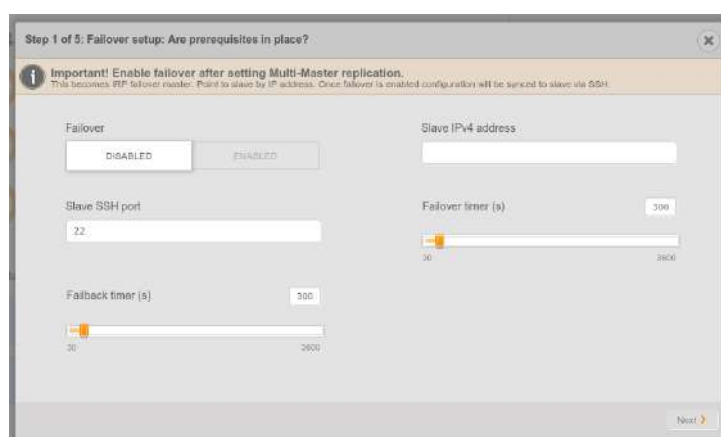


Figure 3.12.80: Configuration editor: Failover setup

Step 2: Probing IP configuration Probing IPs must be assigned to all configured providers for both master and slave nodes (`peer.X.ipv4.master_probing`, `peer.X.ipv4.slave_probing`).

Provider	Master Probing IP	Slave Probing IP
A	192.168.123.44	192.168.123.78
B	192.168.123.190	192.168.123.79
ISP1	192.168.123.94	192.168.123.95
IX1	192.168.123.80 192.168.123.81 192.168.123.82 192.168.123.83 192.168.123.84 192.168.123.85	192.168.123.87 192.168.123.88 192.168.123.89 192.168.123.90 192.168.123.91 192.168.123.92

Figure 3.12.81: Configuration editor: Failover probing IPs for providers

Step 3: Router BGP session settings Master and slave nodes of a failover configuration establish BGP sessions with all configured routers. The following parameters are required:

- Local IP address for both master and slave (`bgpd.peer.X.master_our_ip`, `bgpd.peer.X.slave_our_ip`)
- BGP session password for master and/or slave if any (`bgpd.peer.X.master_password`, `bgpd.peer.X.slave_password`)

Router	Master Local IP Address	Master Password	Slave Local IP Address	Slave Password
s019	192.168.123.42	-	192.168.123.77	-
s019v6	-	-	8.8.8.8	announcements
s020	192.168.123.1	-	192.168.123.2	-
test	192.168.123.104	-	192.168.123.105	-
test2	3.3.3.3	-	3.3.3.4	-

Figure 3.12.82: Configuration editor: Failover BGP session settings

Step 4: BGP LocalPref and Communities Master and slave nodes apply BGP attributes to announcements:

- LocalPref for master and slave (`bgpd.peer.X.master_localpref`, `bgpd.peer.X.slave_localpref`)
- Communities for master and slave (`bgpd.peer.X.master_communities`, `bgpd.peer.X.slave_communities`)

❌ Master's LocalPref value must be greater than slave's LocalPref.

Router	Master Localpref	Master Communities	Slave Localpref	Slave Communities
s019	300	65534:1	130	65503:1
s019v6	200	-	130	-
s020	150	-	130	-
test	140	1111:2222	130	1111:2222
test2	140	222:333	130	222:333

Figure 3.12.83: Configuration editor: BGP announcement attributes

Step 6: Probing interfaces Failover requires configuration of Management and Probing interfaces for both master and slave nodes (`global.master_probing_interface`, `global.slave_probing_interface`). Toggle between master and slave settings to specify values for both nodes. Toggle between master and slave settings to specify values for both nodes.




Figure 3.12.84: Configuration editor: Probing and management and interfaces

Chapter 4

Using IRP Global Management Interface

4.1 Getting started

The **GMI Frontend** represents a web interface with a comprehensive set of reports, graphs and diagnostic information which can reflect the current and historical network state, as well as the network optimization benefits delivered by IRP instances.

4.1.1 Accessing the system

The system frontend is accessible over HTTPS. GMI's Frontend can be accessed using any major browser, via the server IP address.

When opening the GMI Frontend's for the very first time, a login form will pop up asking you to create a user.

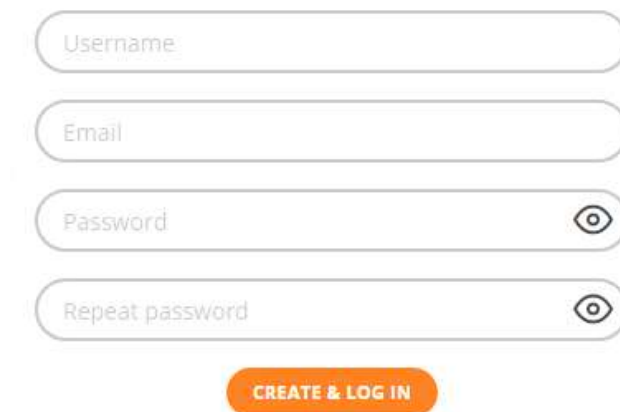
The image shows a web form for creating a new user. It consists of four input fields stacked vertically: 'Username', 'Email', 'Password', and 'Repeat password'. Each field has a light gray border and rounded corners. The 'Password' and 'Repeat password' fields include an eye icon on the right side to toggle password visibility. Below the input fields is a single orange button with the text 'CREATE & LOG IN' in white capital letters.

Figure 4.1.1: Creating initial user in GMI

Next, register IRP instance(s) in GMI to proceed (Introduce any of your IRP instances IP address along with its access credentials):

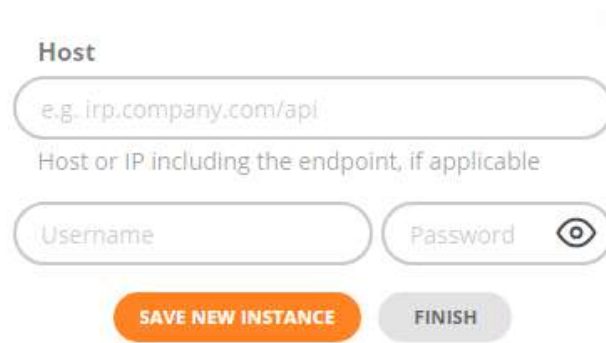


Figure 4.1.2: Registering IRP instance in GMI

On consecutive visits the GMI username and password need to be provided to login.

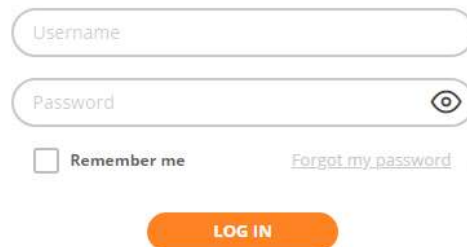


Figure 4.1.3: GMI Frontend login form

4.2 Frontend structure

After a successful login, the system's dashboard will be displayed. The Frontend has several components that operate together and allow a regular user or administrator to access various reports, graphs, or configuration pages.

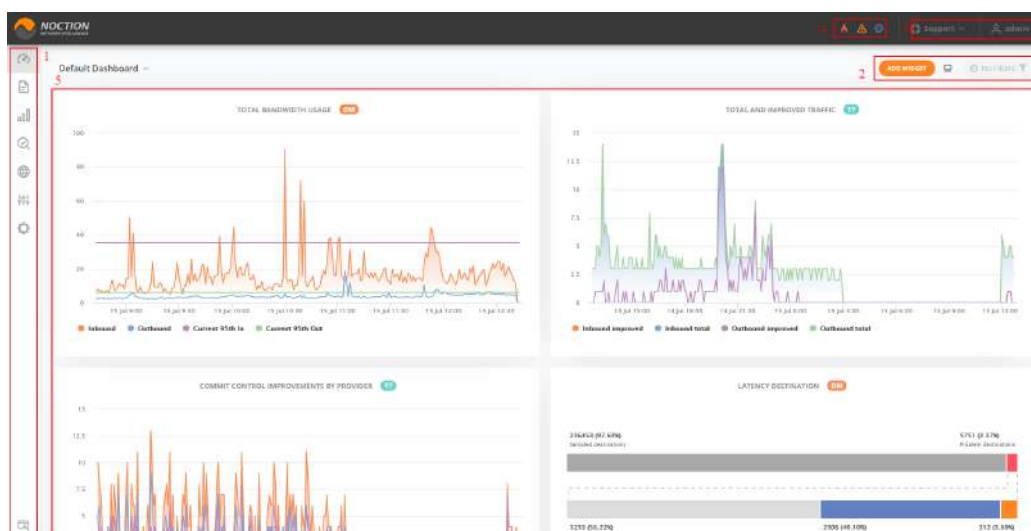


Figure 4.2.1: Frontend Sections

1. The main left sidebar allows quick access to any part of the system via its menu items.

2. A variation of controls appears here in dashboards, reports and graphs depending on the actual view.
Clicking the timestamp/filter icon opens up the menu on the right side (See section 4.2.1). Its contents depend on the current page, and include the time selector as well as specific filters, as explained in the next section.
3. Various links to Support, documentation and the User account administration are located in the top right corner of the Frontend.
4. The status bar allows one to access system events and notifications (See the [Events](#) section).
5. Main dashboard area.
Main dashboard area is editable, widgets can be added, removed or rearranged.

Keyboard shortcuts

- Dashboard - SHIFT+D
- Reports - SHIFT+R
- Graphs - SHIFT+G
- Troubleshooting - SHIFT+T
- Blackholing - SHIFT+B
- Global Commit - SHIFT+M
- Configuration - SHIFT+C
- Management - SHIFT+M
- Search - Ctrl + K
- Instance sidebar - SHIFT+E

4.2.1 Right Sidebar

All dashboards, reports, graphs and events have a sidebar that opens up on click of the timestamp/filter icon at the right side of the screen. Depending on the current page, the sidebar has a different set of controls, allowing users to select a desired timeframe or to filter any report/graph results.

On graphs, reports, events page and dashboards, the time period which the information is being provided for, can be selected by using the “Timestamp” button. An intuitive control will open, and the desired time period can be selected.



Figure 4.2.2: Time period selection

In the case of a report, graph or events page, custom filters can be applied, so that only the specific information is displayed.

The form contains the following fields and options:

- Prefix/ID:
- Prefix size:
- AS number:
- AS name:
- BGPd mode:
- IPv4/IPv6:
- From:
- To:
- Probing source:
- Improvement age:
- RD:
- Improvement type:
- State:
- Page size:
- RESET:
- APPLY:

Figure 4.2.3: Report filters

4.2.2 IRP Instance Performance Stats

Clicking a particular IRP short-name button, opens the right hand side panel with instance-specific information:

- Current amount of inbound and outbound traffic as well as a quick bandwidth reference graph for the last 5 minutes.
- Configured maximum number of improvements (Max IPv6 Improvements will be shown if IPv6 is enabled)
- Number of Routing Policies improvements
- Number of improvements performed today
- The amount of traffic improved today
- Percentage of prefixes improved today
- Number of operating iBGP sessions
- Number of announcements to the edge router
- Announcements to each from configured routers
- Number of BGP sessions with the providers
- Number of operating BGP sessions with the providers

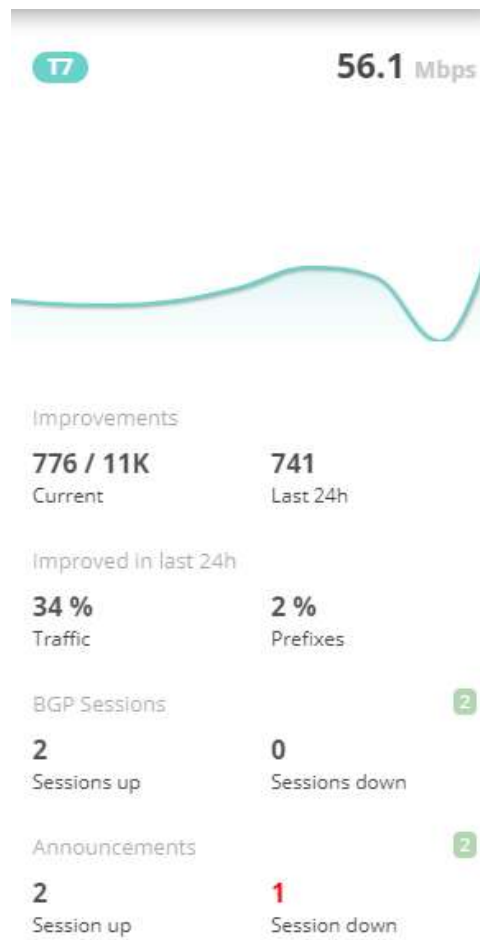


Figure 4.2.4: IRP Instance Performance Stats Quick Reference

See also: [GMI Dashboards](#) .

4.2.3 Global Search

GMI frontend allows to search for historical improvements and the AS Path problems information on prefixes, AS Numbers, AS Names, etc. This function can be accessed by clicking the Search icon in the top bar.




Figure 4.2.5: Global Search

Choose all instances or the IRP instance of interest, enter the needed prefix, AS number or AS name in the search box, then press the SEARCH button.

i When displaying the complete results, specific information for a prefix can be obtained or it can be probed, by clicking the corresponding icons next to each result.

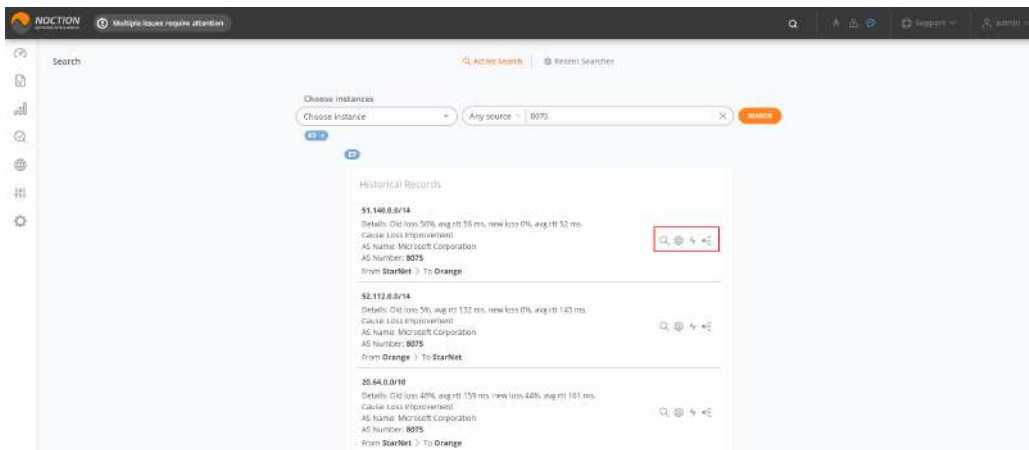


Figure 4.2.6: Global Search results

Visit the Historic Searches tab to see the list of recent search queries.



Figure 4.2.7: Recent Searches

4.2.4 Warning bars

The warning bars are shown at the top of the GMI frontend in case the attention should be drawn to operational/performance issues, some specific IRP instance configuration settings or license/payment issues.




Figure 4.2.8: IRP instance license expiration warning example

4.3 GMI Dashboards

GMI dashboards are the specific sets of interactive visualizations, designed for quick analysis of the IRP instances performance and informational awareness.

Dashboards consist of widgets which contain a reduced version of any given report or a graph. These can be added, edited, deleted or modified as you like. GMI allows users to set up multiple dashboards.

To see a list of existing dashboards, go to “Default Dashboard > ALL”.

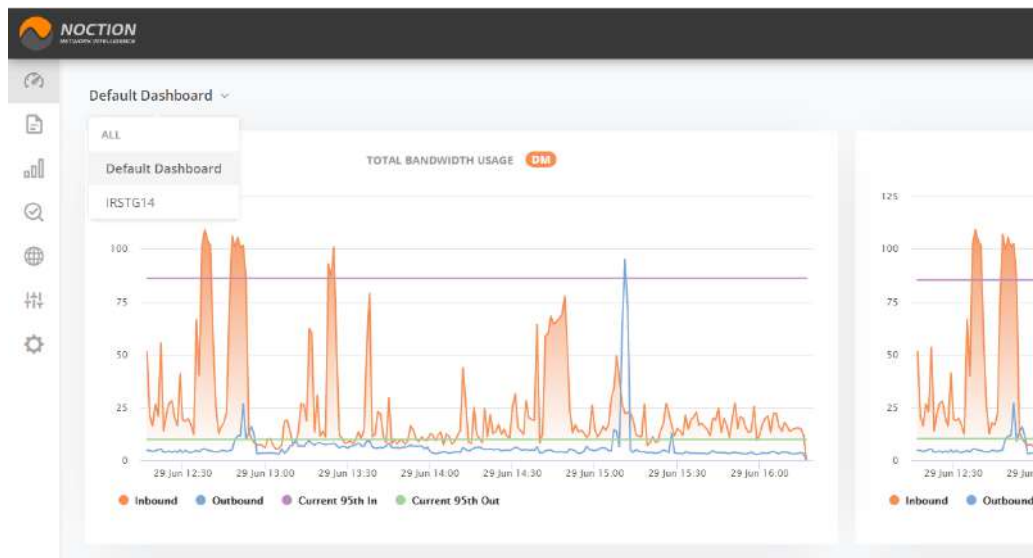


Figure 4.3.1: Dashboard Selection

Dashboards are grouped for easy access into favorite and all. For each dashboard, the directory displays the following information:

Name The name of the dashboard

Description Dashboard user-defined description

Favorite a state marked by a star icon

List of widgets the names of widgets used in the dashboard

Default status the default dashboard the user lands on when logging into GMI

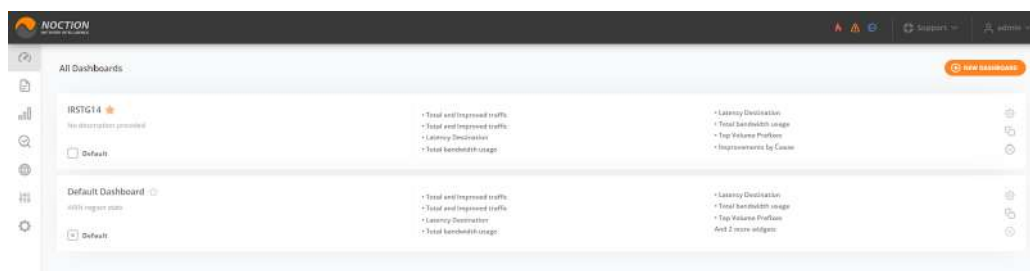


Figure 4.3.2: List of Dashboards

4.3.1 Creating a new dashboard

You can easily create a new dashboard in GMI from the All Dashboards directory.

Click the “NEW DASHBOARD” button at the top right corner of the directory. A pop up will appear. Provide a meaningful name and description for your dashboard.

Mark if you’d prefer it to be a “Favorite” (dashboard will appear at the top of the “All Dashboards” directory) and/or “Default” dashboard.

Press “Create” to continue, or “Close” to return to the directory.

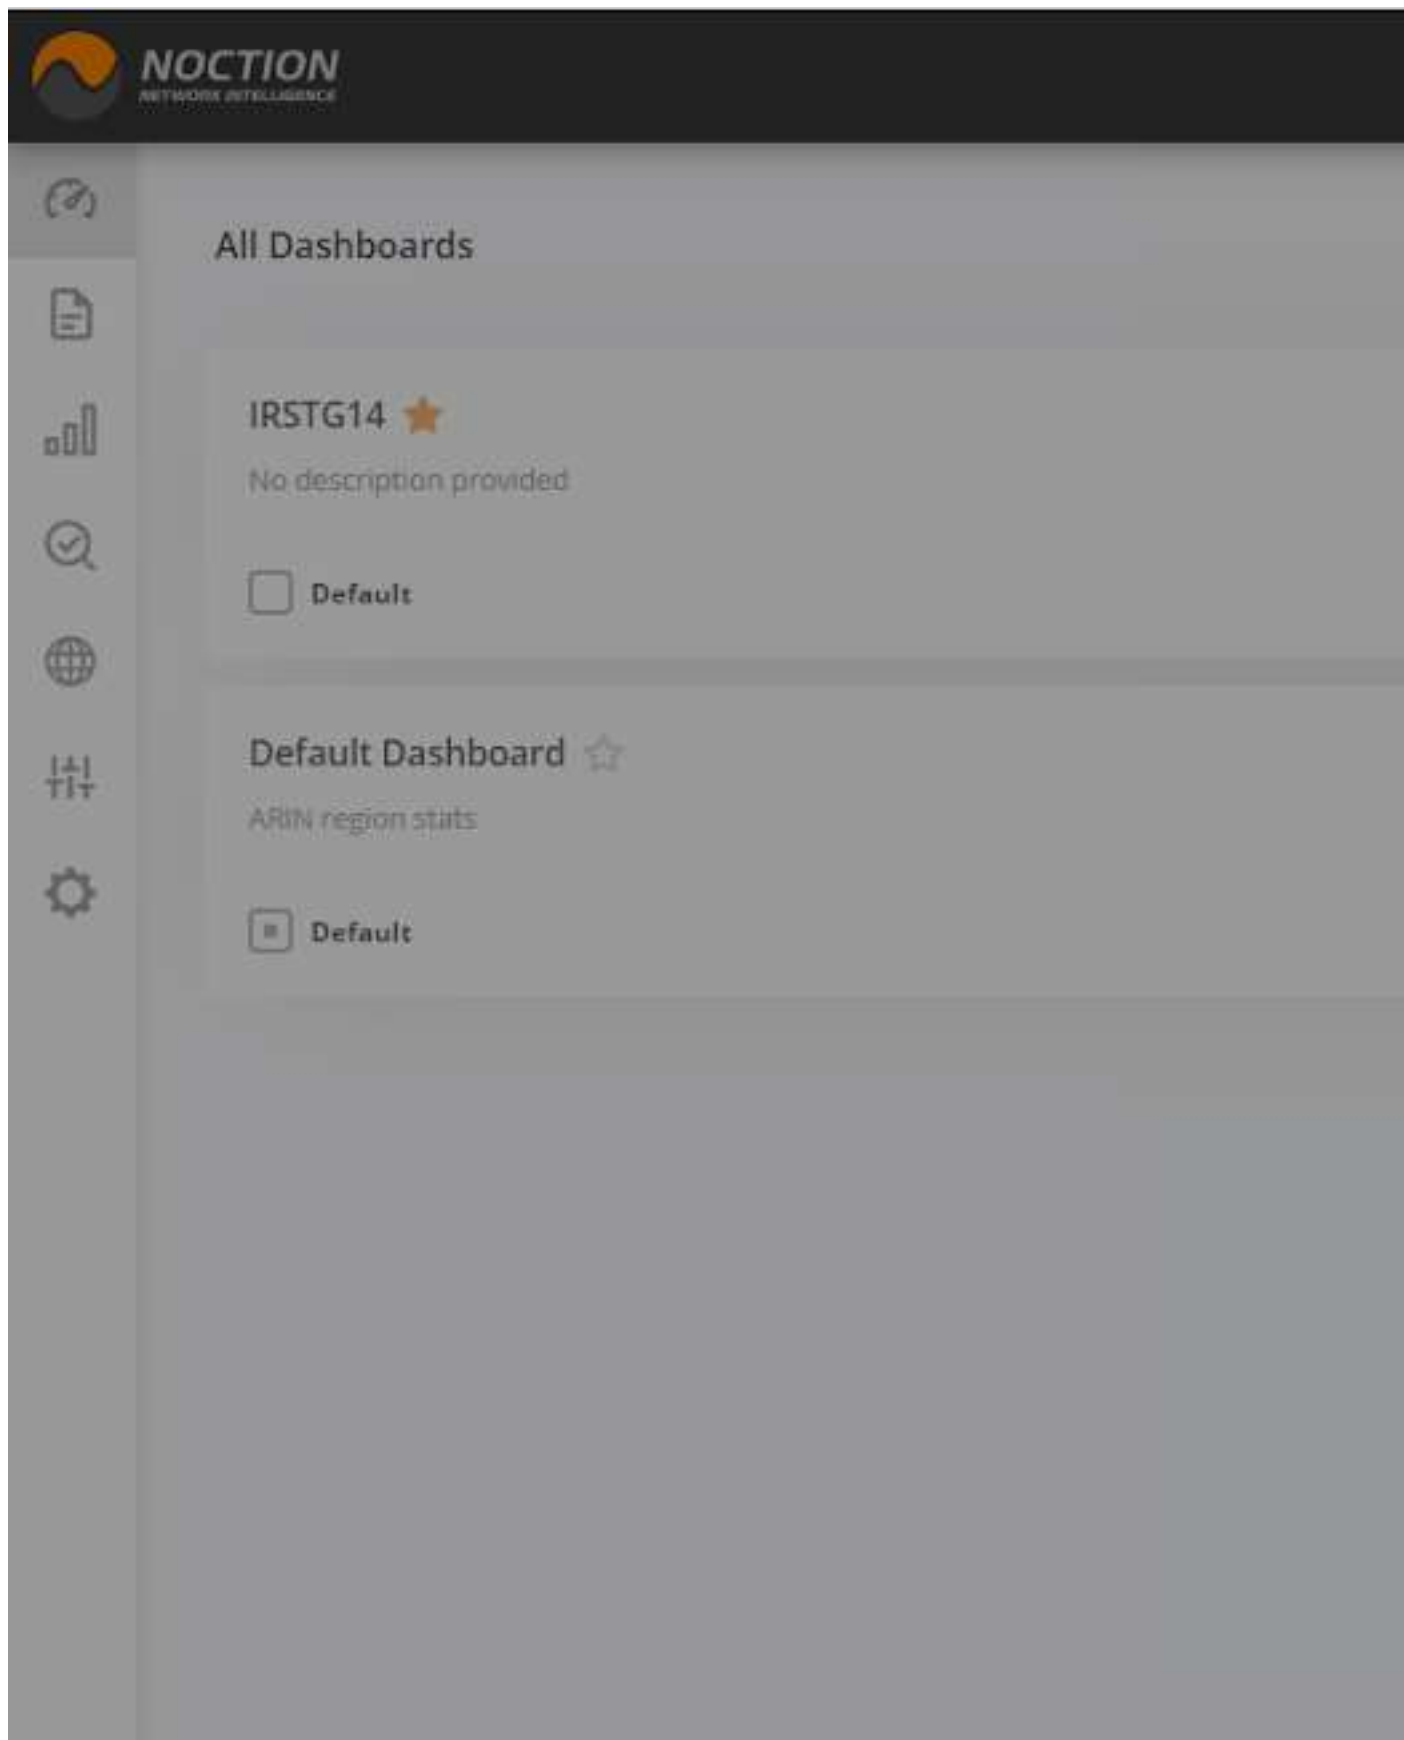


Figure 4.3.3: Creating a new dashboard

Alternatively, you can create a new dashboard by cloning an existing one in the All Dashboards directory. The clone dashboard will be automatically created along with widgets from the original dashboard and added to the directory. Edit the newly created dashboard to change its name and description.

4.3.2 Managing Dashboards

A dashboard can be customized by adding, moving and removing widgets. Hover over any widget to move, edit or delete it. To resize a widget, click the lower right corner and drag it to the desired widget size. Click the display icon in the upper right corner to open a dashboard in full view.

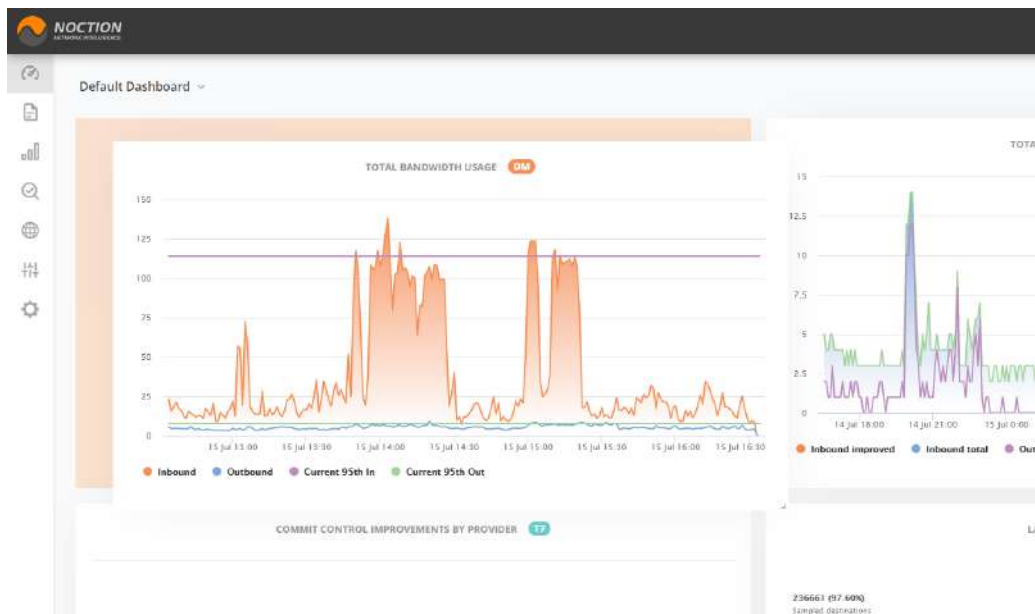


Figure 4.3.4: Rearranging Widgets on Dashboards

Use the "ADD WIDGET" function available on each dashboard to see the library of existing widgets and place the desired ones on a dashboard.

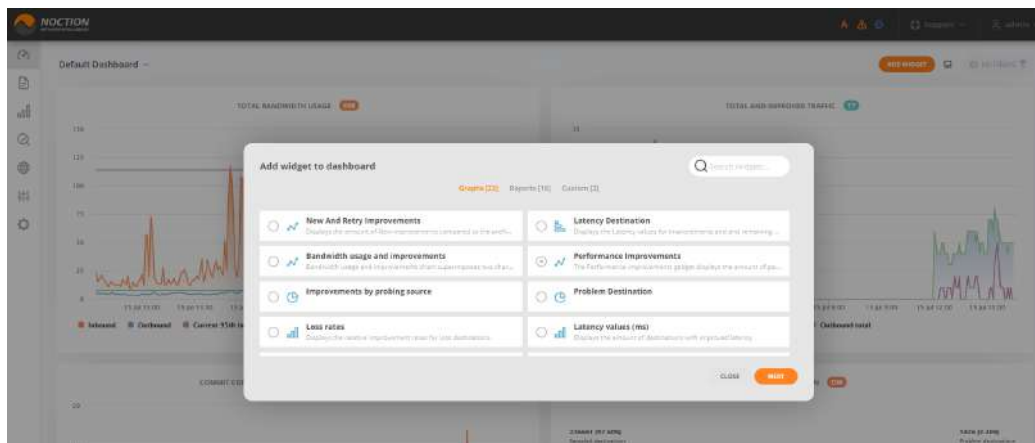


Figure 4.3.5: Widgets Library

You can change the widget's name, refresh interval, timestamp interval as well as the actual IRP instance it is referring to by clicking the Settings icon in the widget's top right corner. Use the rest of the icons to correspondingly refresh, minimize or remove the widget from a dashboard.

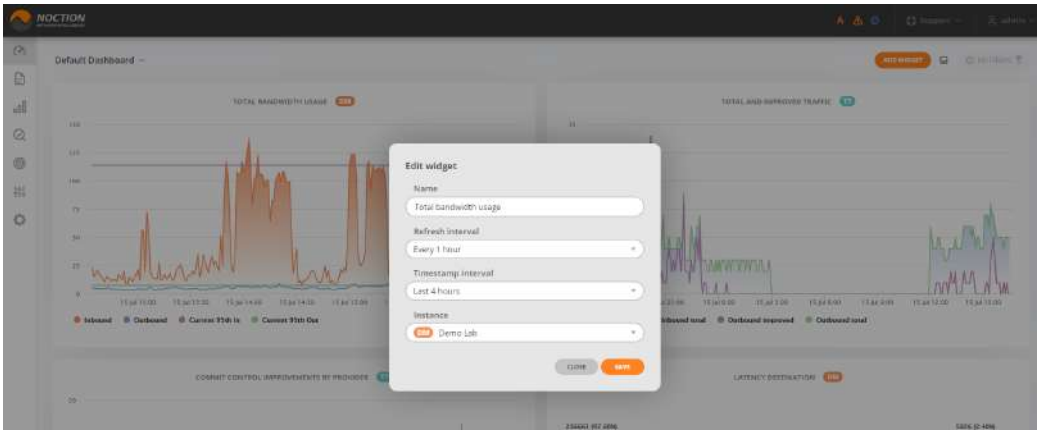


Figure 4.3.6: Editing Widgets

4.3.3 Deleting Dashboards

Click the “Remove” icon on the dashboard you’d like to get rid of in the All Dashboards directory.

Note The default dashboard can not be deleted.




Figure 4.3.7: Deleting Dashboards

4.4 Reports

The GMI Frontend comes with a comprehensive set of reports, reflecting the current state of the network as well as overall statistics on a particular IRP instance performance. Click on the desired report title to open it.

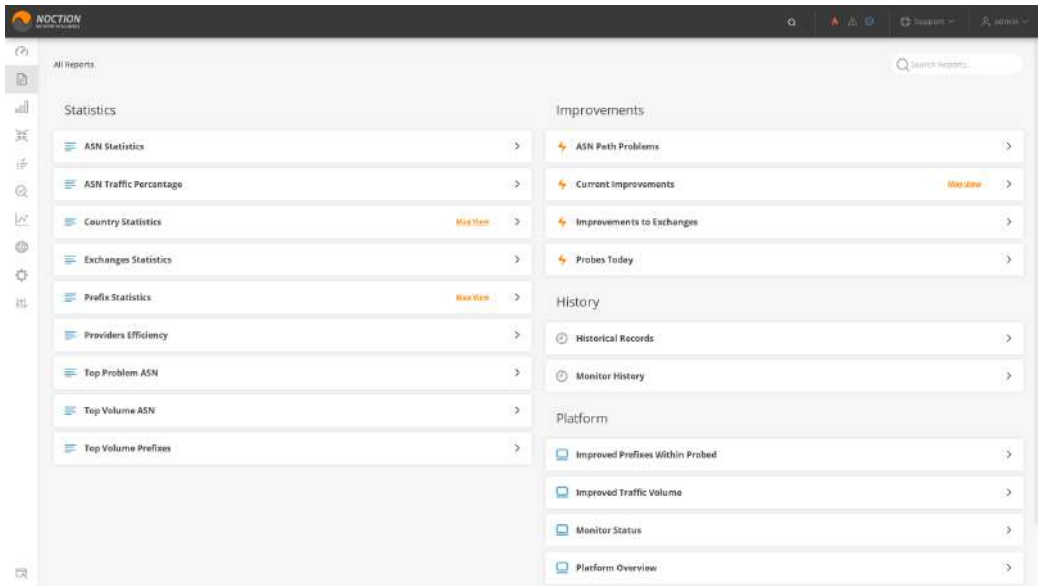


Figure 4.4.1: Reports menu

You can save any report as a widget and add it to any dashboard by clicking the “ADD TO DASHBOARD” button. Select the IRP instance of interest for the report to show relevant data. Filtering options that open up together with a time picker will vary depending on the report type.



Figure 4.4.2: Report options

You can export any report/graph by selecting the desired format (PDF, XLSX, CSV). To print the report, clicking the corresponding print icon.



Figure 4.4.3: Report export and print options

Feel free to subscribe to any report by clicking on the “envelope” icon and filling out fields in the popup form.




Figure 4.4.4: Report email delivery configuration

When there is a need to compare two separate sets of data within a single report or graph (different filters, time periods, or instances), click the “SAVE FOR LATER COMPARE” button.




Figure 4.4.5: Save for Later Compare Button

A saved report/graph view will be placed below the current view, allowing you to subsequently introduce new filtering and/or time frame conditions in the top report and compare the two data sets.




Figure 4.4.6: Save for Later Compare Results

4.4.1 ASN Path Problems

The report is available only when the "Outage detection" is enabled in the IRP Core configuration. The ASN Path Problems report represents the statistics (congestion and outages) related to the specific AS-PATH sequence, which is selected from IRP probing results and categorized as possible outage and/or congestion.

The report highlights the total number of selected, probed, and pending-for-probing prefixes to either confirm or reject the presence of the congestion or outage for an AS-PATH pattern. When the platform confirms the problem's existence, all of the previously probed prefixes with a similar AS-PATH sequence get rerouted away from the affected path.

#	Time	AS-Pattern	Problem	Total	Selected for probing	Pending probing	Affected prefixes		Confirmation rate	Status
							Problem Confirmed	Problem not confirmed		
1	17:24:29	17816 134543	congestion	8	8	8	0	0	0%	🔍
2	17:24:29	17816 17622	congestion	60	60	55	0	5	0%	🔍
3	17:23:51	45796 38794	congestion	9	9	9	0	0	0%	🔍
4	17:23:45	6939 40796	congestion	13	13	8	1	4	0%	🔍
5	17:23:45	8953 4762	congestion	392	117	112	0	5	0%	🔍
6	17:23:45	17816 136959	congestion	6	6	2	1	3	17%	🔍
7	17:23:45	36026 37076	congestion	8	8	8	0	0	0%	🔍
8	17:23:11	35349 16509	congestion	17	12	10	0	2	0%	🔍
9	17:23:10	4837 140726	congestion	27	27	22	2	3	0%	🔍
10	17:22:02	6939 37662	congestion	12	12	8	1	3	9%	🔍
11	17:21:56	6762 4837	congestion	306	106	113	1	22	1%	🔍
12	17:21:55	3356 4837	congestion	112	112	91	6	15	9%	🔍
13	17:21:40	4637 4838	congestion	262	121	101	8	10	7%	🔍
14	17:21:43	3808 56543	congestion	21	11	40	2	13	10%	🔍
15	17:21:31	38827 3608	congestion	24	24	24	8	0	0%	🔍

Figure 4.4.7: ASN Path Problems

4.4.2 ASN Statistics

The “ASN Statistics” report aggregates by Autonomous System the total number of improvements as well as the volume of traffic for a selected time period.

#	Time	ASN	AS Name	Volume	Improvements	95th Inbound (Mbps)	95th Outbound (Mbps)
1	21-08-23	16509	Amazon.com, Inc.	29 GB	87	-	109.1
2	22-08-23	16509	Amazon.com, Inc.	17 GB	25	-	109.1
3	21-08-23	714	Apple Inc.	6.5 GB	26	-	0.0
4	21-08-23	15169	Google LLC	5.3 GB	17	-	0.6
5	21-08-23	8551	Bering International Ltd.	4.0 GB	9574	-	0.3
6	21-08-23	4808	ChinaNicom Hostmaster	3.1 GB	4195	-	0.2
7	21-08-23	4857	CHNGroup Hostmaster	2.8 GB	501	-	0.1
8	21-08-23	55045	APNIC Hostmaster	2.2 GB	3271	-	0.2
9	22-08-23	15169	Google LLC	2.1 GB	3	-	0.6
10	21-08-23	15836	Alcatel-Lucent	2.1 GB	-	-	0.0
11	22-08-23	8551	Bering International Ltd.	1.9 GB	2088	-	0.3
12	21-08-23	50847	APNIC Hostmaster	1.8 GB	1900	-	0.2
13	21-08-23	31252	StartNet S.p.A. SRL	1.7 GB	-	-	0.1
14	21-08-23	23724	ChinaNicom Hostmaster	1.7 GB	1930	-	0.2
15	22-08-23	4808	ChinaNicom Hostmaster	1.6 GB	2241	-	0.2
16	21-08-23	50847	APNIC Hostmaster	1.4 GB	2284	-	0.1

Figure 4.4.8: ASN Statistics

4.4.3 ASN Traffic Percentage

The ASN Traffic Percentage report shows the top sorted Autonomous Systems by the amount of traffic and the calculated percentage of the traffic of these ASes from total traffic.

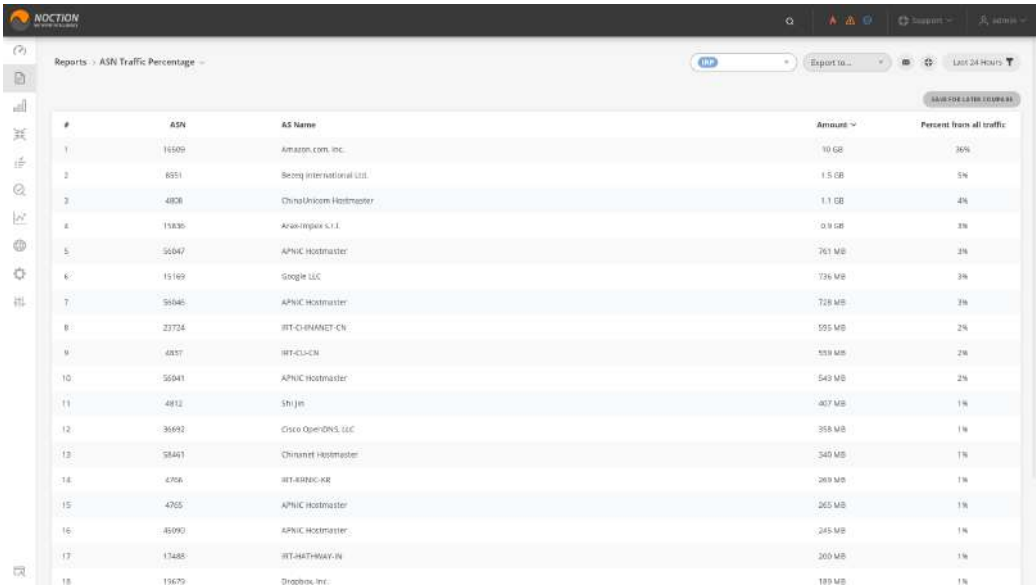


Figure 4.4.9 shows the 'ASN Traffic Percentage' report in the NOCTION interface. The report displays a table with 18 rows, each representing an Autonomous System (ASN). The columns are: #, ASN, AS Name, Amount, and Percent from all traffic. The data is sorted by traffic amount in descending order.

#	ASN	AS Name	Amount	Percent from all traffic
1	16509	Amazon.com, Inc.	10 GB	36%
2	8551	Beemj International Ltd.	1.5 GB	5%
3	4808	ChinaUnicom Hostmaster	1.1 GB	4%
4	15830	Arak-impex s.r.l.	0.9 GB	3%
5	55047	APNIC Hostmaster	761 MB	3%
6	15169	Google LLC	736 MB	3%
7	54046	APNIC Hostmaster	728 MB	3%
8	23734	BIT-CHINANET-CN	555 MB	2%
9	4837	BIT-CLIX-CN	539 MB	2%
10	55041	APNIC Hostmaster	543 MB	2%
11	4812	Shijie	407 MB	1%
12	36692	Cisco OpenDNS, LLC	358 MB	1%
13	58441	Chinunet Hostmaster	340 MB	1%
14	4756	BIT-8080-KR	289 MB	1%
15	4705	APNIC Hostmaster	265 MB	1%
16	45090	APNIC Hostmaster	245 MB	1%
17	17488	BIT-44THWAY-IN	200 MB	1%
18	19679	Dropbox Inc.	189 MB	1%

Figure 4.4.9: ASN Traffic Percentage

4.4.4 Country Statistics

The “Country Statistics” report shows the distribution of traffic volume and the number of improvements by country. It helps see to what regions most of a network’s traffic is addressed and also the relative number of suboptimal routes towards them that a selected IRP instance identified and addressed by injecting improvements.

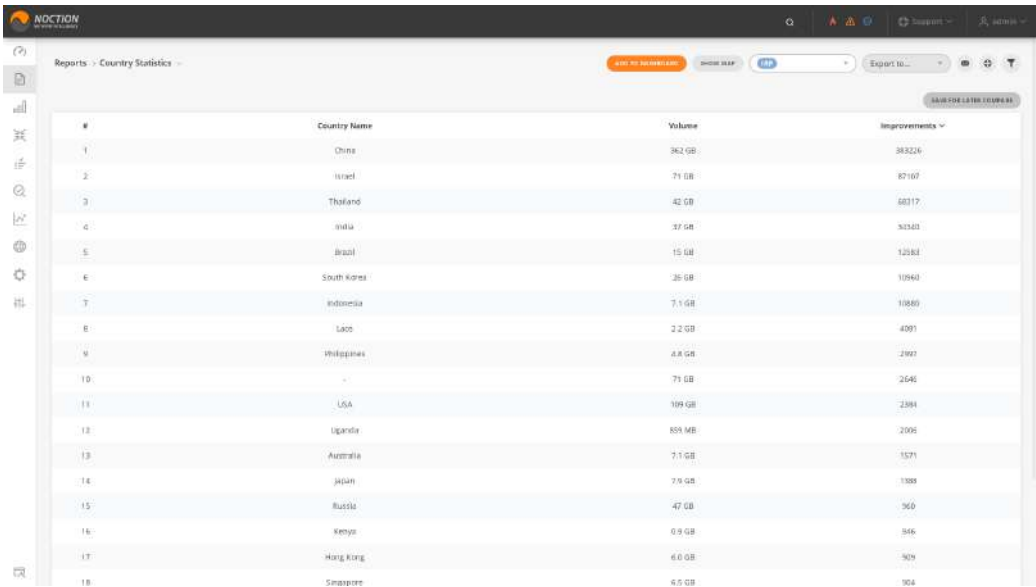


Figure 4.4.10 shows the 'Country Statistics' report in the NOCTION interface. The report displays a table with 18 rows, each representing a country. The columns are: #, Country Name, Volume, and Improvements. The data is sorted by volume in descending order.

#	Country Name	Volume	Improvements
1	China	962 GB	343226
2	Israel	71 GB	87107
3	Thailand	42 GB	68317
4	India	37 GB	50343
5	Brazil	15 GB	12084
6	South Korea	26 GB	10940
7	Indonesia	7.1 GB	10860
8	Laos	2.2 GB	4091
9	Philippines	8.8 GB	2991
10	-	71 GB	2646
11	USA	109 GB	2384
12	Uganda	859 MB	2095
13	Australia	7.1 GB	1571
14	Japan	7.6 GB	1389
15	Russia	47 GB	360
16	Kenya	0.9 GB	346
17	Hong Kong	6.0 GB	309
18	Singapore	6.5 GB	304

Figure 4.4.10: Country Statistics

Country statistics report offers a map view, highlighting countries with most/least number of improvements.




Figure 4.4.11: Country Statistics map view

4.4.5 Current Improvements

The “Current Improvements” report provides a list of the currently active improvements as per specific IRP instance. The report shows the improved prefix and the providers from and to which the traffic was redirected. It also provides the reason for the improvement, along with the performance values before and after the traffic was rerouted.

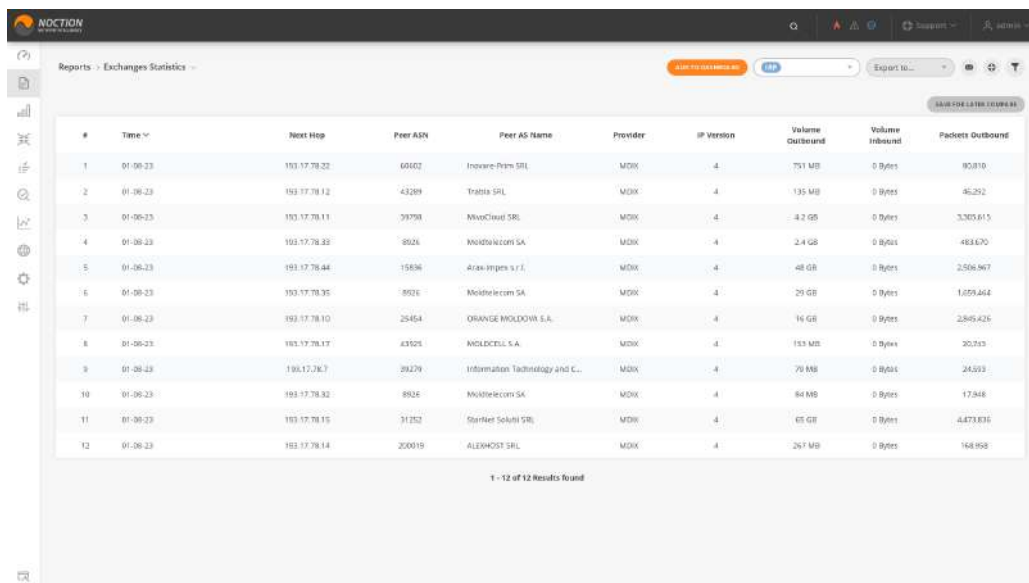
#	Time	Prefix	ASN	AS Name	Loss	Latency	From	To	Type	Rd
1	16:39:09	78.177.60.0/24	8551	Beez International Ltd.	49 → 0	96 → 90	ORANGE	StarnetMD	ORANGE	-
2	16:39:09	115.231.28.0/23	58461	Chinanet Hostmaster	29 → 19	272 → 273	ORANGE	StarnetMD	ORANGE	-
3	16:39:02	183.134.58.0/23	58461	Chinanet Hostmaster	30 → 43	276 → 304	StarnetMD	ORANGE	ORANGE	-
4	16:38:56	117.27.234.0/24	133774	Chinanet Hostmaster	37 → 96	307 → 272	StarnetMD	ORANGE	ORANGE	-
5	16:38:56	49.230.212.0/22	131445	APNIC Hostmaster	0 → 0	412 → 353	StarnetMD	ORANGE	ORANGE	-
6	16:38:56	283.135.112.0/22	58461	Chinanet Hostmaster	40 → 81	292 → 278	ORANGE	StarnetMD	ORANGE	-
7	16:38:56	49.230.208.0/22	131445	APNIC Hostmaster	0 → 0	376 → 272	StarnetMD	ORANGE	ORANGE	-
8	16:38:56	70.177.42.0/24	8551	Beez International Ltd.	42 → 0	88 → 89	ORANGE	StarnetMD	ORANGE	-
9	16:38:53	212.58.224.0/21	2818	RBC	0 → 0	9 → 48	ORANGE	StarnetMD	ORANGE	-
10	16:38:50	183.131.216.0/21	58461	Chinanet Hostmaster	24 → 20	295 → 257	ORANGE	StarnetMD	ORANGE	-
11	16:38:50	122.224.192.0/24	58461	Chinanet Hostmaster	26 → 0	290 → 267	ORANGE	StarnetMD	ORANGE	-
12	16:38:50	122.225.192.0/22	58461	Chinanet Hostmaster	22 → 0	276 → 279	StarnetMD	ORANGE	ORANGE	-
13	16:38:50	184.302.0/22	7470	NET-TELENET-TELENET	68 → 30	179 → 180	ORANGE	StarnetMD	ORANGE	-
14	16:38:50	122.224.194.0/22	58461	Chinanet Hostmaster	25 → 0	272 → 276	ORANGE	StarnetMD	ORANGE	-
15	16:38:42	101.42.224.0/19	49260	APNIC Hostmaster	19 → 9	271 → 263	StarnetMD	ORANGE	ORANGE	-
16	16:38:37	208.194.128.0/21	9729	APNIC Hostmaster	94 → 46	273 → 284	ORANGE	StarnetMD	ORANGE	-
17	16:38:37	27.158.122.0/23	133774	Chinanet Hostmaster	52 → 0	301 → 298	ORANGE	StarnetMD	ORANGE	-

Figure 4.4.12: Current Improvements

i The “ASN” field can be empty in case there is no information about the AS number (a prefix belongs to) in the ASN dictionary. As soon as the ASN information is available in the dictionary, the field will be filled accordingly.

4.4.6 Exchanges Statistics

The “Exchanges statistics” report lists specific information regarding communication with peers on Internet Exchanges. Filters help identify the required details.



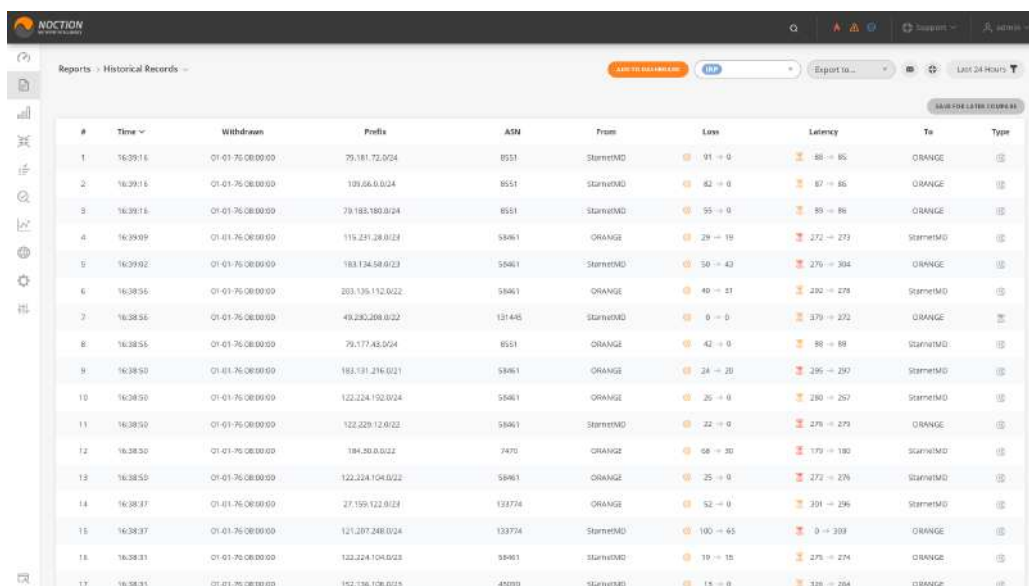
The screenshot shows the 'Exchanges Statistics' report in the NOXION interface. The table lists 12 results, showing details for each exchange hop. The columns include: #, Time, Next Hop, Peer ASN, Peer AS Name, Provider, IP Version, Volume Outbound, Volume Inbound, and Packets Outbound. The data shows various providers like MOK, NIXOCLOUD SRL, and NIXOCLOUD S.A. with different IP versions and volumes.

#	Time	Next Hop	Peer ASN	Peer AS Name	Provider	IP Version	Volume Outbound	Volume Inbound	Packets Outbound
1	01-06-23	193.17.78.22	60602	Inovare-Prem SRL	MOK	4	751 MB	0 Bytes	80,818
2	01-06-23	193.17.78.12	43289	Trabita SRL	MOK	4	135 MB	0 Bytes	45,252
3	01-06-23	193.17.78.11	39798	Nixocloud SRL	MOK	4	4.2 GB	0 Bytes	3,305,815
4	01-06-23	193.17.78.23	8926	Nixocloud SA	MOK	4	2.4 GB	0 Bytes	483,670
5	01-06-23	193.17.78.44	15836	Alcan-IPNet s.r.l.	MOK	4	48 GB	0 Bytes	2,596,967
6	01-06-23	193.17.78.35	8926	Nixocloud SA	MOK	4	29 GB	0 Bytes	1,059,462
7	01-06-23	193.17.78.10	25454	ORANGE HOLDING S.A.	MOK	4	16 GB	0 Bytes	2,845,425
8	01-06-23	193.17.78.17	43525	NIXOCLOUD S.A.	MOK	4	153 MB	0 Bytes	30,743
9	01-06-23	193.17.78.7	29279	Information Technology and C...	MOK	4	79 MB	0 Bytes	24,593
10	01-06-23	193.17.78.32	8926	Nixocloud SA	MOK	4	84 MB	0 Bytes	17,848
11	01-06-23	193.17.78.15	31252	StarNet Soluții SRL	MOK	4	65 GB	0 Bytes	4,473,835
12	01-06-23	193.17.78.14	200019	ALEXHOST SRL	MOK	4	267 MB	0 Bytes	168,958

Figure 4.4.13: Exchanges Statistics

4.4.7 Historical Records

The “Historical Records” report contains a complete list of active and inactive improvements that were made by a particular IRP instance.



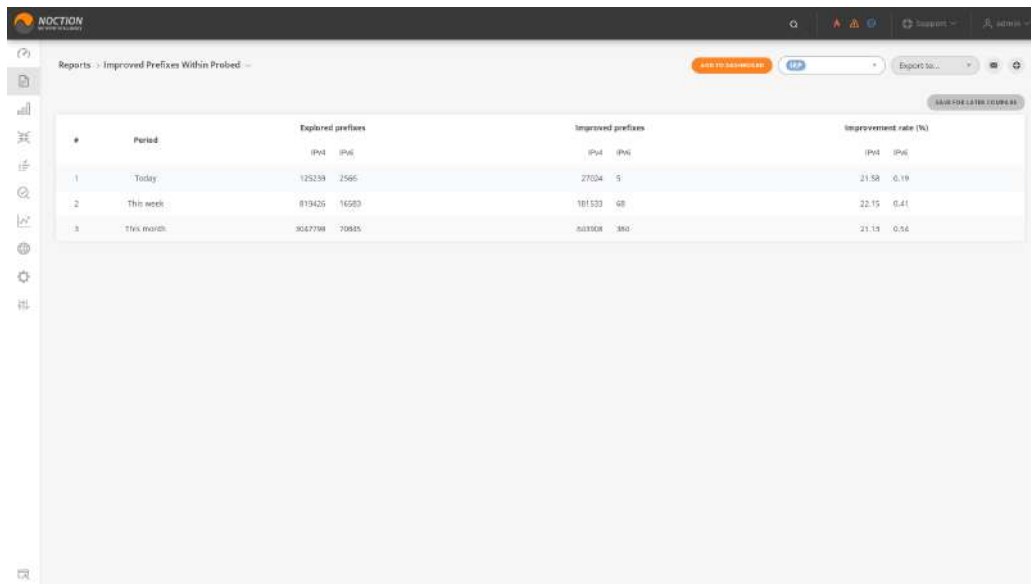
The screenshot shows the 'Historical Records' report in the NOXION interface. The table lists 17 records, showing details for each historical improvement. The columns include: #, Time, Withdrawn, Prefix, ASN, From, Loss, Latency, To, and Type. The data shows various prefixes and ASNs with associated loss and latency values.

#	Time	Withdrawn	Prefix	ASN	From	Loss	Latency	To	Type
1	16:39:15	01-01-76:08:00:00	75.181.72.0/24	8551	StarNetMD	91 → 0	88 → 85	ORANGE	
2	16:39:15	01-01-76:08:00:00	109.66.9.0/24	8551	StarNetMD	82 → 0	87 → 85	ORANGE	
3	16:39:15	01-01-76:08:00:00	79.183.180.0/24	8551	StarNetMD	55 → 0	89 → 86	ORANGE	
4	16:39:09	01-01-76:08:00:00	115.231.28.0/24	58461	ORANGE	29 → 19	272 → 273	StarNetMD	
5	16:39:02	01-01-76:08:00:00	183.134.58.0/23	58461	StarNetMD	50 → 43	276 → 304	ORANGE	
6	16:38:56	01-01-76:08:00:00	203.136.152.0/22	58461	ORANGE	40 → 21	230 → 278	StarNetMD	
7	16:38:56	01-01-76:08:00:00	49.230.208.0/22	131445	StarNetMD	0 → 0	379 → 272	ORANGE	
8	16:38:56	01-01-76:08:00:00	75.177.45.0/24	8551	ORANGE	42 → 0	88 → 88	StarNetMD	
9	16:38:50	01-01-76:08:00:00	183.131.216.0/21	58461	ORANGE	24 → 20	295 → 297	StarNetMD	
10	16:38:50	01-01-76:08:00:00	122.224.152.0/24	58461	ORANGE	25 → 0	280 → 257	StarNetMD	
11	16:38:50	01-01-76:08:00:00	122.229.12.0/22	58461	StarNetMD	22 → 0	276 → 273	ORANGE	
12	16:38:50	01-01-76:08:00:00	184.50.0.0/22	7470	ORANGE	58 → 30	179 → 180	StarNetMD	
13	16:38:50	01-01-76:08:00:00	122.224.104.0/22	58461	ORANGE	25 → 0	272 → 276	StarNetMD	
14	16:38:37	01-01-76:08:00:00	27.159.122.0/24	133774	ORANGE	52 → 0	304 → 296	StarNetMD	
15	16:38:37	01-01-76:08:00:00	121.207.248.0/24	133774	StarNetMD	100 → 45	0 → 309	ORANGE	
16	16:38:31	01-01-76:08:00:00	122.224.104.0/23	58461	StarNetMD	19 → 15	275 → 274	ORANGE	
17	16:38:31	01-01-76:08:00:00	152.136.106.0/23	40009	StarNetMD	13 → 0	328 → 254	ORANGE	

Figure 4.4.14: Historical Records

4.4.8 Improved Prefixes Within Probed

This report offers overall statistics on the number of explored prefixes and the rate of the improved ones per day, week and month.



#	Period	Explored prefixes		Improved prefixes		Improvement rate (%)	
		IPv4	IPv6	IPv4	IPv6	IPv4	IPv6
1	Today	125239	2565	27034	5	21.58	0.19
2	This week	819425	16263	181533	68	22.15	0.41
3	This month	3047798	70925	603908	385	21.13	0.54

Figure 4.4.15: Improved Prefixes Within Probed

4.4.9 Improved Traffic Volume

This report offers overall statistics on the amount of total traffic and improved traffic on a daily, weekly and monthly basis.

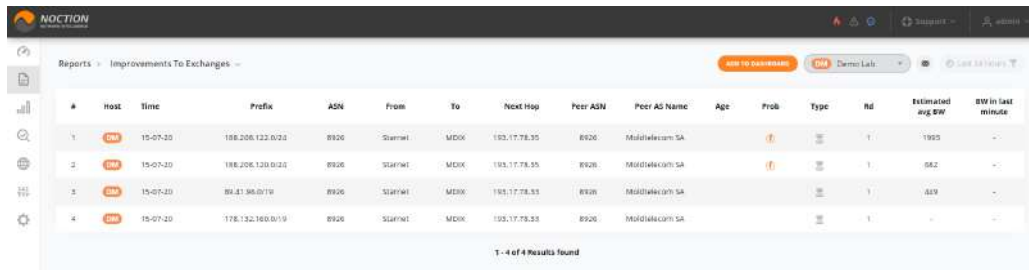


#	Period	Total traffic		Improved traffic		Improved (%)	
		IPv4	IPv6	IPv4	IPv6	IPv4	IPv6
1	Today	45 GB	1.7 GB	19 GB	215 MB	40.38	12.08
2	This week	140 GB	6.2 GB	45 GB	352 MB	30.62	5.58
3	This month	1.7 TB	143 GB	509 GB	14 GB	29.53	9.87

Figure 4.4.16: Improved Traffic Volume

4.4.10 Improvements to Exchanges

The “Improvements to Exchanges” report contains a complete list of current improvements where the new provider is an Internet Exchange. The report includes data about improved prefixes and their relative traffic by means of last-minute and average bandwidth usage. The average bandwidth is estimated based on current hour statistics.



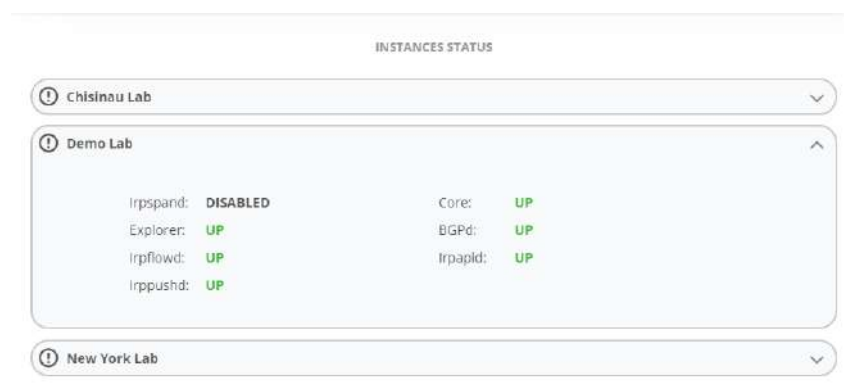
#	host	time	Prefix	ASN	From	To	Next Hop	Peer ASN	Peer AS Name	Age	Prob	Type	ID	Estimated avg BW	BW in last minute
1	198.208.122.0/24	15-07-20	198.208.122.0/24	8926	StarNet	MDIX	193.17.78.35	8926	MultiMediacom SA				1	1995	-
2	198.208.120.0/24	15-07-20	198.208.120.0/24	8926	StarNet	MDIX	193.17.78.35	8926	MultiMediacom SA				1	842	-
3	89.81.88.0/19	15-07-20	89.81.88.0/19	8926	StarNet	MDIX	193.17.78.35	8926	MultiMediacom SA				1	459	-
4	178.152.160.0/19	15-07-20	178.152.160.0/19	8926	StarNet	MDIX	193.17.78.35	8926	MultiMediacom SA				1	-	-

T - 4 of 4 results found

Figure 4.4.17: Improvements to Exchanges

4.4.11 Instances Status (available as a widget only)

The “Instances Status” offers a quick view of the state of IRP instances and their corresponding components.

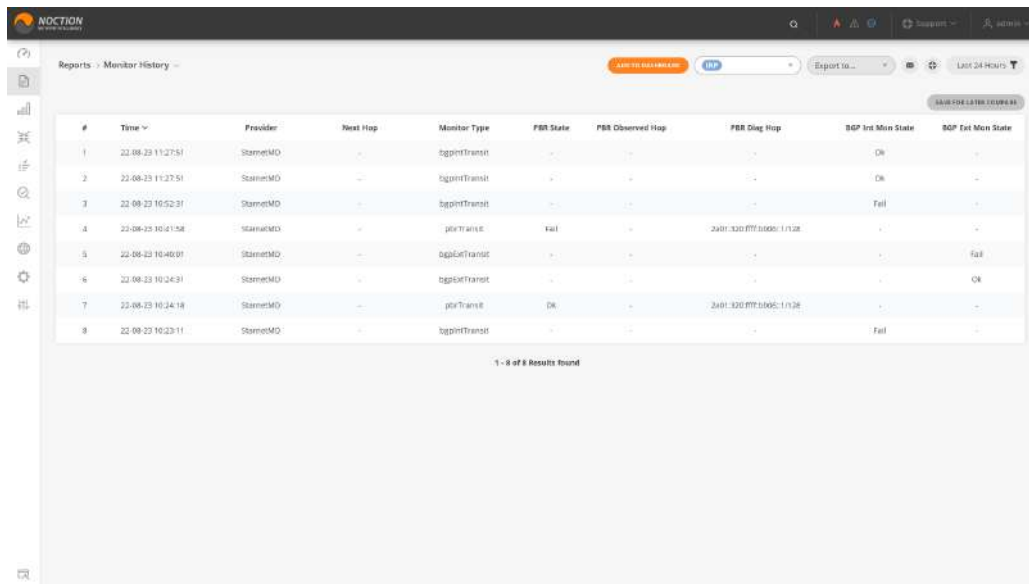


INSTANCES STATUS	
Chisinau Lab	
Demo Lab	
Irpspand:	DISABLED
Explorer:	UP
Irpflowd:	UP
Irppushd:	UP
Core:	UP
BGPd:	UP
Irpapid:	UP
New York Lab	

Figure 4.4.18: Instances Status

4.4.12 Monitor History

The “Monitor History” report highlights the time, monitor details, and the monitor state per provider, facilitating the tracing of various issues on the network and in any particular IRP instance configuration.



The screenshot shows the 'Monitor History' page in the NOCTION interface. It features a table with 10 columns: #, Time, Provider, Next Hop, Monitor Type, PBR State, PBR Observed Hop, PBR Diag Hop, BGP Int Mon State, and BGP Ext Mon State. The table contains 8 rows of data. The first three rows show 'bgpIntTransit' monitors with 'OK' states. The next two rows show 'bgpIntTransit' monitors with 'Fail' states. The last three rows show 'pbrTransit' monitors with 'OK' and 'Fail' states. A '1 - 8 of 8 Results Found' message is displayed at the bottom of the table.

#	Time	Provider	Next Hop	Monitor Type	PBR State	PBR Observed Hop	PBR Diag Hop	BGP Int Mon State	BGP Ext Mon State
1	22-08-23 11:27:51	StarnetMO	-	bgpIntTransit	OK	-	-	OK	-
2	22-08-23 11:27:51	StarnetMO	-	bgpIntTransit	OK	-	-	OK	-
3	22-08-23 10:52:31	StarnetMO	-	bgpIntTransit	OK	-	-	Fail	-
4	22-08-23 10:41:58	StarnetMO	-	pbrTransit	Fail	-	2401:320:FFF:5005:1/1/32	-	-
5	22-08-23 10:40:01	StarnetMO	-	bgpIntTransit	OK	-	-	-	Fail
6	22-08-23 10:24:31	StarnetMO	-	bgpIntTransit	OK	-	-	-	OK
7	22-08-23 10:24:18	StarnetMO	-	pbrTransit	OK	-	2401:320:FFF:5005:1/1/32	-	-
8	22-08-23 10:23:11	StarnetMO	-	bgpIntTransit	OK	-	-	Fail	-

Figure 4.4.19: Monitor History

4.4.13 Monitor Status

The “Monitor Status” will report the status of the IRP BGP Internal and External Monitors for each of the configured providers. BGP Monitors monitor the provider BGP session using SNMP protocol and the ICMP/UDP control packets.

When one of the monitors is reported as failed, IRP will withdraw the improvements made towards this peer.

PBR states as a Policy Based Routing, is the mechanism used to ensure that IRP probing packets follow the preconfigured peer and therefore are correct.



The screenshot shows the 'Monitor Status' page in the NOCTION interface. It features a table with 7 columns: #, Provider, Provider Status, PBR State, PBR Diag Hop, BGP Internal Monitor, and BGP External Monitor. The table contains 3 rows of data. The first two rows show 'Provider 1' and 'Provider 2' with 'OK' status. The third row shows 'Internal Exchange' with 'Unknown' status.

#	Provider	Provider Status	PBR State	PBR Diag Hop	BGP Internal Monitor	BGP External Monitor
1	Provider 1	OK	OK	2401:320:FFF:5005:1/1/32	OK	OK
2	Provider 2	OK	OK	2401:320:FFF:5005:1/1/32	OK	OK
3	Internal Exchange	Unknown	Unknown	16.17.38.204:2011:8888:8888	Unknown	Unknown

Figure 4.4.20: Monitor Status

4.4.14 Platform Overview

The “Platform Overview” provides information on the rate of improvements based upon latency, loss and cost for a particular IRP instance. This report can be used for a quick overview of an IRP instance performance and improvements. The report displays statistics for IPv4/IPv6 separately, while widget added to a dashboard shows aggregated data.

One can see here that almost all the improved IPv4 routes had a loss drop of 20% or more. For 61% of the improved routes the loss was fully eliminated. There are also 16% of the improved routes, which were redirected from a complete blackout. The report also provides the amount of the accomplished route improvements made, based upon latency, cost and commit level.



Figure 4.4.21: Platform Overview

4.4.15 Prefix Improvement Rate by Cause

This report offers overall statistics on improvements rate, based upon performance reasons, commit control, and cost reasons on a daily, weekly, and monthly basis.




Figure 4.4.22: Prefix Improvement Rate by Cause

4.4.16 Prefix Statistics

This report lists specific prefixes with relevant monthly values such as traffic volume and number of improvements or current unique and top hosts values for a selected IRP instance.

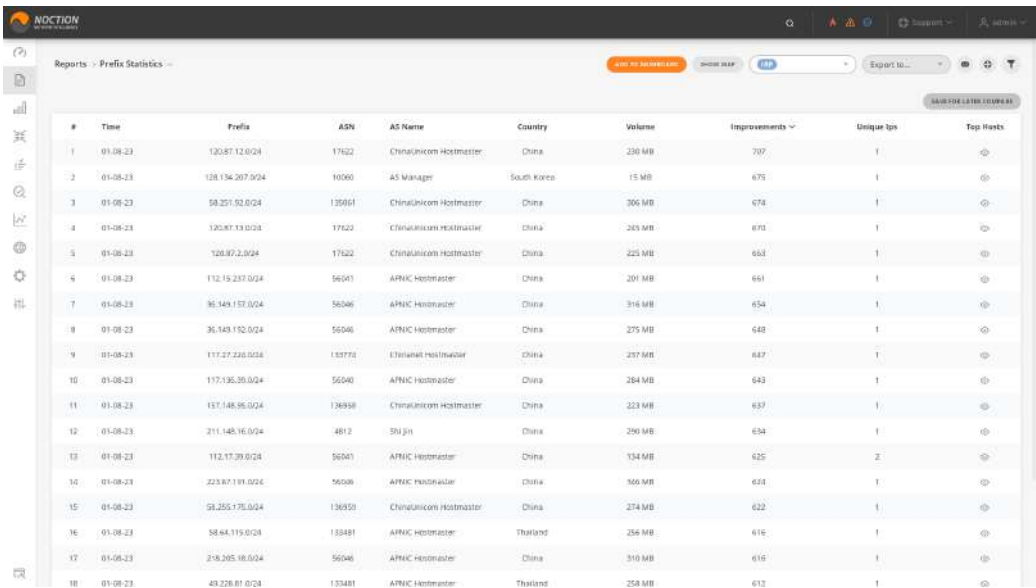


Figure 4.4.23: Prefix Statistics

Note that Prefix Statistics report offers a map view with the regions navigation options, highlighting the parts of the world where either most traffic or the highest number of improvements were made.



Figure 4.4.24: Prefix Statistics map view

4.4.17 Probes Today

The “Probes Today” report provides probing details, including probes that did not warrant an improvement or failed completely. The report shows probed prefixes and the selected probing IPs in that prefix with details about probe state and actual measurements across all providers. Missing details for some providers, for example for partial providers or Internet exchanges, indicate that either the provider was stopped, the prefix is not serviced by the provider, or a probing IP was not identified and the probe has failed completely.

Icons and coloring are used in the report to highlight data for example if an improvement exists for the prefix, or whether there’s a 100% loss. Filtering can be used to quickly analyze the provided data.

See also: [Probes Today](#).

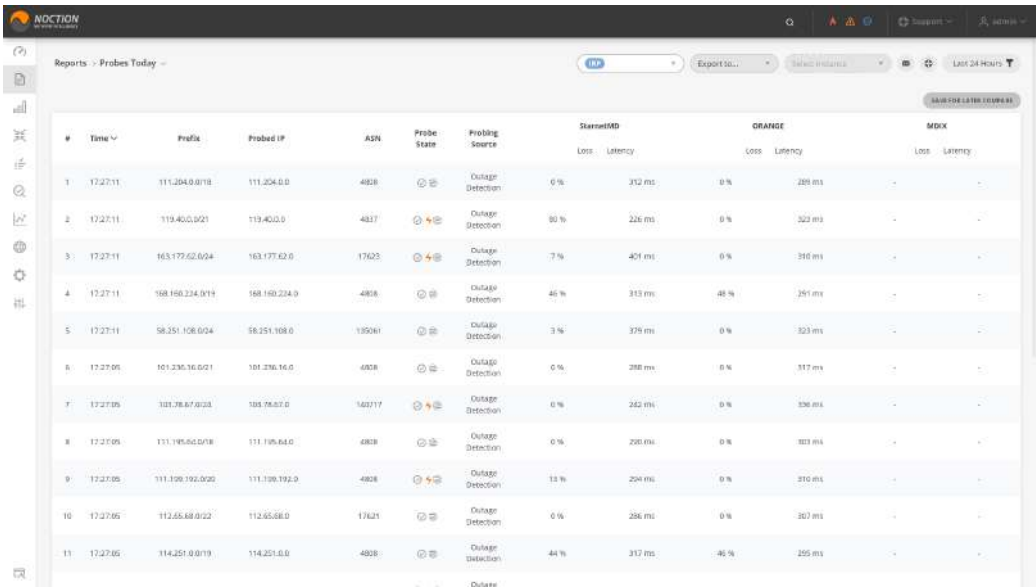


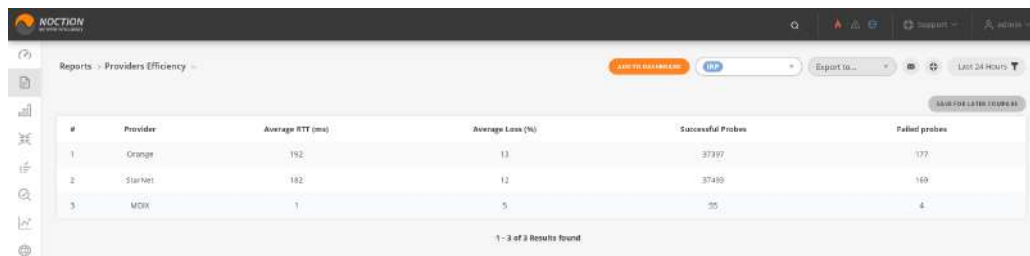
Figure 4.4.25: Probes Today

4.4.18 Providers Efficiency

When operating in any given network, IRP makes numerous probes to determine the performance characteristics of alternative routes.

Provider performance report presents an aggregated view of these measurements. The data on the report highlights the average Packet Loss and average Latency as well as the number of successful and failed probes performed by the platform per each provider.

See also: [Providers Efficiency](#)



#	Provider	Average RTT (ms)	Average Loss (%)	Successful Probes	Failed probes
1	Orange	192	13	37397	177
2	StarNet	182	12	37489	169
3	MOK	1	5	35	4

1 - 3 of 3 Results found

Figure 4.4.26: Providers Efficiency

4.4.19 Providers Performance and Probes (last 24 hours)

During the day IRP instances make many probes to determine performance characteristics of alternative routes. The provider performance report presents an aggregated view of these measurements for providers. The data on the report highlights Packet loss and Average latency per provider as well as the number of unique probed prefixes.

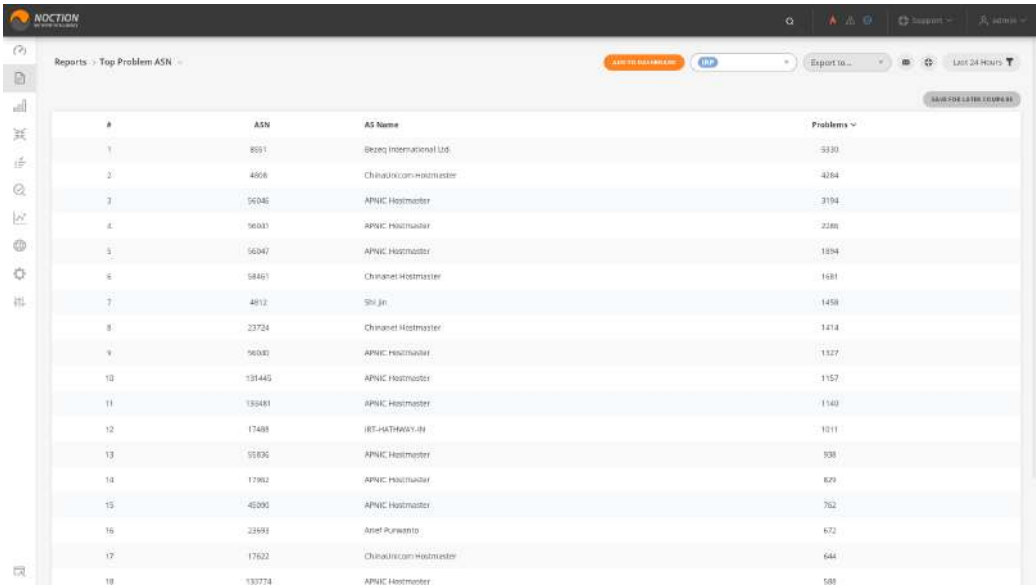


#	Provider	Average Performance		Unique probed prefixes
		Loss (%)	Latency (ms)	
1	ORANGE	39.15	223	3565
2	MOK	35.29	3	17
3	StarNetMD	35.60	235	3577

Figure 4.4.27: Providers Performance and Probes (last 24 hours)

4.4.20 Top Problem ASN

The “Top Problem ASN” report reveals the most problematic Autonomous Systems that the monitored networks are sending traffic to. It does that by showing how many times the traffic going to that AS was rerouted by a particular IRP instance.

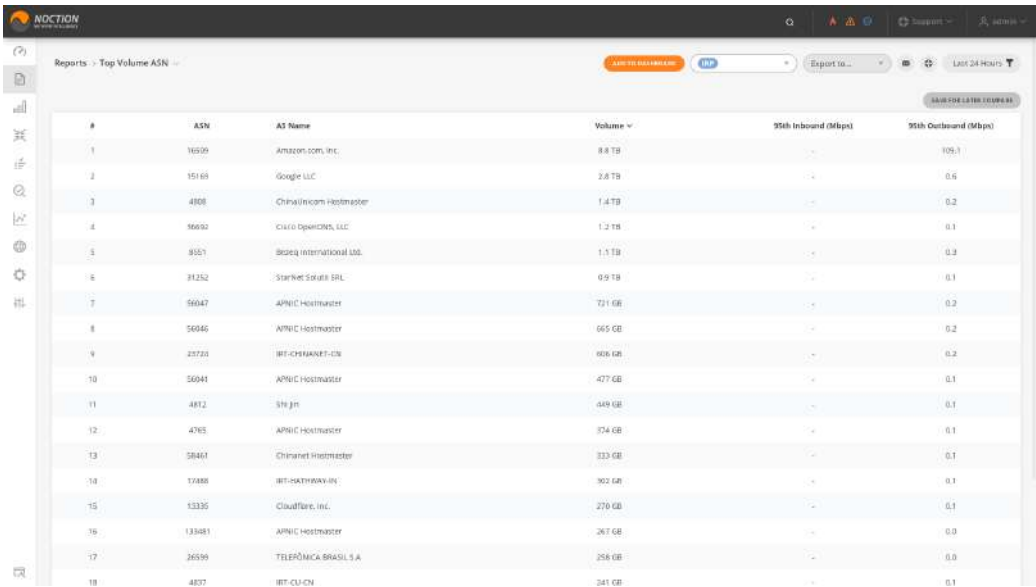


#	ASN	AS Name	Problems
1	8531	Bezeq International Ltd.	5330
2	4806	Chinaunicom Hostmaster	4264
3	56346	APNIC Hostmaster	3194
4	56343	APNIC Hostmaster	2386
5	56347	APNIC Hostmaster	1894
6	58461	Chinanet Hostmaster	1681
7	4812	SH JIN	1458
8	23724	Chinanet Hostmaster	1414
9	56340	APNIC Hostmaster	1327
10	131445	APNIC Hostmaster	1157
11	133481	APNIC Hostmaster	1140
12	17488	IBT-HATHAWAY-IR	1011
13	55836	APNIC Hostmaster	938
14	17862	APNIC Hostmaster	829
15	45090	APNIC Hostmaster	762
16	23593	Anet Punweto	672
17	17622	Chinaunicom Hostmaster	644
18	133774	APNIC Hostmaster	588

Figure 4.4.28: Top Problem ASN

4.4.21 Top Volume ASN

The “Top Volume AS” report shows the Autonomous Systems with the highest volume of traffic coming from the monitored network. It can help you understand your traffic flow directions. If according to this report a specific AS has significantly more traffic than the others, then the network operator may consider getting a direct connection to this AS, to reduce traffic costs.

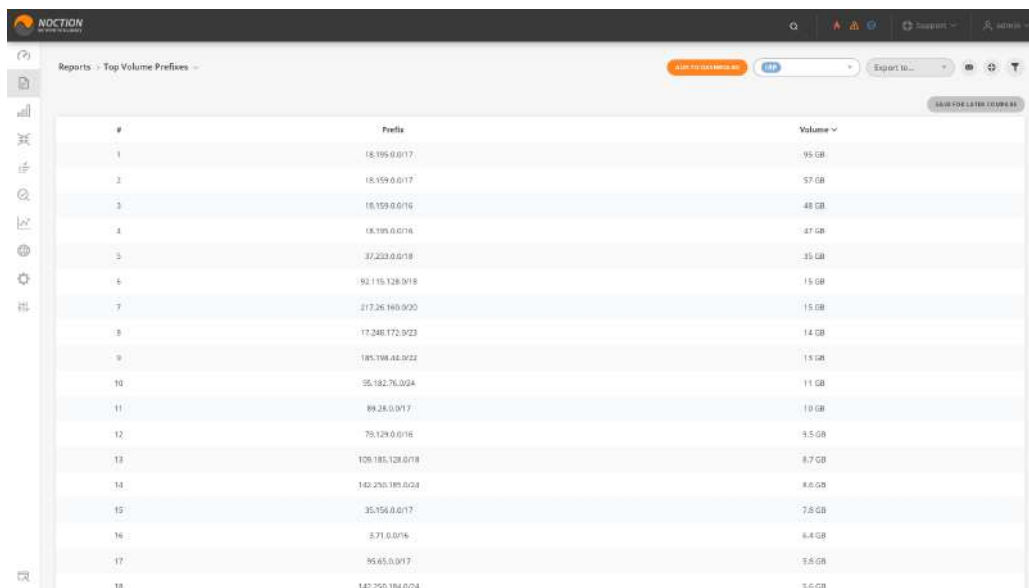


#	ASN	AS Name	Volume	95th Inbound (Mbps)	95th Outbound (Mbps)
1	16509	Amazon.com, Inc.	8.8 TB	-	109.1
2	15168	Google LLC	2.8 TB	-	0.6
3	4808	Chinaunicom Hostmaster	1.4 TB	-	0.2
4	56690	Cloud OpenDNS, LLC	1.2 TB	-	0.1
5	8531	Bezeq International Ltd.	1.1 TB	-	0.3
6	31252	StarNet Solusi SRL	0.9 TB	-	0.1
7	56047	APNIC Hostmaster	721 GB	-	0.2
8	56046	APNIC Hostmaster	665 GB	-	0.2
9	23724	IBT-CHINANET-CN	608 GB	-	0.2
10	56041	APNIC Hostmaster	477 GB	-	0.1
11	4812	SH JIN	449 GB	-	0.1
12	4765	APNIC Hostmaster	374 GB	-	0.1
13	58461	Chinanet Hostmaster	333 GB	-	0.1
14	17488	IBT-HATHAWAY-IR	302 GB	-	0.1
15	13335	Cloudflare, Inc.	270 GB	-	0.1
16	133481	APNIC Hostmaster	267 GB	-	0.0
17	26599	TELEFONICA BRASIL SA	258 GB	-	0.0
18	4837	IBT-CU-CN	241 GB	-	0.1

Figure 4.4.29: Top Volume ASN

4.4.22 Top Volume Prefixes

The “Top volume prefixes” displays a list of prefixes sorted by the total volume transferred to these remote networks.



#	Prefix	Volume
1	18.195.0.0/17	95 GB
2	18.159.0.0/17	57 GB
3	18.159.0.0/16	48 GB
4	18.195.0.0/16	27 GB
5	37.223.0.0/18	35 GB
6	92.115.128.0/18	15 GB
7	217.26.160.0/20	15 GB
8	17.246.172.0/23	14 GB
9	185.188.48.0/22	15 GB
10	55.182.76.0/24	11 GB
11	89.28.0.0/17	10 GB
12	75.129.0.0/16	9.5 GB
13	108.185.128.0/18	8.7 GB
14	142.250.185.0/24	8.6 GB
15	35.156.0.0/17	7.8 GB
16	8.71.0.0/16	6.4 GB
17	95.45.0.0/17	5.8 GB
18	142.250.184.0/24	5.6 GB

Figure 4.4.30: Top Volume Prefixes

4.5 Graphs

The GMI graphs are visual representations of data found in the reports. They provide an opportunity to quickly identify problems in the network, as well as forecast periodic changes in the traffic patterns of specific IRP instances and the overall network behavior. For each graph you can adjust the time period which you want to get the results for.

i Unless otherwise noted, all graphs are plotted using 5 minute intervals.

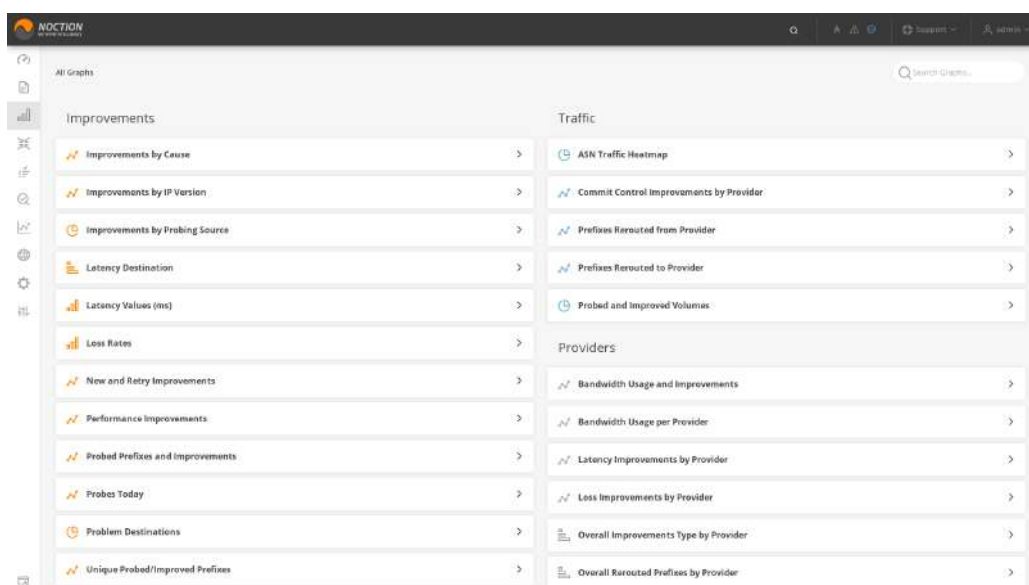


Figure 4.5.1: Graphs menu

To add any given graph to a dashboard, export it or print, consult the instructions provided in section 4.4

4.5.1 ASN Traffic Heatmap

ASN Traffic Heatmap is the graphical illustration of the AS traffic percentage which highlights the total amount of traffic by ASN and the calculated percentage from total traffic.

Each ASN is represented by a colored heatmap cell, with darker shades indicating higher traffic volume and lighter shades indicating lower traffic volume. The graph legend explains the color coding scheme used in the heatmap further.

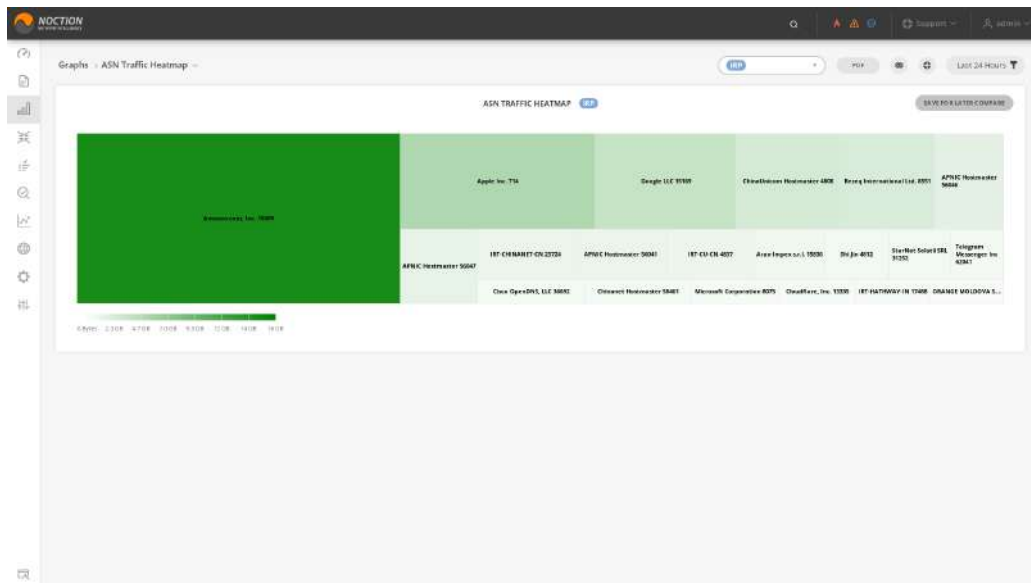


Figure 4.5.2: ASN Traffic Heatmap

4.5.2 Bandwidth Usage and Improvements

Bandwidth usage and improvements chart superimposes graphs for easy analysis and comparison of how improvement counts correlate with bandwidth.

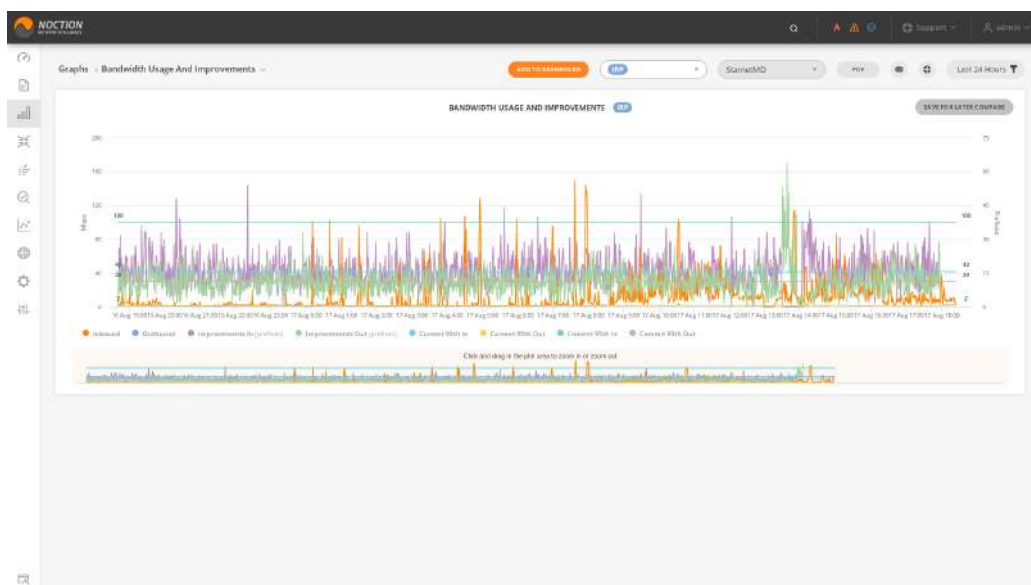


Figure 4.5.3: Bandwidth Usage and Improvements

4.5.3 Bandwidth Usage per Provider

The Bandwidth Usage graph shows the total usage for each of the providers. It allows you to compare the current traffic volume to the current 95th percentile and the commit 95th. The graph provides average, maximum and minimum traffic usage values for each of the providers during the selected time period.

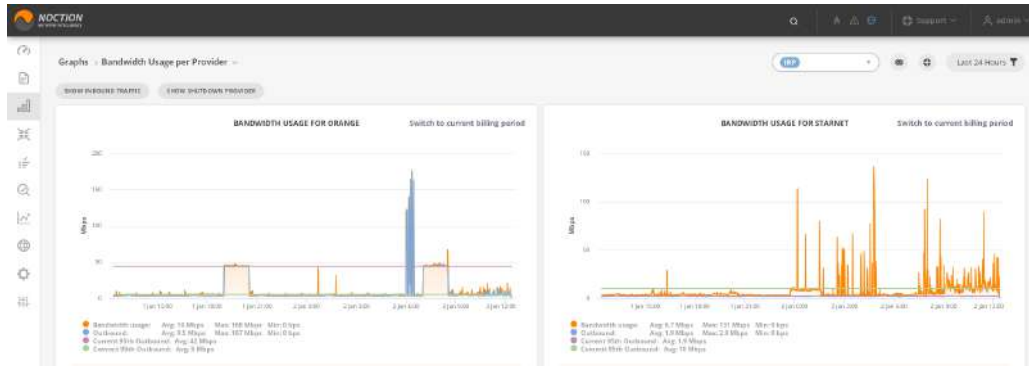


Figure 4.5.4: Bandwidth Usage per Provider

4.5.4 Commit Control Improvements by Provider

If the Commit Control algorithm is enabled on a selected IRP instance, this graph will display an overview of all the commit improvements that were enforced by the platform for a specific time period.



Figure 4.5.5: Commit Control Improvements by Provider

4.5.5 Improvements by Cause

The “Improvements by Cause” graph shows the improvements made by a specific IRP instance based upon performance, cost, commit, and outage detection. Various problem patterns can be detected in compliance with the platform’s activity shown in this graph. Different patterns can be noticed depending on the improvement mode that is currently running.



Figure 4.5.6: Improvements by Cause

4.5.6 Improvements by IP Version

This graph shows the number of improved prefixes by either IPv4 or IPv6 protocol.



Figure 4.5.7: Improvements by IP Version

4.5.7 Improvements by Probing Source

The graph displays the percentage of improvements that are differentiated by the probing source including Commit Control, Outage Detection, VIP Probing, Regular, and Retry Probing.

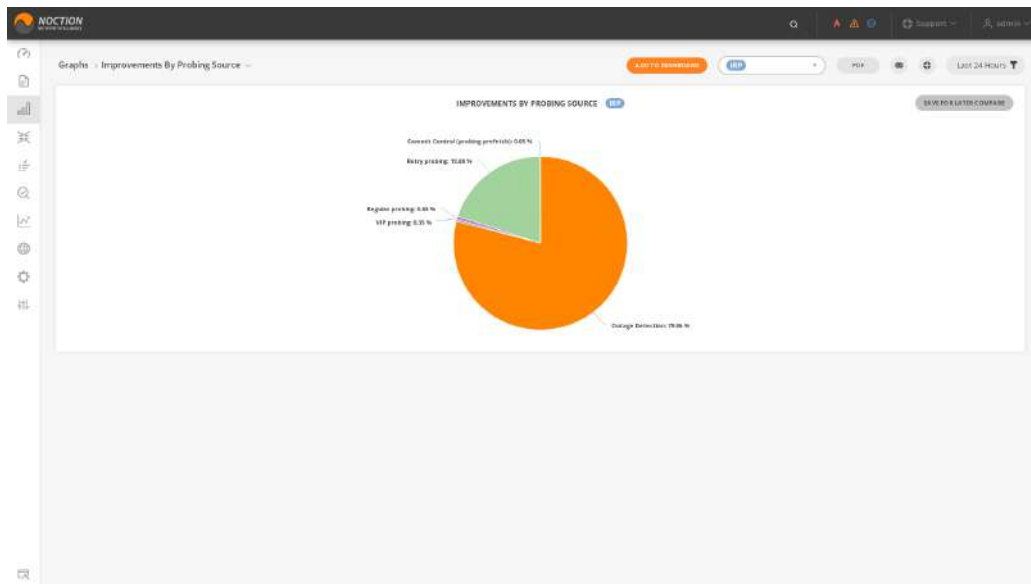


Figure 4.5.8: Improvements by Probing Source

4.5.8 Latency Destination

The Latency Improvements report presents an overview for average latency values before and after specific IRP instance optimization during a selected time period and depicts them across the following populations:

- Sample destinations that cover good and problematic destinations that an IRP instance probed during the selected time period
- Problematic destinations cover those routes/prefixes that IRP assessed to have excessive latency rates and identified a better route via a different upstream provider
- 50% or more cover those destinations where IRP was able to identify a route with latency improvements of 50% or more
- 20% or more cover those destinations where IRP was able to identify a route with latency improvements of 20% or more




Figure 4.5.9: Latency Destination

4.5.9 Latency Improvements by Provider

Graph displays a representation of the average number of latency-based improvements for each provider per selected IRP instance in five minute intervals. Specific provider lines can be removed from the graph by clicking the provider name in the graph's legend.




Figure 4.5.10: Latency Improvements by Provider

4.5.10 Latency Values (ms)

Latency values bar-chart highlights the latency averages before and after improvement for the designated groups of destinations.

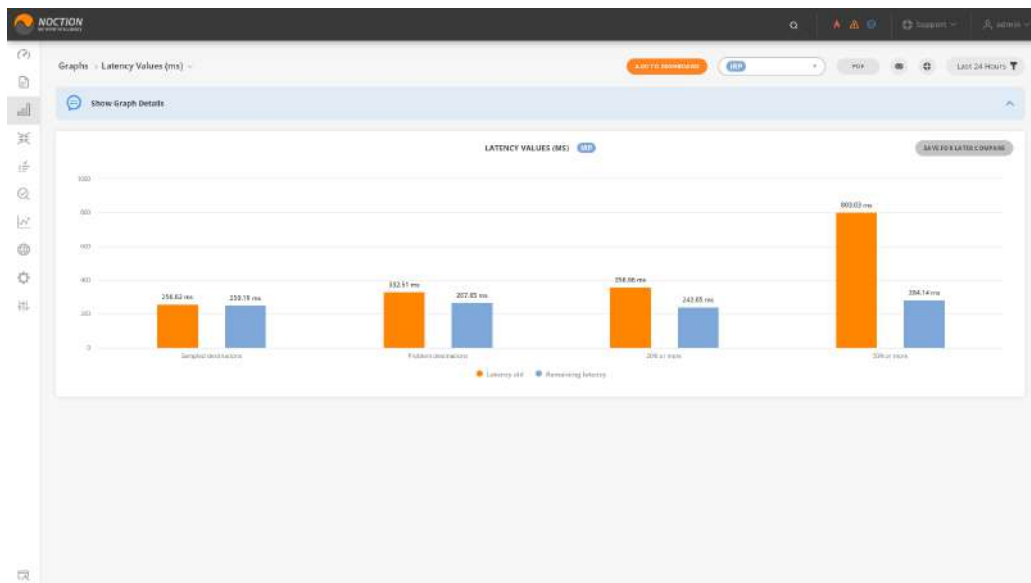


Figure 4.5.11: Latency Values (ms)

4.5.11 Loss Improvements by Provider

Graph displays a representation of the average number of loss-based improvements for each provider per selected IRP instance in five minute intervals.




Figure 4.5.12: Loss Improvements by Provider

4.5.12 Loss Rates

Loss values bar-chart highlights the packet loss averages before and after improvement for the designated groups of destinations.

Sampled destinations cover good and problematic destinations that registered traffic flowing to during a selected timeframe. Problem destinations include only the routes with packet loss improved by specific IRP instances.

Loss eliminated displays Problem destinations that an IRP instance was able to improve to a better route with zero loss.

Loss reduced: displays Problem destinations that IRP was able to improve but could not reduce loss altogether.

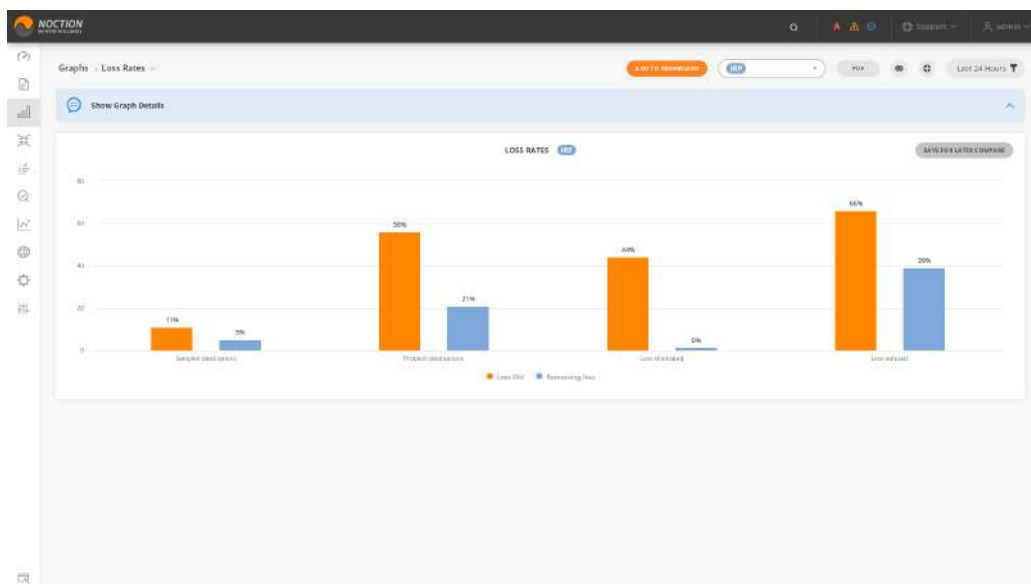


Figure 4.5.13: Loss Rates

4.5.13 New and Retry Improvements

After a particular improvement is made, IRP analyzes the path periodically. If the improvement is still valid, the platform leaves it as it is. The graph shows the number of improvements made by a selected IRP platform for the first time as related to those, which were simply confirmed after re-probing.

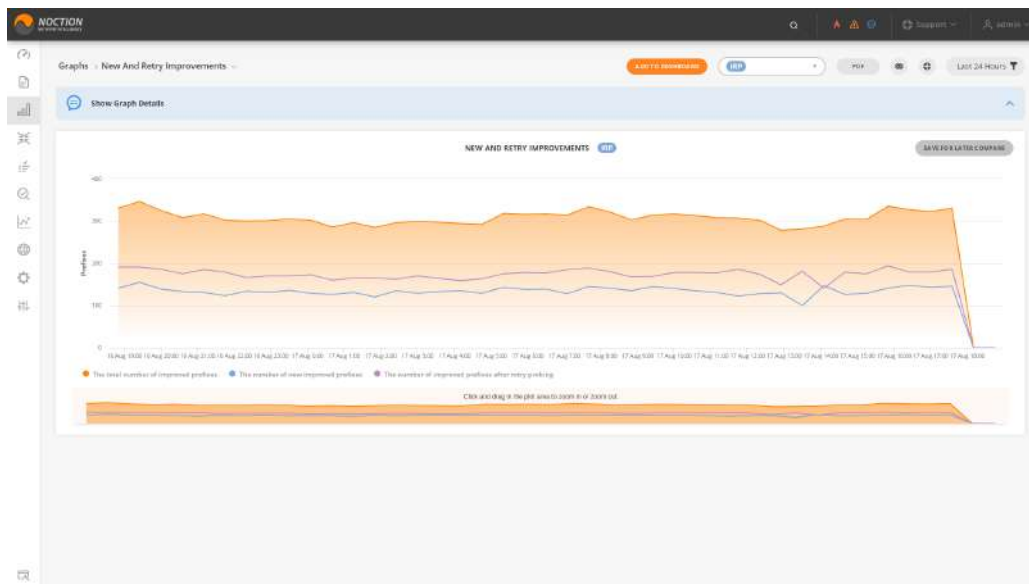


Figure 4.5.14: New and Retry Improvements

4.5.14 Overall Improvements Type by Provider

Graph shows the number of performance improvements vs other types of improvements per provider for a selected IRP instance.




Figure 4.5.15: Overall Improvements Type by Provider

4.5.15 Overall Rerouted Prefixes by Provider

The “Overall rerouted prefixes by provider” graph provides the number of prefixes that were rerouted from and to each of the providers. In the example below, the traffic exchange between the providers is

balanced. However, if the difference between the outgoing and incoming traffic is significant, then the quality of that provider's network should be questioned. "In" represents prefixes rerouted to a provider, "Out" represents prefixes rerouted from a provider. The "Excess In/Out" values represent the difference between the number of prefixes rerouted to a particular provider and the number of prefixes rerouted from this provider.

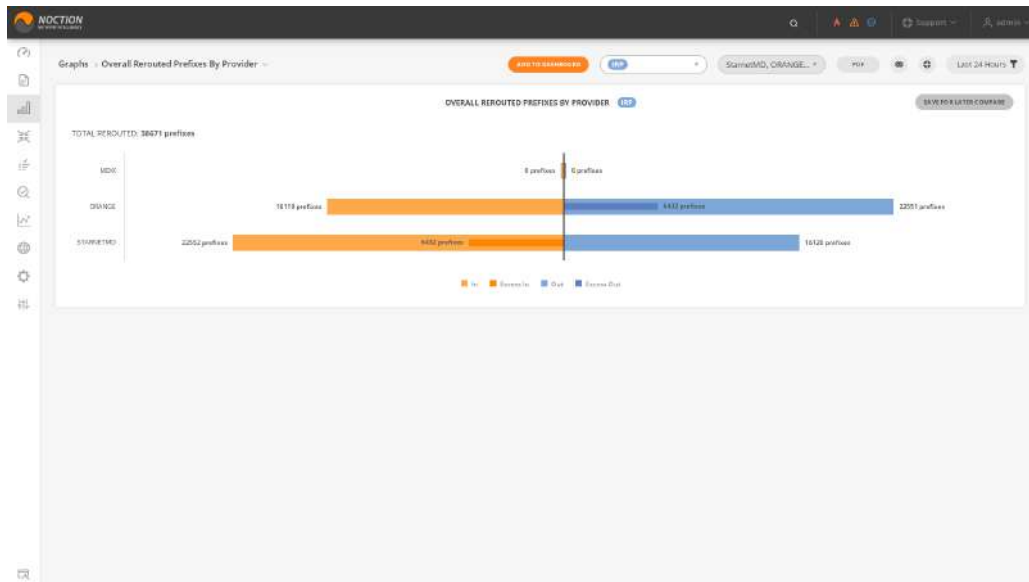


Figure 4.5.16: Overall Rerouted Prefixes by Provider

4.5.16 Performance Improvements

The "Performance Improvements" graph displays prefixes which were improved based upon a performance reason. By following this graph, the solved loss and latency issues, occurring in the network, can be monitored.

This graph provides the average, maximum and minimum number of improvements for each of the improvement reasons during the selected time frame. The maximum peaks are indicators of a loss or latency problem in the network.

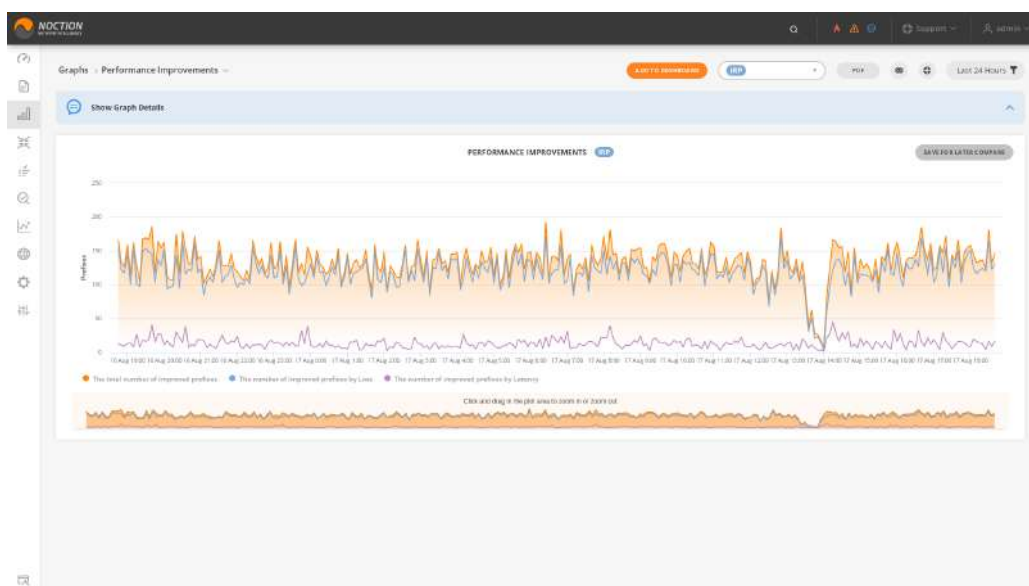


Figure 4.5.17: Performance Improvements

4.5.17 Prefixes Rerouted from Provider

The current graph displays the number of prefixes that were rerouted from a specific provider in order to resolve performance or cost issues.



Figure 4.5.18: Prefixes Rerouted from Provider

4.5.18 Prefixes Rerouted to Provider

The graph displays the number of prefixes that were rerouted to a specific provider in order to resolve performance or cost issues.



Figure 4.5.19: Prefixes Rerouted to Provider

4.5.19 Probed and Improved Volumes

The pie chart graph displays the amount of traffic volume improved by a particular IRP instance out of the total traffic volume.




Figure 4.5.20: Probed and Improved Volumes

4.5.20 Probed Prefixes and Improvements

The "Probed prefixes and Improvements" graph shows the number of improvements (Loss, Latency, Commit Control and Other) per selected time period as well as the amount of total probed prefixes.




Figure 4.5.21: Probed Prefixes and Improvements

4.5.21 Probes per Provider

The graph shows the number of probes for each provider for categories such as packet loss and latency, as well as the total count of probes.

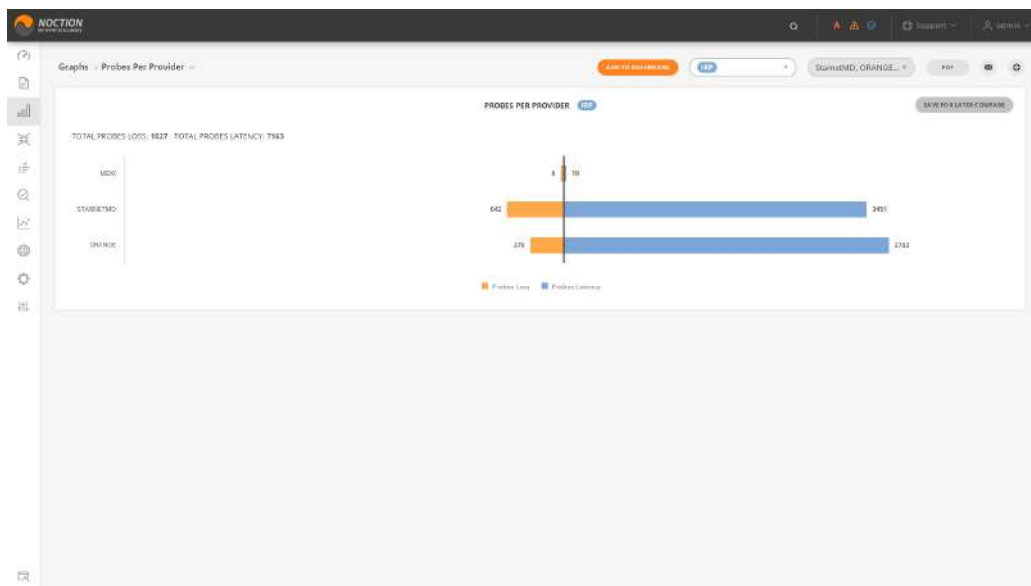


Figure 4.5.22: Probes per Provider

4.5.22 Probes Today

The “Probes Today” graph provides a visual representation of the probes performed within the last 24 hours, including probes that did not warrant an improvement or failed completely.

See also: [Probes Today](#).




Figure 4.5.23: Probes Today

4.5.23 Problem Destinations

Represents a pie chart with the distribution of improved destinations based on Packet Loss reason. It shows for what percentage of destinations Packet Loss was reduced and for what percentage of problematic destinations Loss was eliminated completely.




Figure 4.5.24: Problem Destinations

4.5.24 Provider Performance

During the day IRP instances make many probes to determine performance characteristics of alternatives routes.

Provider performance chart presents an aggregated view of this measurements for providers. The data on the report highlights Packet loss and Average latency per provider putting the provider with best packet loss data at the top.




Figure 4.5.25: Provider Performance

4.5.25 Providers Bandwidth Usage

The Providers bandwidth usage graph illustrates inbound, outbound, and total bandwidth details on the same page as per a specific IRP instance. It allows cross-comparison between providers. The chart includes options to hide/show providers of interest during a selected period via the graph legend.



Figure 4.5.26: Providers Bandwidth Usage

4.5.26 Providers Efficiency

The horizontal bar charts are the graphical representation of the specific measurement results (packet loss, latency, successful and failed probes) obtained by IRP while performing numerous probes via all the providers.

Note:

- The provider with the best packet loss data is always placed on top of the others
- Filtering by date range and providers allows reviewing of past data or average performance over longer time intervals as well as showing/hiding specific providers.

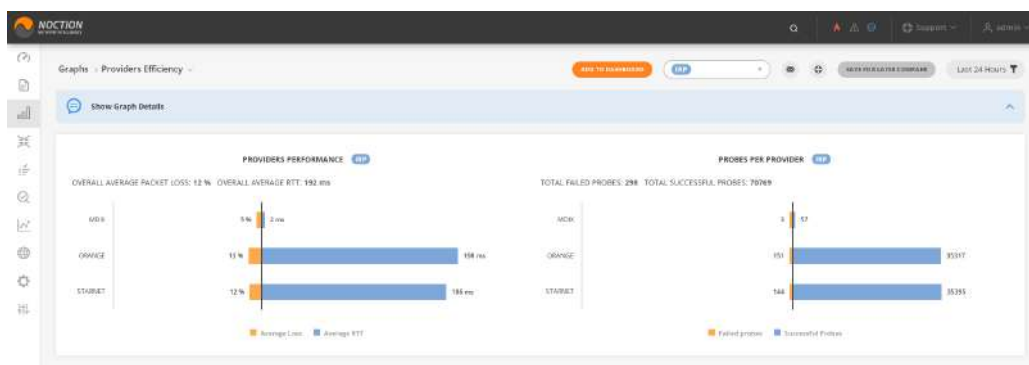


Figure 4.5.27: Providers Efficiency

While some providers seem better performing (for example, MDIX in the provided screen capture), the fact that they display a much smaller number of probes clearly indicates they are Internet Exchanges with only a few peers interchanging data with a given network.

A table view with exact details of the data is available in the report form as well.

See also: [Providers Efficiency](#)

4.5.27 Total and Improved Traffic

This graph allows users to see the positive impact of IRP instances on the network. It provides the amount of inbound/outbound improved traffic, as related to the total inbound/outbound traffic.




Figure 4.5.28: Total and Improved Traffic

4.5.28 Total Bandwidth Usage

The graph shows total inbound, outbound bandwidth as well as the In/Out 95th value as per specific IRP instance.




Figure 4.5.29: Total Bandwidth Usage

4.5.29 Unique Probed/Improved Prefixes

The "Unique Probed/Improved Prefixes" graph shows the number of improvements per time unit as reported to the total probed unique prefixes.

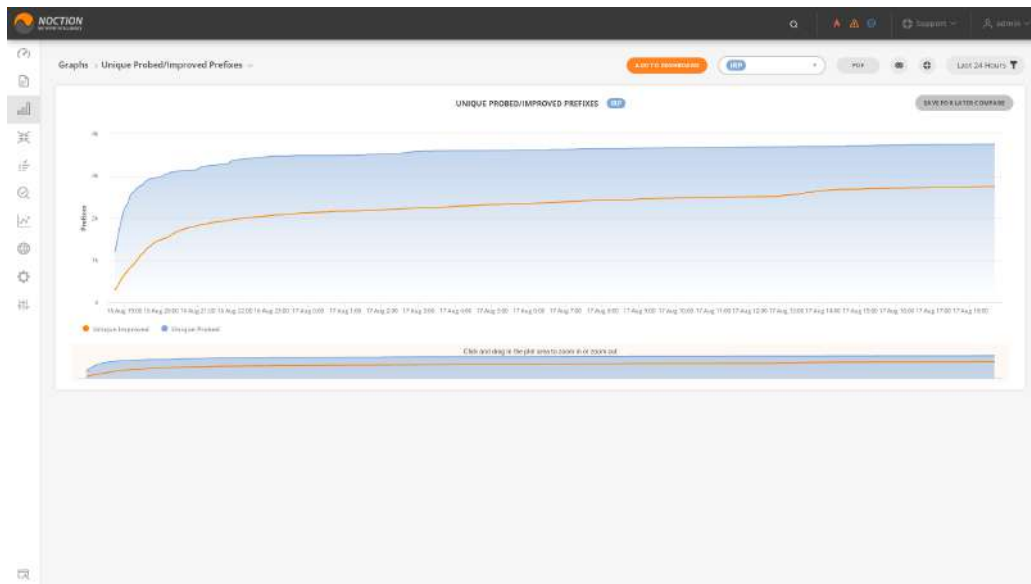


Figure 4.5.30: Unique Probed/Improved Prefixes

4.6 Inbound

The GMI's Inbound section is visually divided into two parts, each containing a set of links to reports, graphs, rules, and configuration options for both the Inbound Performance Optimization and the Inbound Commit Control features.

i The details on the Inbound Commit Control and Inbound Performance Optimization configuration are provided in the following sections: [Inbound Commit Control](#) and [Inbound performance optimization configuration](#).




Figure 4.6.1: Inbound section

4.6.1 Inbound Commit Control

Starting with version 3.4 IRP introduced the Inbound bandwidth optimization feature, the configuration and reports for which are available in GMI.

4.6.1.1 Current Inbound Improvements

The “Current Inbound Improvements” report provides a list of the currently active improvements. The report shows the improved prefix and the providers from and to which the traffic was redirected. It also provides the reason, which the improvement was based upon.

If necessary, specific improvements can be manually deleted by clicking the check-box next to the needed prefix, afterward clicking the “Remove selected” button.

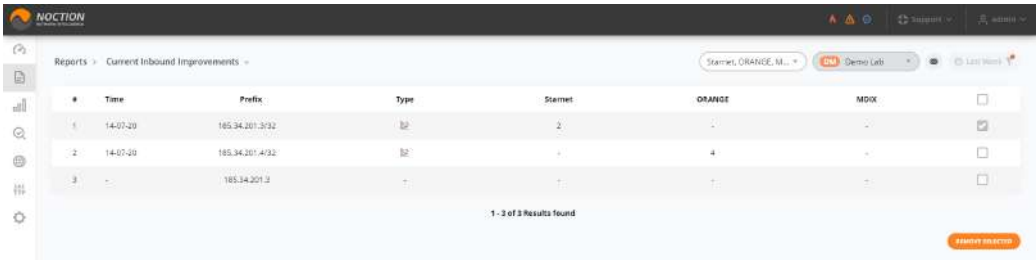


Figure 4.6.2: Current Inbound Improvements

4.6.1.2 Historical Inbound

The Historical Inbound traffic report provides an overview and analysis of the incoming traffic over a selected period of time per prefixes and the associated providers.

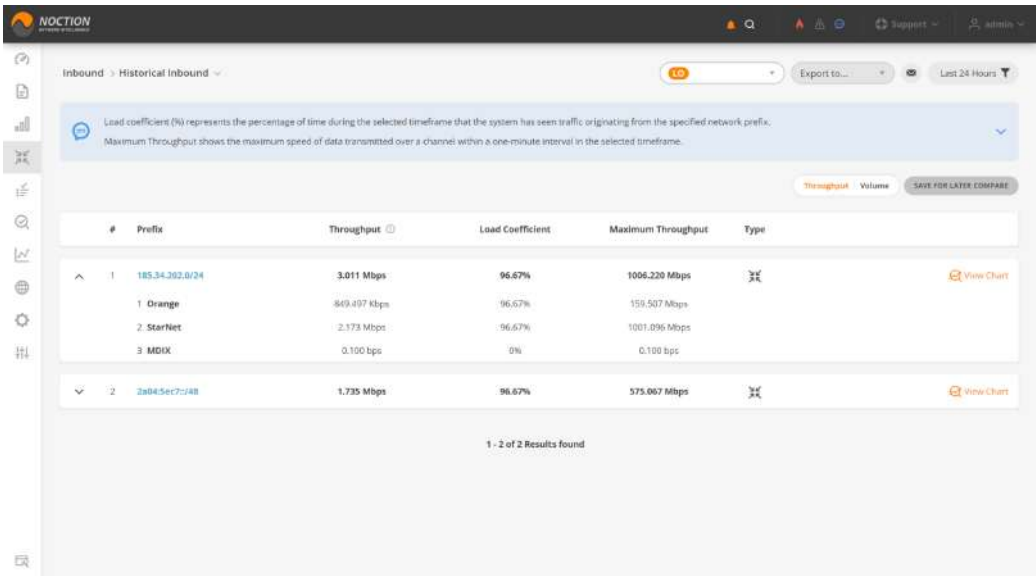


Figure 4.6.3: Historical Inbound

The report includes a range of relevant throughput and volume metrics related to inbound traffic, such as:

- Prefix** the network prefix associated with traffic
- Throughput** the average rate at which data is transmitted over a given period of time for a specific prefix
- Load coefficient** the percentage of time during the selected timeframe that the system has seen traffic per the specified network prefix
- Maximum throughput** shows the maximum speed of data transmitted over a channel within a one-minute interval in the selected timeframe
- Volume** this metric denotes the total amount of data transmitted or received for a specific prefix during the analyzed period
- Count of updates** the number of updates or changes made to the traffic data related to a specific prefix

Maximum volume similar to maximum throughput, this metric indicates the highest volume of data transmitted or received for a specific prefix during the analyzed period

Type categorizes traffic based on its direction. Inbound refers to data coming into a network from external sources, while Transit represents data passing through the network without being the final destination

View Chart displays a graphical representation of the throughput or volume of the inbound traffic for each specific prefix.

4.6.1.3 Inbound Improvements History

The “Inbound Improvements History” report provides a detailed list of past inbound improvement actions. The report shows the improved prefix and the new prepends count towards a provider.



#	Host	Time	Prefix	Provider Name	Type	Prepends	Action
1	ORANGE	10-07-20	195.34.201.0/32	ORANGE		4	Added
2	TEARINC	10-07-20	195.34.201.0/32	TEARINC		2	Added

Figure 4.6.4: Inbound Improvements History

4.6.1.4 Inbound Traffic Distribution

The graph highlights distribution of inbound traffic across inbound prefixes and providers. Inbound traffic distribution is available in the report form as well and includes a feature to cross-compare current values with a saved inbound traffic distribution from the recent past.



Figure 4.6.5: Inbound Traffic Distribution

4.6.1.5 Moderated Inbound Improvements

IRP includes a moderation feature that allows review and approval of Inbound improvements in GMI. If this feature is enabled, Inbound improvements are not automatically announced but instead, they are queued for review. If an inbound improvement has not been approved in time for a subsequent cycle of improvements an IRP instance will review the proposal itself and will adjust as needed thus making another recommendation.

The screenshot shows the 'Moderated Inbound Improvements' section of the NOXION interface. It features a table with columns: #, Time, Orange, Startset, MDIX, Prefix, and Type. There are three rows of data, each with a checkbox and a 'confirm' button. The interface also includes a sidebar with navigation icons, a top navigation bar with search and user options, and a bottom status bar indicating '1 - 3 of 3 results found'.

#	Time	Orange	Startset	MDIX	Prefix	Type
1	03/09/21	-	6 (B)	-	24945ec7:0a8	S2
2	03/09/21	0	4 (B)	-	185.34.200.0/24	S2
3	-	-	-	-	185.34.201.0/24	-

Figure 4.6.6: Moderated Inbound Improvements

Inbound Improvements moderation highlights the number of additional prepends proposed by IRP and allows submitting or rejecting individual improvements as well as a batch of selected improvements.

4.6.2 Inbound Performance Optimization

The Inbound Performance Optimization capability was added to the Intelligent Routing Platform starting with version 4.0. The feature is based on IRP's automated network performance metrics identification algorithm and the use of customer traffic engineering (TE) capabilities of the client's providers.

How Inbound performance works:[Inbound performance optimization](#)

4.6.2.1 Inbound Performance Feed

Inbound Performance Feed displays a list of inbound network performance improvements (per single IRP instance) that have been suggested, automatically applied, or manually enabled by a GMI user according to the pre-established set of Inbound Optimization rules. Each active/suggested inbound improvement entry contains the following details:

Status indicating whether the improvement is active/inactive

Improvement Type indicating if the improvement is Automated (suggested and activated by the system), Moderated (suggested by the system but not yet enabled by the user) or Manual (manually enabled by the user)

Description contains the name of the rule which triggered a particular action

TE Mark represents the Traffic Engineering community used within a particular rule

AS Path Match / BGP Community offers details on the selection criteria of the prefixes that are due to be improved as per the rule

Confirmation shows the time until the next improvement actuality verification

Initial Stats displays the packet loss and latency values prior to the improvement

Improvement Results displays the packet loss and latency values post the improvement activation

Enabled On displays the time and date the improvement was activated

Actions provides the options to enable/disable or completely delete an individual improvement.

Status	Improvement Type	Description	TE Mark	AS Path Match / BGP Community	Confirmation	Initial Status	Improvement Results	Enabled On	Actions
Active	Manual	Rule 1	39.49	as-path matches 25454 *	-	-	-	05-11-23 16:43:01	DISABLE
Active	Manual	Rule 2	50.60	community contains 31252-999	-	-	-	06-11-23 11:30:08	DISABLE
Active	Automatic	Rule 3	65300.111	as-path matches 65311 *	-	Loss 21% RTT 180ms	Loss 0% RTT 45ms	06-11-23 11:01:58	DISABLE
Active	Moderated	Rule 5	1200.3333	community contains 5880-888	-	Loss 48% RTT 220ms	Loss 2% RTT 53ms	06-11-23 11:05:41	DISABLE
Active	Manual	Rule 7	777.1300	as-path matches 71300	-	-	-	06-11-23 11:18:07	DISABLE

Figure 4.6.7: Inbound Performance Feed

4.6.2.2 Inbound Performance History

The Inbound Performance History tab displays a list of past (non-actual) improvements, with details of the time and date the improvement was deactivated.

Improvement Type	Description	TE Mark	AS Path Match / BGP Community	Enabled On	Disabled On
Moderated	Rule 5	1200.3333	community contains 5880-888	06-11-23 11:05:41	06-11-23 11:30:08
Manual	Rule 1	39.49	as-path matches 25454 *	05-11-23 16:43:01	05-11-23 16:43:01

Figure 4.6.8: Inbound Performance History

4.6.2.3 Inbound Performance Rules

Inbound Performance Rules tab displays a list of the inbound network performance rules (per single IRP instance) created by a GMI user according to the user's network specific needs and requirements. Each rule can be turned on/off, to enable or disable the subsequent validation of inbound traffic characteristics and the detection of possible inbound routing improvements. There are also options to instantly apply the improvement as per the rule manually, edit or delete it.

Description	Provider	TE Marker	AS Path Match / BGP Community	Active
Rule 1	Orange	39.49	as-path matches 25454 *	ON OFF ✓ Applied Manually ✖
Rule 2	StarNet	50.60	community contains 31252-999	ON OFF ✓ Applied Manually ✖
Rule 3	MOX	65300.111	as-path matches 65311 *	ON OFF ✓ Apply Manually ✖

Figure 4.6.9: Inbound Performance Rules

To create a rule, click the NEW RULE button. Introduce details in the corresponding fields:

- **Description** - introduce the meaningful details for the rule you are about to create
- **Provider** - select the provider that the rule will apply to
- **Provider TE marker** - indicate the traffic engineering marker (a list of BGP communities) which has special meaning in context of traffic engineering capabilities of a specific provider
- **Improvement Activation Threshold (Packet Loss)** - indicate the packet loss value threshold (% delta worsening from the initial statistical model values)

- **Improvement Activation Threshold (Packet Latency)** - indicate the packet latency value threshold (% delta worsening from the initial statistical model values)
- **AS Path Match / BGP Community** specify the selection criteria (**filter expression**) of the prefixes that are due to be improved as per the rule

Filter expression

Filter expression is a boolean expression of sub-expressions:

- `as-path matches as-path-regex`; for example: `as-path matches 1 2 10-20 [40 50-60]+ (. 200)*`
- `community contains bgp-community`; for example: `community contains 55:66`

Boolean expression supports just 'and' and 'or' operators.

For example, filter expression `"as-path matches .* 20 [^40 50] and community contains 55:66"` would match an external prefix, the RIBIN record for which denotes:

- `as-path "90 20 60"`
- `community set "2:3 55:66 40:50"`

As-path regex

As-path regex implements subset of juniper as path regex, relaxing group operator '()' to support any expression within it. EBNF grammar:

```
regex = expr;
expr = seq | subexpr;
subexpr = repeat | term | group;
seq = (subexpr, seq) | subexpr;
repeat = ((group | term), "*")
| ((group | term), "+")
| ((group | term), "{", uint8, ",", uint8, "}");
any_of = "[", {range | atom}-, "]";
none_of = "[^", {range | atom}-, "]";
group = "(", expr, ")";
term = range | atom | wildcard | any_of | none_of;
wildcard = ".";
range = atom, "-", atom;
atom = uint32;
```

Examples

Regex matches input's start and end, as if `"^regex$"`. If you need to skip numbers from start or end, use combination of 'repeat zero or more times' with 'wildcard', for example

- `"9 5"` matches only entire as-path `"9 5"`, whereas
- `".* 9 5 .*"` can match `"1 2 10 9 5 30"`
- Group and repeat on or more times operators `"(. 5 6)+"` can match `"1 5 6 2 5 6 3 5 6"`
- Repeat at least m and at most n times operator `"2{3,4}"` can match `"2 2 2"`
- Range `"5-10"` can match `"6"`
- Any of operator `"5 [6 8]"` can match `"5 8"`
- None of operator `"5 [^6 8]"` can match `"5 9"`
- The complex regex `"5-10 [1 20-30]{1,2} (5 7{3,4})+ [^10]"` can match `"6 1 2 5 7 7 7 11"`, or `"6 1 25 5 7 7 7 5 7 7 7 11"`, but not `"6 1 25 5 7 7 7 10"` because the path can't end with 10

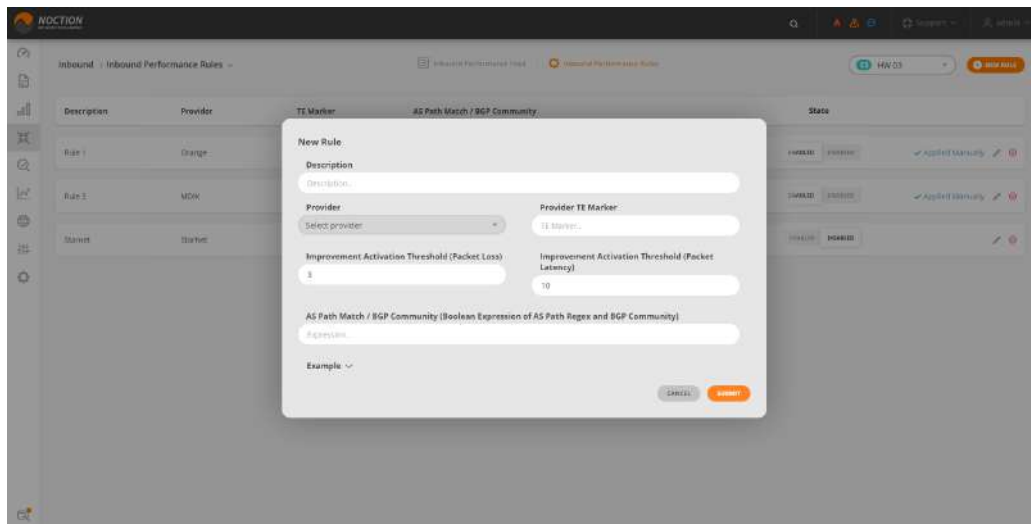


Figure 4.6.10: Creating an Inbound Performance Rule

Refer to: [Inbound performance optimization configuration](#) for the Inbound Performance configurations.

4.6.2.4 Inbound Performance Report

Inbound Performance Report contains two graphs showing the effect of the inbound improvements on bandwidth as well as packet loss and packet latency values. The presented data is displayed in accordance with a single inbound rule and IRP instance selected.



Figure 4.6.11: Inbound Performance Report Bandwidth Graph



Figure 4.6.12: Inbound Performance Report Loss/Latency Graph

4.7 Routing Policies

To see the full list of routing policies configured for a particular IRP instance, select the Policies option from the main menu and click on the Routing Policies option. The list displays all the enabled and disabled routing policies.

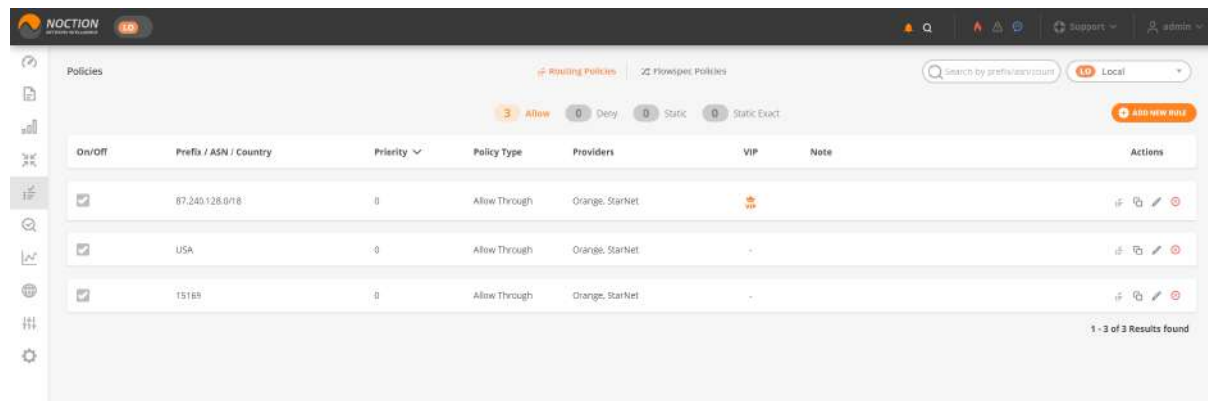


Figure 4.7.1: Routing Policies

As shown in the screenshot above, the list provides the following details:

- The ON/OFF state of each policy
- The country, prefix or ASN to which the policy applies
- Priority of a policy

i Priority specifies what policy to use in case of overlapping prefixes as might be the case of a large aggregate prefix policy extending over an ASN policy

- The type of the policy (allow/deny/static/static_exact)
- Providers which are affected by the policy
- The status of the VIP reprobng (enabled/disabled)
- A note to describe the policy
- The available actions for the policy. These include:
 - Show detailed routing policy information
 - Duplicate the routing policy
 - Edit the routing policy
 - Remove the routing policy

When opening up the detailed routing policy information, the following details become available:

- Time passed after the last probe was performed towards this prefix
- The time left before the next scheduled probe towards this prefix
- From and To details of an existing improvement relating to this policy.

The screenshot shows the NOCION IRP Global Management Interface. The top navigation bar includes the NOCION logo, a search icon, and a user profile icon. The main content area is titled "Routing Policies" and shows a policy for "87.240.128.0/18" with a status of "Local Active". Below the title, there is a "SUMMARY" section with fields for "Policy type: Allow Through", "Priority: [0]", "VIP improvement: [icon]", "Providers: Orange, StarNet", and "Local improvement only: DISABLED". A "Note" field is also present. The "Probed prefixes" section contains a table with the following data:

Prefix	AS Number	AS Name	Country	Last probed	Next probe	Details
87.240.166.0/24	47542	Vkontake Ltd	-	03/29/22 12:23:32 PM	03/29/22 1:23:32 PM	Old loss 100%, avg rtt 0 ms, new loss 0%, avg rtt 57 ms.
87.240.167.0/24	47542	Vkontake Ltd	-	03/29/22 12:23:32 PM	03/29/22 1:23:32 PM	Old loss 100%, avg rtt 0 ms, new loss 0%, avg rtt 60 ms.
87.240.128.0/18	47541	Vkontake Ltd	-	03/29/22 12:23:06 PM	03/29/22 1:23:06 PM	

Figure 4.7.2: Detailed routing policy view

A new routing policy can be added by pressing the "ADD NEW RULE" button. The following window will open up:

The screenshot shows the "New Routing Policy" window. It has a title bar with "New Routing Policy" and a status icon. The form contains the following fields:

- Policy by:** A dropdown menu with "Prefix" selected.
- IP prefix:** A text input field with a character count of "0/500".
- Policy type:** A dropdown menu with "allow" selected.
- Policy notes:** A large text area with a character count of "0/200".

At the bottom of the window, there is a "Step 1 from 3" indicator, a "CANCEL" button, and a "NEXT" button.

Figure 4.7.3: Routing policy window

Depending on the "Policy by" field selection, some of the following parameters should be configured to add a routing policy:

- Prefix: The prefix to which the policy is applied
- ASN: The ASN to which the policy is applied
- Country: The country to which the policy is applied. Such a policy applies to all prefixes known to be located in the country.
- Note: A note to describe the policy
- Policy Type: The type of the policy (allow/deny/static/static_exact)
- A community to mark improvements
- A flag to set if policy is VIP
- The policy status (enabled/disabled)
- A flag to indicate if an ASN policy should propagate to downstream ASN.
- A flag indicating if a global improvement is allowed.
- Priority slider or text-box will set a policy's priority.
- The Providers for which the policy should be respected

- Cascade flag for ASN policies that indicate if the policy should be enforced only for this AS or also for downstream AS. If cascading is enabled it ensures that the policy is enforced for all traffic that will eventually be routed through this AS.

Refer policy configuration parameters [Routing Policy](#) for details.

Click any of the “Allow”, “Deny”, “Static” or “Static Exact” options to see a list of policies of a particular type.




Figure 4.7.4: Routing Policy Type

4.8 Flowspec Policies

To review Flowspec policies, select the Routing Policies option from the main menu and proceed to the Flowspec Policies tab.

⚠ Flowspec must be enabled globally for each IRP instance under Configuration > Global as well as each router under BGP > Router Name > Advanced tab.

The list displays all the enabled and disabled Flowspec policies configured for a particular IRP instance. Click any of the “Redirect”, “Throttle”, “Drop” or “Redirect IP” options to see a list of policies of a particular type. Flowspec policies are also grouped as follows: Prefix Policies, ASN Policies, Country Policies, and Other Policies (Protocol/DSCP).

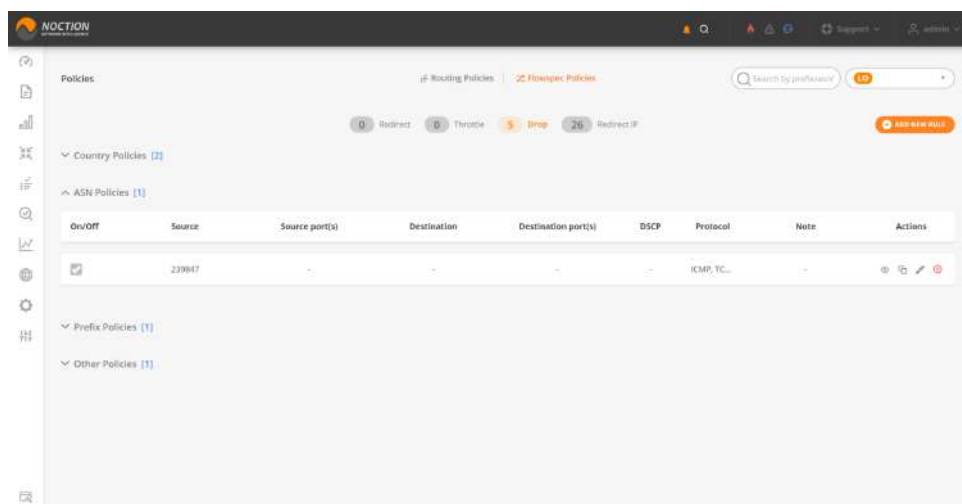


Figure 4.8.1: Flowspec policies

Depending on the selected policy type tab, the list highlights:

- The ON/OFF state of each policy
- A source ASN/prefix/country and port(s) for matching packets
- A destination prefix and port(s) for matching packets
- DSCP traffic classification value

- Protocols of matching packets, e.g., TCP, UDP, or ICMP
- Redirect IP or Provider for the redirect to VRF policies
- A note to describe the policy
- The available actions for the policy:
 - Display detailed information about the FlowSpec policy
 - Duplicate the FlowSpec policy
 - Edit the FlowSpec policy
 - Remove the FlowSpec policy

A FlowSpec policy is added by clicking on the designated “ADD NEW RULE” button.

Figure 4.8.2: FlowSpec policy popup

The first step prompts you to choose the IRP instance as well as the type of FlowSpec policy you are about to create. There are 4 options:

- Policy by Prefix
- Policy by ASN
- Policy by Country
- Protocol/DSCP policy

Depending on the selection, some of the following parameters should be configured:

- Source Prefix/ASN/Country/Port(s): The source ASN/prefix and port(s) of the IP packets that match. A prefix in CIDR notation or a single IP address should be provided. Multiple valid TCP/UDP ports can be provided, as well as port ranges.

i ASN/Country FlowSpec policies can be set up only on the source of IP packets

- Ensure that routers are capable of processing a large number of FlowSpec rules before setting policies for an AS/Country consisting of a very large number of network segments (prefixes)

- Destination Prefix/Port: The destination prefix/port attribute of the IP packets that match. Same rules as for Source Prefix/Port(s) apply
- Protocols: packet protocols that match the policy. Can be filtered down to one or a combination of the following protocols: TCP, UDP, ICMP
- DSCP traffic classification value

- Policy Type: The type of the policy (Throttle, Drop, Redirect, or Redirect IP)
- Provider specifies one of the provider identifiers where traffic will be redirected. The provider is set only for Redirect policies

i In order for Redirect policies to be implemented on routers, an extended community value must be assigned to them, and VRF must be configured on the router itself. Consult Noction support for details

- Rate limits the allowed bandwidth usage for matching traffic. The value is set only for Throttling policies. The rate specifies a number in the range of 1-4200 Mbps
- Exempted Prefixes/ASNs are the lists excluded from country policies. The fields are relevant to the country policies only.

w When creating an ASN Flowspec rule and selecting the Throttle policy type, the value introduced under “Rate limit” applies to individual prefixes within the AS, and not the Autonomous System as a whole

i If the provided Flowspec policy attributes are incomplete or invalid on attempting to submit it, a warning will be raised

For more details, refer to Flowspec configuration parameters, for example: `global.flowspec`, `core.flowspec.max`, `core.flowspec.max_ipv6`, `bgpd.peer.X.flowspec`, `peer.X.flowspec.ipv4.redirect_community`, `peer.X.flowspec.ipv6.redirect_community`

4.8.1 Policies by Country

Policies by Country, as part of the Flowspec Policies functionality, provides IRP users with the geo-blocking capability at the BGP level. Such policies allow network operators to restrict internet traffic by manipulating routing decisions based on geographic regions, particularly countries. Network administrators can define packet filtering rules based on additional parameters, such as protocols, port numbers, and destination prefixes, enabling a more granular approach. Moreover, users can access lists of affected prefixes and ASNs associated with each specific country for every policy, facilitating more informed routing decisions.

On/Off	Source	Source port(s)	Destination	Destination port(s)	Protocol	Affected Prefixes	Exempted Prefixes	Exempted ASN(s)	Note	Actions
<input checked="" type="checkbox"/>	Albania	-	3.7.7.0/2	777	TCP, UDP	249	1	1	Sanctions/embargos	
<input checked="" type="checkbox"/>	Austria	-	-	-	ICMP, TC...	2544	0	0	-	
<input checked="" type="checkbox"/>	Angola	-	-	-	ICMP, TC...	123	0	0	-	
<input checked="" type="checkbox"/>	Albania	-	-	-	ICMP, TC...	233	0	0	Law compliance	
<input checked="" type="checkbox"/>	Andorra	-	-	-	ICMP, TC...	54	0	0	Excessive traffic	
<input checked="" type="checkbox"/>	Afghanistan	-	195.34.262.0/25	22	ICMP, TC...	122	0	0	-	
<input checked="" type="checkbox"/>	Bahamas	22	-	-	ICMP, TC...	53	0	0	-	
<input checked="" type="checkbox"/>	Costa Rica	22	-	-	ICMP, TC...	205	0	1	Potential security threat	

Figure 4.8.3: Policies by Country

Once enabled, each policy contains statistics on the number of affected prefixes that can be viewed for specific details on the actual list of prefixes as well as ASNs that the prefixes belong to. The history of changes is maintained for each policy.

ASN	ASN Name	Number of Prefixes	Actions
6752	ANDORRA TELECOM, S.A.U.	17	Exempt ASN
1	89.136.2.0/23		Exempt Prefix
2	94.125.180.0/22		Exempt Prefix
3	185.341.26.0/23		Exempt Prefix
4	46.175.156.0/22		Exempt Prefix
5	80.80.84.0/22		Exempt Prefix
6	80.80.92.0/22		Exempt Prefix
7	89.136.4.0/22		Exempt Prefix
8	185.4.52.0/22		Exempt Prefix
9	185.33.5.0/22		Exempt Prefix
10	185.87.36.0/22		Exempt Prefix
3009	M247 Europe SRL	1	Exempt ASN
15224	Adobe Inc.	13	Exempt ASN
37518	Undefined	5	Exempt ASN
64614	Foraspant Cloud Ltd	1	Exempt ASN
48070	PPTECHNOLOGY LIMITED	1	Exempt ASN

Figure 4.8.4: Policies by Country - Affected Prefixes view

i To avoid routers overflowing with multiple policies, there are some limitations to the number of Country/ASN policies one can create. A maximum of three policies for a particular Country/ASN can be created, with each of them being different from the other two by having an IPv4 destination prefix, an IPv6 destination prefix, or no destination prefix at all

4.9 Troubleshooting

The GMI Frontend provides several troubleshooting tools that can offer you quick information on the specific remote networks.

4.9.1 Looking glass

A ping, traceroute or MTR can be run from the GMI Frontend via different providers available in particular IRP instances. The results are displayed on the same page in specific query results sections that can be expanded.




Figure 4.9.1: Looking glass

4.9.2 Prefix probing

GMI features a manual prefix probing function. One can manually submit a specific IP address for probing and optimization by particular IRP instances. Valid hostnames and IP addresses can be entered. The probing results will be displayed on the same page. Please note that the selected IRP instances probes the exact entered IP address.

In case the IP address doesn't respond, an IRP instance runs the indirect probing process and optimizes the whole prefix based on the indirect probing results.

Prefix probing allows to automatically or manually choose the best path for a particular IP address or prefix.

In order to check what are the possible paths for a prefix, enter the IP address or prefix into the field. Checkmark the "Fast Probing" option to send only a small batch of 5 control packets (instead of a 100 packets full probe) to establish performance characteristics through different providers. If you prefer the prefix to be improved automatically, tick the check box "Improve automatically" and press the "Submit for probing" button. Otherwise, just press the "Submit for probing" button and wait until the system returns the probing results.

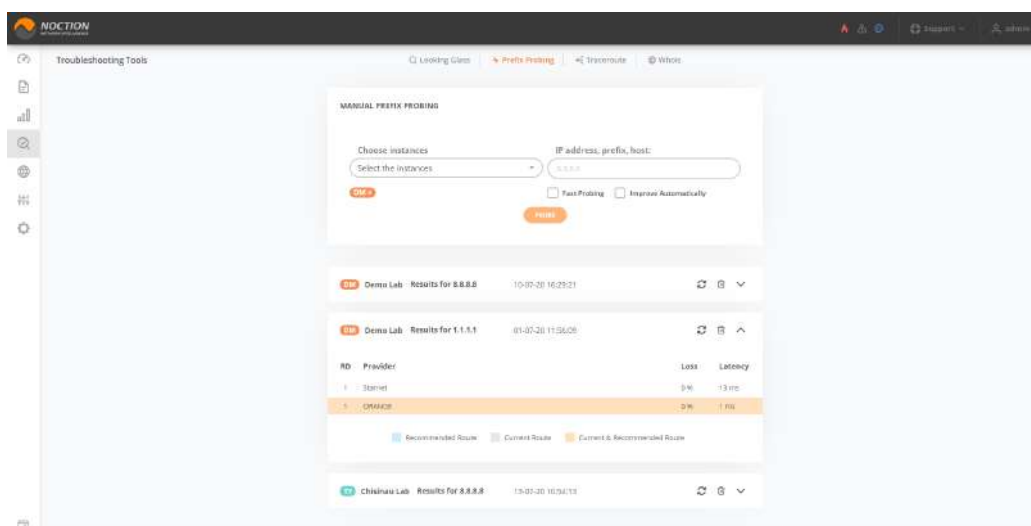


Figure 4.9.2: Prefix Probing

The system returns the probing results by displaying the loss and latency values for each of the providers. After the probing results are received, choose the path you consider the best one. Note, that the current and the best path are highlighted.

4.9.3 Traceroute

A traceroute to a specific IP address or hostname can be run from any IRP instance. Results will be displayed in the graphical and detailed table formats, as in the examples below. The graph can be presented in full view and exported as an image.




Figure 4.9.3: Traceroute

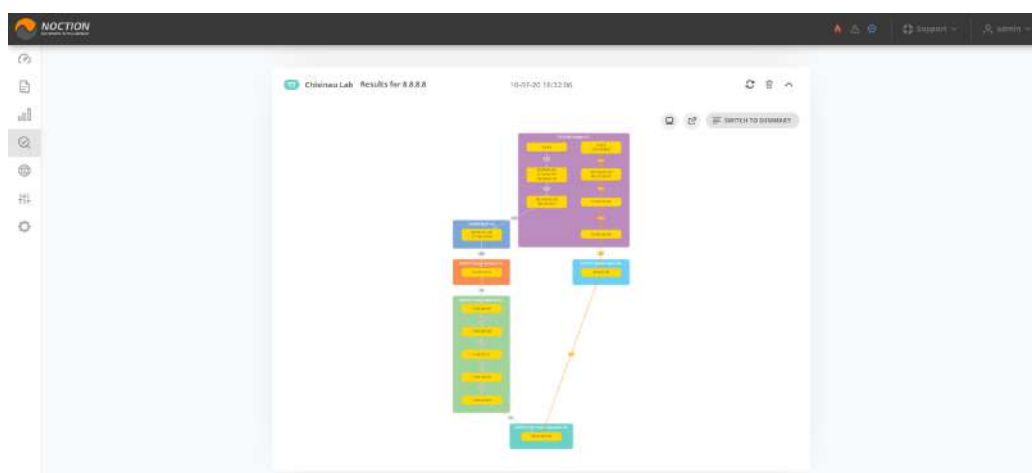


Figure 4.9.4: Traceroute results

TTL	Address	AS	AS Name	RTT Min	RTT Max	RTT Avg	Loss
11	8.8.8.8	13335	Google LLC	41	42	41	0%
10	66.249.24.240	13335	Google LLC	37	41	39	0%
	71.14.236.115	13335	Google LLC				
	216.239.40.225	13335	Google LLC				
9	106.170.251.185	13335	Google LLC	40	42	41	0%
	106.170.252.1	13335	Google LLC				
8	96.104.125.120	8030	W247 Ltd	40	49	43	0%
	217.136.112.94	8030	W247 Ltd				
7	99.121.125.91	8030	Orange Romania S.A.	13	13	13	0%
6	77.88.182.187	25054	Orange Moldova S.A.	1	1	1	0%
5	77.88.185.129	25054	Orange Moldova S.A.	1	27	9	0%
4	77.88.192.17	25054	Orange Moldova S.A.	1	3	1	0%
3	77.88.182.240	25054	Orange Moldova S.A.	1	18	1	0%
2	77.88.223.201	25054	Orange Moldova S.A.	1	1	1	0%
1	185.24.202.126	80208	Forn Trade Corporation Ltd.	1	1	1	0%

Figure 4.9.5: Traceroute exploring details results

4.9.4 Whois

The “Whois” tool can query IP addresses, hostnames, network prefixes, or AS Numbers. The following valid values can be entered in the search box:

- Hostname: fully qualified domain name, e.g: “domain.com”
- IP: valid IP address, e.g: “10.20.30.40”
- IP Block: any valid prefix, in CIDR format, e.g: “10.20.30.40/24”
- ASN: valid ASN, using the following format: “AS1234”

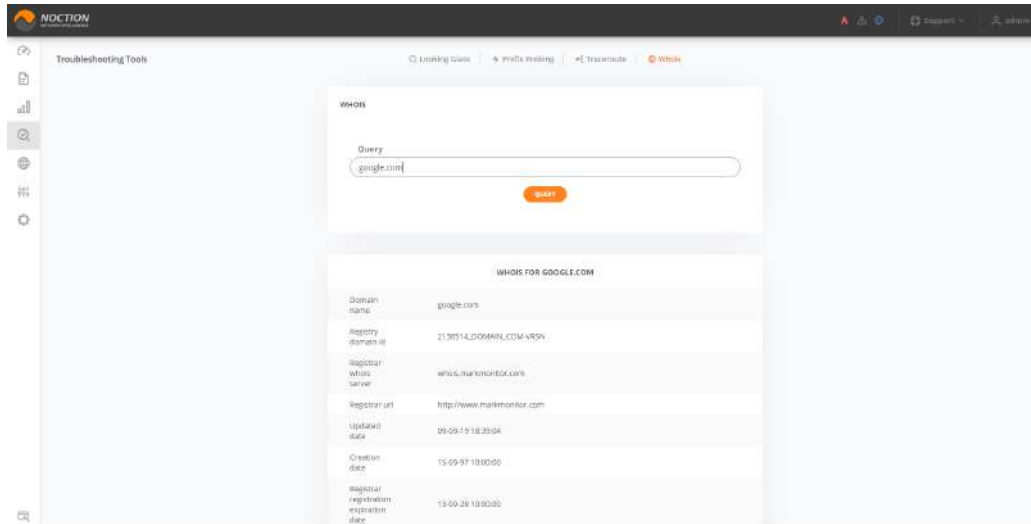


Figure 4.9.6: Whois

4.10 Threat Mitigation

The Threat Mitigation feature can be specifically used to better understand and automatically mitigate the effects of the Distributed Denial of Service (DDoS) attacks. The feature uses standard telemetry and control (NetFlow/sFlow and BGP) capabilities of routers to automatically block disruptive volumetric denial of service attacks. The solution leverages standard features of modern routing hardware to scale easily to large high-traffic networks. It employs a set of threshold-based user-defined rules for potential attack detection as well as the use of FlowSpec and Blackholing mechanisms to mitigate the detected attack.

4.10.1 Monitor

The Monitoring tab offers a 30-minute router(s) traffic graph view, as per specific IRP instance, providing the total Bytes and Packets statistics along with the filter down capability on packet types, including Syn, Fin, UDP, ICMP.

The graph gets updated automatically every minute. To pause the auto-update, click the data point of interest on the graph. This will populate the Historical Data table on the lower left side as well. Click the RESET button to enable the graph auto-update.




Figure 4.10.1: Threat Monitor

4.10.2 Mitigation Rules

All the configured rules are listed under the Mitigation Rules tab.

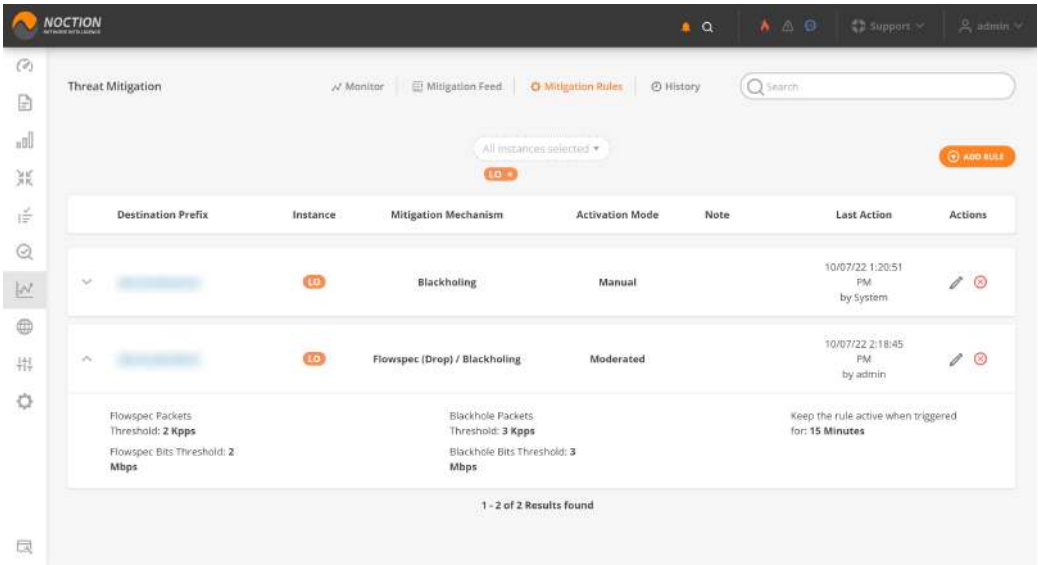


Figure 4.10.2: Mitigation Rules

To add the new rule, click the NEW RULE button. Fill out the required fields and hit CREATE. Based on the chosen Activation mode, a particular threat mitigation action will either be presented as a suggestion under the **Mitigation Feed** tab or started automatically.

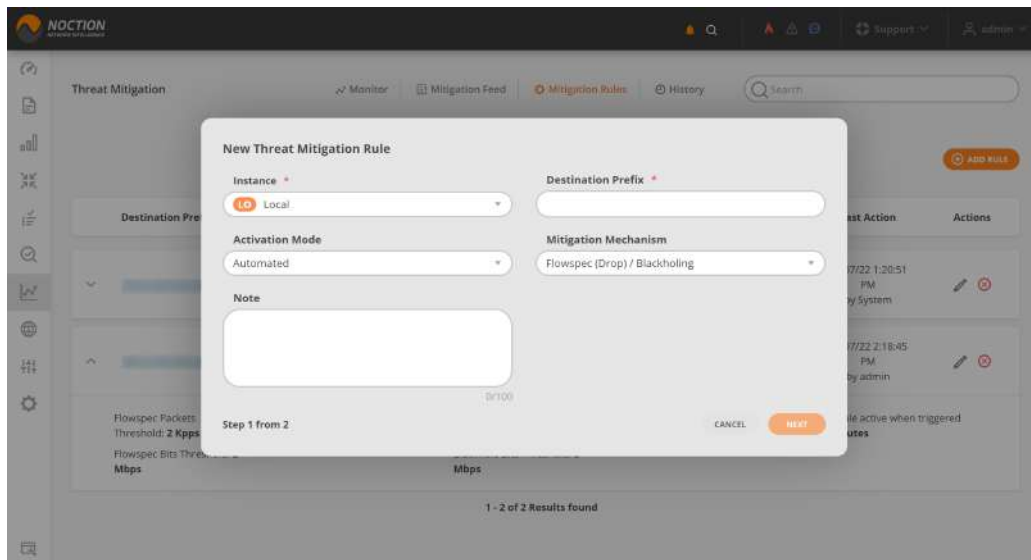


Figure 4.10.3: Add new rule

i To be able to set up a specific Flowspec and/or Blackholing mitigation mechanism, details should be first provided at the Router and Provider level under the Configuration editor section.

In case the individual rule threshold values are not provided, the default values provided in the configuration section get used. Default values for Threat Mitigation rules are set in these configuration parameters: [Threat mitigation](#)

! Deletion of an existing mitigation rule does not affect existing mitigation feed entries. They should be removed separately.

4.10.2.1 Mitigation Mechanisms

BGP Blackholing BGP Blackholing mechanism in Threat Mitigation allows network operators to selectively block or drop traffic to a specific prefix.

Once a rule with the corresponding mechanism gets triggered, the Blackholing community (`peer.X.blackholing.community`) is attached to an announced prefix by the Blackhole Routers (`peer.X.blackholing.bgp_peer`). The purpose here consists of signaling to a provider network that a neighbor router should discard all traffic destined for the received prefix with the attached BGP blackhole community.

It's worth noting that BGP Blackholing should be used with caution, as it will drop all the traffic to the blackholed prefix. This can cause serious consequences, such as dropping legitimate traffic. That's why it's advised to use it carefully and with a clear plan to avoid any collateral damage.

FlowSpec Drop The FlowSpec Drop mechanism allows network administrators to drop packets towards a destination prefix.

Once the FlowSpec Drop Threat Mitigation rule gets triggered, it is propagated to all the configured routers in IRP that support FlowSpec. When a router receives a FlowSpec drop rule, it is installed in its forwarding plane and used to match incoming packets. If a packet matches the rule, it is dropped and not forwarded to its destination.

BGP redirect The BGP redirect mechanism allows network administrators to manually or automatically tag a particular route with a BGP community value and subsequently signal client or provider routers to apply specific actions or policies to such routes.

To redirect traffic to a specific IP address or prefix using the BGP Redirect feature, a network administrator would first define a redirect community(s), then set up the custom threat mitigation rule(s), which, when triggered, would associate the community(s) with a route by adding it as an attribute to the route. Next, depending on the case, the network administrator would configure the router to apply a routing policy that matches the community value and redirects traffic to the desired IP address or readvertises it to the upstream providers.

It's important to note that, the use BGP redirect mechanism requires a good understanding of BGP protocol, network topology and routing policies.

FlowSpec Redirect IRP's Threat Mitigation feature offers the FlowSpec Redirect to IP mechanism, allowing network operators to automatically redirect specific traffic to an IP address for further processing or inspection.

During a DDoS attack, an attacker floods a targeted network or service with a large amount of traffic in an attempt to overload it and make it unavailable to legitimate users. To protect against this type of attack, network administrators can configure the default/custom rules and employ the FlowSpec redirect to IP mechanism to divert traffic that violates predefined thresholds and matches a particular destination prefix to a DDoS mitigation/scrubbing service.

The DDoS mitigation/scrubbing service can then analyze traffic and block or rate-limit the malicious traffic while allowing legitimate traffic to pass through. By redirecting traffic to a DDoS mitigation service using FlowSpec, network administrators can effectively protect their network and services from DDoS attacks without having to block all incoming traffic.

i FlowSpec mitigation method can be chosen to filter out smaller attacks while the Remote Triggered Blackhole should be sent to providers to block large volume attacks. Both Flowspec and Blackholing mechanisms can be enabled for a single rule, with the thresholds for Flowspec being indicated lower than the Blackholing values. In this case, when the FlowSpec Mitigation is not capable of handling the attack the BGP Mitigation comes into play.

4.10.3 Mitigation Feed

The Mitigation Feed represents a list of currently active rules or suggested actions once a particular rule gets triggered. Users can accept the suggestions or simply delete any entry from the list.

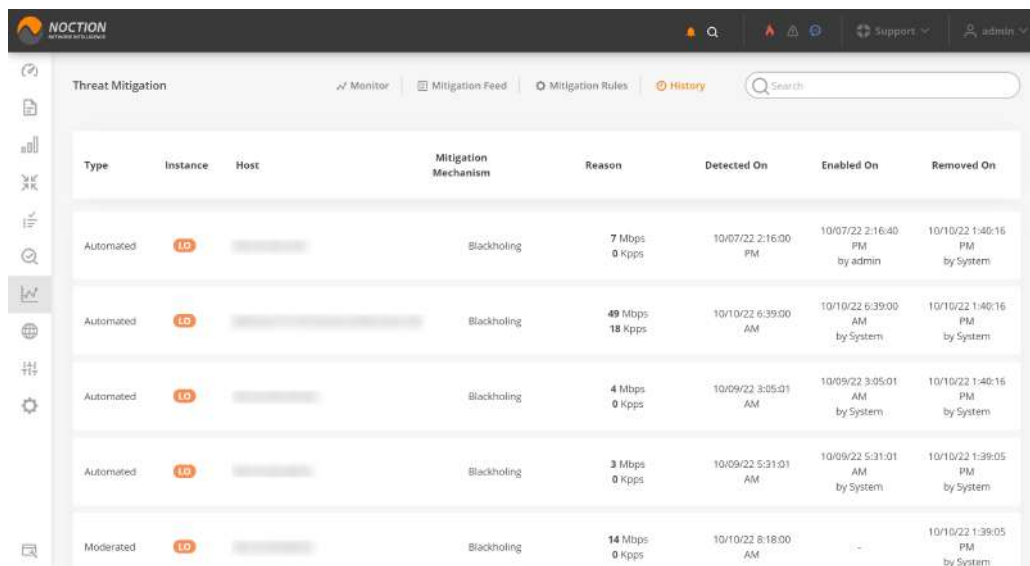
Status	Type	Instance	Host	Mitigation Mechanism	Reason	Detected On	Enabled On	Disabling On	Actions
●	Moderated	10		Blackholing	16 Mbps 0 Kpps	10/11/22 9:06:01 AM	-	10/11/22 9:30:00 AM	ACCEPT
●	Automated	10		Blackholing	9 Mbps 0 Kpps	10/11/22 9:12:00 AM	10/11/22 9:12:00 AM by System	10/12/22 3:12:00 PM	
●	Automated	10		Blackholing	4 Mbps 0 Kpps	10/11/22 9:06:01 AM	10/11/22 9:06:01 AM by System	10/12/22 3:06:01 PM	
●	Manual	10		Blackholing	-	-	10/07/22 1:20:51 PM by System	-	

1 - 4 of 4 Results found

Figure 4.10.4: Mitigation feed

4.10.4 History

The “History” tab highlights the time and the past threat mitigation action details, facilitating the tracing of various issues on the network and in any particular rule configuration per each IRP instance.



Type	Instance	Host	Mitigation Mechanism	Reason	Detected On	Enabled On	Removed On
Automated	LD		Blackholing	7 Mbps 0 Kpps	10/07/22 2:16:00 PM	10/07/22 2:16:40 PM by admin	10/10/22 1:46:16 PM by System
Automated	LD		Blackholing	49 Mbps 18 Kpps	10/10/22 6:39:00 AM	10/10/22 6:39:00 AM by System	10/10/22 1:46:16 PM by System
Automated	LD		Blackholing	4 Mbps 0 Kpps	10/09/22 3:05:01 AM	10/09/22 3:05:01 AM by System	10/10/22 1:46:16 PM by System
Automated	LD		Blackholing	3 Mbps 0 Kpps	10/09/22 5:31:01 AM	10/09/22 5:31:01 AM by System	10/10/22 1:39:05 PM by System
Moderated	LD		Blackholing	14 Mbps 0 Kpps	10/10/22 8:18:00 AM	-	10/10/22 1:39:05 PM by System

Figure 4.10.5: History

4.11 Global Commit

Customers can deploy network configurations with many actual links going to a single ISP from different points of presence. The additional links can serve various purposes such as to provision sufficient capacity in case of very large capacity requirements that cannot be fulfilled over a single link, to interconnect different points of presence on either customer (in a multiple routing domain configuration) or provider sides, or for redundancy purposes. Individually all these links are configured in IRP instances as separate providers. When the customer has an agreement with the ISP that imposes an overall limitation on bandwidth usage, these providers (from various IRP instances) can be grouped together in GMI so that the whole group can be optimized.

Global Commit options can be configured and results visualised by navigating to the corresponding tabs within the Global Commit section.

4.11.1 Global Commit Configuration

Go to the Configuration tab, drag and drop the providers of choice from the right side instances/providers options available to create Global groups. Providers will get distributed and marked by the IRP instance they belong to automatically within the Global group.




Figure 4.11.1: Global Commit groups configuration

Next, provide the “Global Group local member 95th percentile” (95th for the providers from a specific IRP instance within a Global group) as well as “BW reserved for local instance operation” (the amount of summed up bandwidth reserved to providers on specific IRP instance) values.

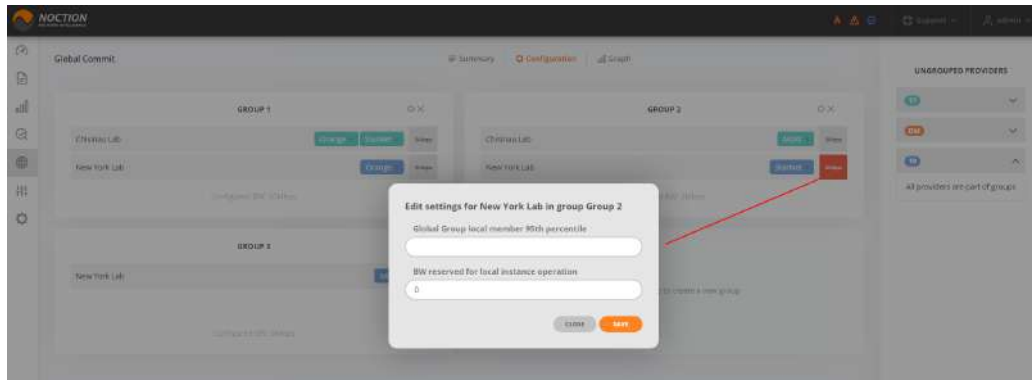


Figure 4.11.2: IRP instance individual settings within the Global Group

Click on the Global Group gear icon to change the group name, add the global group’s members (hosts), Global Group high load bandwidth limit parameter similar in function to (5.8.24) and Global Group low load bandwidth limit parameter similar to (5.8.25)




Figure 4.11.3: Global Commit Group Settings

4.11.2 Global Commit Summary

This tab offers current bandwidth statistics for Global Groups as well as instances and providers within groups. The configured 95th and Reserve Bandwidth values as well as a refresh and graph view link is offered for each group.

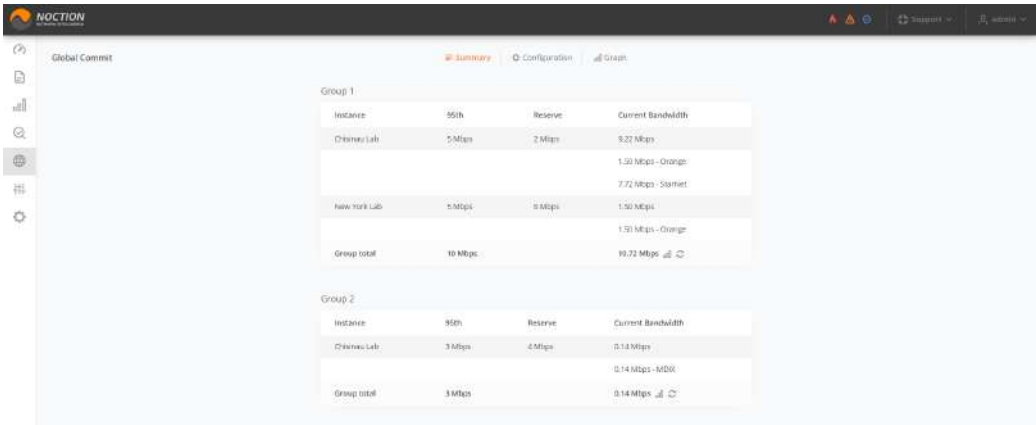


Figure 4.11.4: Global Commit summary

4.11.3 Global Commit Graph

This tab offers graphical representation of the bandwidth usage for a selected group, instances within the selected group as well as individual providers.




Figure 4.11.5: Global Commit Graphs

4.12 Configuration editor

The individual IRP instance system parameters can be modified from the Frontend, using the “Configuration” menu.



Figure 4.12.1: System events

When multiple GMI users are editing specific configuration settings for a given IRP instance, a notification shows up in the top right corner, highlighting their usernames. Such information helps to better coordinate the application of any changes to the configuration settings.




Figure 4.12.2: Concurrent editing of the Configuration Settings

4.12.1 Global configuration

Global settings can be adjusted under the Global tab.

Several global parameters can be configured by selecting their desired values and clicking on the “Save” button.

Available parameters:

- **Improvement mode** (`global.improve_mode`)
- **Management interface** (`global.master_management_interface`)
- **BGP mode** (Intrusive / Non-intrusive, see `global.nonintrusive_bgp`)
- **Probing interface(s)**. See `global.master_probing_interface` and others.

Click “Show advanced” to open a complete list of options. See section 5.2 for details on each parameter type.

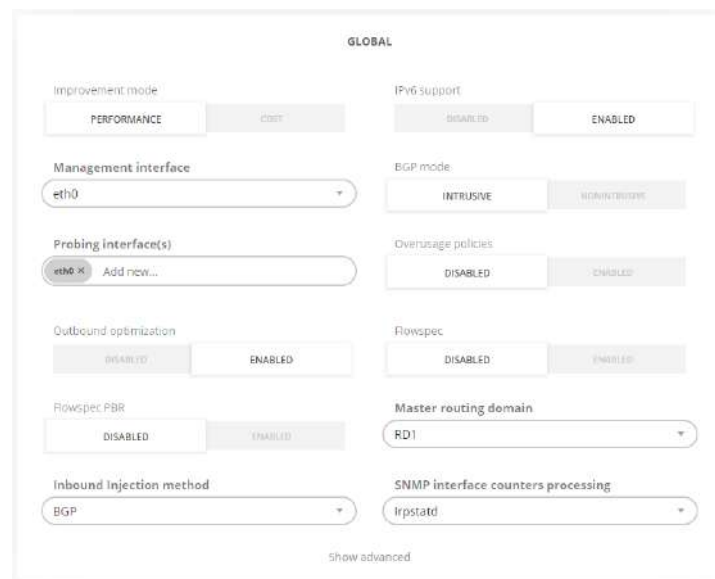


Figure 4.12.3: Configuration editor: Global settings

The list of ignored networks can be modified on the same page. The options allow listing as appropriate:

- prefixes(`global.ignorednets`),
- ASNs (`global.ignored.asn`),
- BGP Community attributes that mark prefixes/routes to ignore (`global.ignored_communities`).

Setting prefixes and ASNs is possible either manually, or by importing a text file containing one IP/network in CIDR format or ASN per line.




Figure 4.12.4: Configuration editor: Ignored prefixes

The Routing Domains can be configured from the Global tab as well.




Figure 4.12.5: Configuration editor: Routing Domains

4.12.2 BGP and Routers configuration

Specific IRP instance BGP daemon-related parameters can be configured under the “Configuration→BGP” tab.

- Bgpd parameters:
 - BGP mode intrusive/non-intrusive (`global.nonintrusive_bgp`)
 - BGP monitor enabled/disabled, this option is useful during DoS/DDoS attacks. Bgpd ICMP / SNMP monitors can be disabled on the fly, without restarting the Bgpd daemon.
 - AS-Path attribute a restore priority (`bgpd.as_path`)
 - Remove prefixes on next-hop update (`bgpd.improvements.remove.next_hop_eq`)
 - Remove prefixes on aggregate withdraw (`bgpd.improvements.remove.withdrawn`)
- BGP monitoring settings:
 - Guard time (`bgpd.mon.guardtime`)
 - Holdtime (`bgpd.mon.holdtime`)
 - Keepalive (`bgpd.mon.keepalive`)

The screenshot shows the BGP configuration editor. At the top, there are several toggle switches for BGP settings:

- 300 announcements to preserve original route attributes:** Disabled (toggle to Enabled).
- Remove next hop update:** Disabled (toggle to Enabled).
- Strip non-IP communities:** Disabled (toggle to Enabled).
- Remove an aggregate withdrawal:** Disabled (toggle to Enabled).
- Re-probe on new path change:** Disabled (dropdown menu).
- Re-probe on old path change:** Disabled (dropdown menu).

 Below these is a "Show advanced" link. The main section shows two tabs: "R1" and "R1_V6". The "R1" tab is active and shows:

- Local IPv4 inbound next hop:** Input field.
- Announced inbound localpref value:** Input field with value 102.
- Transit SNMP hosts:** Input field with "Add new..." button.
- Transiting traffic:** Disabled (toggle to Enabled).

 The "R1_V6" tab is also visible, showing:

- Autonomous System:** Input field with value 60230.
- Improvement localpref:** Input field.
- IRP's IPv4 address:** Input field.
- Router IPv4 address:** Input field.

Figure 4.12.6: Configuration editor: BGP settings

BGP routers can be added, removed, activated, shutdown or modified on the same page.

A new BGP neighbor can be added using the “Add BGP Peer” button. Provide a name and configure the following parameters:

- General settings:
 - ASN to be used for the iBGP session (`bgpd.peer.X.as`)
 - Announced local-pref value (`bgpd.peer.X.master_localpref`)
 - Local IP address (`bgpd.peer.X.master_our_ip`, `bgpd.peer.X.master_our_ipv6`)
 - Neighbor IP address, usually the router’s IP (`bgpd.peer.X.master_peer_ip`, `bgpd.peer.X.master_peer_ipv6`)

The screenshot shows the configuration editor for a new router, "R1". The "General" tab is active, showing:

- Autonomous System:** Input field with value 60230.
- Improvement localpref:** Input field with value 190.
- IRP's IPv4 address:** Input field with value 185.34.202.184.
- Router IPv4 address:** Input field with value 185.34.202.246.

 The "Advanced" tab is also visible. The router status is "ACTIVE".

Figure 4.12.7: Configuration editor: New Router, general settings

- Advanced settings:
 - BGP session password
 - Improvement communities to be appended to the Communities attribute (`bgpd.peer.X.master_communities`)
 - Keepalive (`bgpd.peer.X.keepalive`)
 - Announced MED value (`bgpd.peer.X.med`)
 - BGP Router ID, mandatory for IPv6 only (`bgpd.peer.X.slave_router_id`)

Figure 4.12.8 shows the 'Advanced' settings tab for a new router configuration. The interface includes tabs for 'General', 'Advanced', 'Inbound', and 'Blackholing'. The 'Advanced' tab is active, displaying fields for 'BGP session password', 'Keepalive interval (sec)' (set to 60), 'BGP Router ID', 'Flowspec status' (set to DISABLED), 'Improvement communities' (with an 'Add new...' button), and 'Improvement MED value' (with a red note 'should be integer').

Figure 4.12.8: Configuration editor: New Router, advanced settings

- Inbound settings.
 - Local inbound next_hop (5.5.8)
 - Announced inbound LocalPref value (5.5.10)
 - Transit SNMP hosts (5.5.34)
 - Transiting traffic toggle (5.2.28)

Figure 4.12.9 shows the 'Inbound' settings tab for a new router configuration. The interface includes tabs for 'General', 'Advanced', 'Inbound', and 'Blackholing'. The 'Inbound' tab is active, displaying fields for 'Local IPv4 inbound next_hop', 'Announced inbound localpref value' (set to 102), 'Transit SNMP hosts' (with an 'Add new...' button), and 'Transiting traffic' (set to DISABLED).

Figure 4.12.9: Configuration editor: New Router, Inbound settings

- Blackholing.
 - Announced blackholing localpref value
 - Blackholing IPv4 next hop
 - Blackholing IPv6 next hop

Figure 4.12.10 shows the 'Blackholing' settings tab for a new router configuration. The interface includes tabs for 'General', 'Advanced', 'Inbound', and 'Blackholing'. The 'Blackholing' tab is active, displaying fields for 'Announced blackholing localpref value' (set to 100), 'Blackholing IPv4 next hop', and 'Blackholing IPv6 next hop'.

Figure 4.12.10: Blackholing configuration settings

4.12.3 Collector configuration

Global collector parameters, as well as Span and Flow-specific settings can be adjusted under “Configuration→Collector” tab.

- Global collector settings:
 - Top volume prefixes per cycle (`collector.export.volume.high.top_n`)
 - Minimal traffic volume (`collector.export.volume.min`)
 - Minimal traffic volume (%) (`collector.export.volume.min_pct`)
- Flow collector settings:
 - UDP port to listen for NetFlow or sFlow packets (`collector.flow.listen.nf`, `collector.flow.listen.sf`)
 - Flow sources (`collector.flow.sources`)
- SPAN collector settings:
 - Capture interfaces (`collector.span.interfaces`)
 - Mindelay algorithm enable/disable (`collector.span.min_delay`)
 - Mindelay probing queue slots (`collector.span.min_delay.probing_queue_size`)

Figure 4.12.11: Configuration editor: Collector settings

The list of prefixes announced by the monitored network can be edited in the same form. Its contents can be also imported from a text file.

See also: [collector.ournets](#).

- Inbound Prefixes collector settings:
 - prefix belonging to the network (5.15.6),
 - router(s) where announcements are sent (5.15.1),
 - next hop used by improvements for this prefix (5.15.5)
 - an option to instruct IRP whether to fully control the prefix or to only announce improvements when there are any.
 - providers for which this inbound prefix can be optimized (5.15.7). It must be noted that if no provider is selected then the prefix is optimized for all providers. Otherwise, the prefix will be optimized only through selected providers.
 - prefix status as each prefix can be disabled in order to exclude from further inbound optimization (5.15.3)

The screenshot shows the 'INBOUND PREFIXES' configuration editor. On the left, there is a search bar with two results: '185.34.202.0/24' and '2a04:5ec7::/48'. The main form on the right contains the following fields: 'Prefix' (185.34.202.0/24), 'Router' (R1), 'Inbound Prefix Next Hop' (empty), 'Control' (If improved), 'Providers' (Orange, startnet), and 'Prefix status' (DISABLED, ENABLED).

Figure 4.12.12: Configuration editor: Collector Inbound Prefixes settings

4.12.4 Core configuration

All Core parameters affecting the decision-making algorithms can be modified on the next tab (“Configuration→Core”). The following configuration parameters can be adjusted:

- **Max IPv4 improvements** (`core.improvements.max`, `core.improvements.max_ipv6`)
- **Standard reprobing period** (`core.improvements.ttl.retry_probe`) - the time period after which a specific improvement is being scheduled for re-probing.
- **VIP reprobing period** (`core.vip.interval.probe`) - defines the VIP prefixes probing interval, in seconds.
- **Minimal probe lifetime** (`core.probes.ttl.min`) - Ordinary probing will not be performed for a specific prefix if its probe age is lower than this value.
- **Allowed latency worsening** (`core.cost.worst_ms`)
- **Exploring queue slots** (`core.eventqueue.limit`)
- **Relevant loss** for loss-based improvements (`core.performance.loss_pct`) - a prefix will be improved based on a loss cause only if the packet loss can be decreased by this value (in %)
- **Relevant RTT difference** (`core.performance.rtt.diff_ms`)
- **Relevant RTT difference (%)** (`core.performance.rtt.diff_pct`)
- **Maximum probe lifetime** (`core.probes.ttl.max`)

Figure 4.12.13 shows the 'CORE' configuration page. It contains the following fields:

- Max IPv4 Improvements: 10000
- Max IPv6 Improvements: 1000
- Standard reprobing period (sec): 14400
- Minimal probe lifetime (sec): 7200
- VIP reprobing period (sec): 3600
- Max IPv4 Flowspec rules: 100
- Max IPv6 Flowspec rules: 100
- A 'Show advanced' link at the bottom.

Figure 4.12.13: Configuration editor: Core configuration

Outage detection algorithm can be enabled and configured on the same page. Outage detection algorithm can be enabled or disabled using the toggle buttons at the top right.

Figure 4.12.14 shows the 'OUTAGE DETECTION' configuration page. It includes:

- Toggle buttons: 'DISABLED' (selected) and 'ENABLED'.
- Outage prefix confirmation rate: Slider set to 90 (range 0 to 100).
- Outage round trip rate (ms): Slider set to 90 (range 0 to 100).
- Outage confirmation timeout (sec): Text input field with value 600.

Figure 4.12.14: Configuration editor: Outage detection

Throttling with Flowspec can be configured on the same page.

Figure 4.12.15 shows the 'OVERUSAGE' configuration page. It includes:

- Overusage interval (sec): Text input field with value 60.
- Overusage rule retention (sec): Text input field with value 60.
- Prefix BW average time (hours): Slider set to 1 (range 1 to 24).
- Prefix relevant BW (Mbps): Text input field with value 50.
- Overusage throttle multiplier: Text input field with value 2.
- Overusage threshold multiplier: Text input field with value 10.

Figure 4.12.15: Configuration editor: Overusage

The following configuration parameters can be adjusted:

- **Overusage interval** (`core.overusage.check_interval`) sets the frequency in seconds of checking for excessive bandwidth use.
- **Overusage rule retention** (`core.overusage.hold_timer`) sets how much time a rule is kept after bandwidth use returns to normal.
- **Prefix BW average time** (`core.overusage.out.average.period`) sets the number of hours used to determine the average bandwidth use of a prefix.

- **Prefix relevant BW** (`core.overusage.out.average.relevant_min`) sets the relevant bandwidth use in Mbps by a prefix before considering any rules for it.
- **Overusage throttle multiplier** (`core.overusage.out.threshold.throttle`) sets the multiplier to apply to average prefix bandwidth use when setting a rule.
- **Overusage threshold multiplier** (`core.overusage.out.threshold.trigger`) sets the threshold multiplier used to determine excessive bandwidth use.

➖ Prefix relevant BW and Overusage threshold multiplier parameters control the number of Flowspec rules that will be generated automatically. Adjust these values to control the number of Throttling Flowspec policies.

Circuit Issues Detection parameters are listed and can be adjusted:

Figure 4.12.16: Configuration editor: Circuit issues detection

The following configuration parameters can be adjusted:

- **Delta loss to shutdown** (`core.circuit.high_loss_diff`) sets threshold to initiate shutting down a provider with circuit issues as a difference between examined provider's average loss and overall average loss during the given time horizon. Shutdown is attempted only for providers marked accordingly.
- **Delta loss to warn** (`core.circuit.warn_loss_diff`) sets threshold to raise a warning regarding issues with a provider as difference between examined provider's average loss and overall average loss during the given time horizon. Usually this threshold is significantly lower than the shutdown threshold.
- **Delta loss to restore** (`core.circuit.recover_loss_diff`) sets low loss level threshold when IRP will restore full functionality over provider that had circuit issues.
- **Issues time horizon** (`core.circuit.hist_interval`) sets how many minutes in the past IRP looks for probes with loss over both analyzed provider and all the other providers on the network.
- **Withdraw improvements on warn** (`core.circuit.withdraw_on_warn`) instructs IRP to withdraw outbound improvements over provider with circuit issues when warning threshold is reached. By default improvements are withdrawn only when shutdown threshold is exceeded.
- **Restore after** (`core.circuit.recover_hold_time`) sets the interval in seconds after which IRP should re-evaluate a provider's circuit loss and attempt restoring it to normal function. This will be performed only for providers that are configured to attempt to restore after shutdown.

- **Restore interval** (`core.circuit.recover_monitored_intervals`) sets the interval in minutes during which IRP will periodically re-evaluate a provider's circuit loss and attempt restoring it to normal function. This will be performed only for providers that are configured to attempt to restore after shutdown.

❌ The configuration parameters above are applied for all providers even though for individual providers different level of control can be set starting from disabling and ending with attempting both to automatically shutdown and later restore connectivity with it. See also [Providers configuration](#)

4.12.5 Commit Control configuration

Commit Control algorithm can be enabled and configured on the Commit Control group of pages. It can be enabled or disabled using the toggle On/Off buttons at the top of the form.

⚠️ If NetFlow is used to collect statistics Routers **MUST** be configured to export statistics every minute (or as often as possible). Some router models have default export intervals for either inactive or active flows of up to 1800 seconds. Big delays cause IRP to react very slowly to increased load and reduce the effectiveness of Commit Control feature.

The most important parameters for Commit Control are, as follows:

- **Commit Control probing queue slots** (`core.commit_control.probing_queue_size`)
- **Minimal prefix bandwidth** in Mbps (`core.commit_control.agg_bw_min`)

Figure 4.12.17: Configuration editor: Commit Control

Providers Overall block displays Commit Control configuration for all providers including their precedence. See details in the figure below.

Individual provider parameters can be adjusted as well as detail regarding current Commit Control changes applied by an IRP instance can be accessed by following the provided links.

PROVIDERS OVERALL											
Provider	Status	95th Mode	Configured 95th for inbound	Group Configured 95th for inbound	Configured 95th	Group Configured 95th	Current BW Usage	Monthly 95th	Cost	Load Balancing	Precedence
Orange	CC Enabled	95th from greater of in, out	3	-	5	-	0.00 Mbps	0.00 Mbps	2	Enabled	2
StarNet	CC Enabled	95th from greater of in, out	6	-	10	-	0.00 Mbps	0.00 Mbps	4	Enabled	3
MDIX	CC Enabled	95th from greater of in, out	-	-	100	-	0.00 Mbps	0.00 Mbps	-	Enabled	10

Figure 4.12.18: Providers Overall

4.12.5.1 Inbound Commit Control

Inbound bandwidth control is part of the Commit Control feature. You need to enable Commit Control and further to enable Inbound Commit control to make use of this feature.

Figure 4.12.19: Configuration editor: Inbound commit control

Commit control for inbound highlights the relevant configuration parameters:

- whether inbound commit control is on or off
- whether review and moderation feature is enabled or disabled
- what is the bandwidth estimation algorithm for inbound and others

Refer also to `peer.X.95th.mode`, `core.commit_control.inbound.enabled`, `core.commit_control.inbound.moderated`.

Besides the above mentioned settings for Inbound Commit Control separate low and high limits are configured to instruct IRP when to start and when to stop improving and also how to estimate the effects of inbound improvements.

Estimating inbound traffic Internet traffic has high variability and adjacent measurements can differ significantly between them. In these conditions a leveling function helps mitigate the risk of generating excessive numbers of Inbound improvements due to higher traffic variability of some networks. This estimation can be fine tuned by the Inbound bandwidth estimation algorithm parameter that tells an IRP instance what value to use as the basis of the calculation. See the possible base values below:

Figure 4.12.20: Configuration editor: Inbound BW estimation algorithms

4.12.5.2 Optimization of transiting traffic

Optimization of transiting traffic is setup as part of [Inbound Commit Control](#).

The screenshot shows the 'INBOUND COMMIT CONTROL' configuration interface. At the top right, there are 'DISABLED' and 'ENABLED' toggle buttons, with 'ENABLED' being selected. The interface is divided into two columns of settings:

- Left Column:**
 - 'Commit Control moderated inbound improvements' with 'DISABLED' and 'ENABLED' toggles; 'ENABLED' is selected.
 - 'Transiting traffic' with 'DISABLED' and 'ENABLED' toggles; 'ENABLED' is selected.
 - 'Transit matching at egress' with 'DISABLED' and 'ENABLED' toggles; 'ENABLED' is selected.
 - 'Transit ASNs' with an 'Add new...' button and an upward arrow icon.
 - 'Transit Improvement TTL min (sec)' with a text input field containing '14400'.
 - 'Transit top N prefixes' with a text input field containing '100'.
 - 'Provider's high load inbound bandwidth limit (%)' with a slider ranging from 50 to 99, currently set at 90.
- Right Column:**
 - 'Inbound bandwidth estimation algorithm' with a dropdown menu showing 'Greatest of last minute and current hour average'.
 - 'Transit Improvements Max' with a text input field containing '100'.
 - 'Transit prefixes' with a text input field containing '185.34.203.0/24', an 'Add new...' button, and an upward arrow icon.
 - 'Transit Improvement TTL max (sec)' with a text input field containing '86400'.
 - 'Provider's low load inbound bandwidth limit (%)' with a slider ranging from 30 to 99, currently set at 70.

Figure 4.12.21: Configuration editor: Inbound optimization of transiting traffic parameters

The following configuration parameters can be adjusted:

- Transiting traffic toggle that enables or disables the feature ([global.inbound_transit](#))
- Transit Improvements Max sets the maximum number of transit improvements ([core.improvements.inbound_transit.](#)
- Transit ASNs and Transit prefixes specify those segments of the Internet that IRP monitors and optimizes ([bgpd.prefixlist.asn](#), [bgpd.prefixlist.prefixes](#))
- Transit Improvements TTL min and max set the lower and upper bounds in seconds to keep a transit improvement ([core.improvements.inbound_transit.ttl.min](#), [core.improvements.inbound_transit.ttl.max](#))

i Note that traffic belonging to a specific prefix that fragments a large transit prefix is collected as belonging only to the specific prefix and excluded from the larger prefix traffic statistics. If the specific prefix is also filtered from IRP configuration as for example is the case of a /26 prefix, then the specific prefix will be excluded from optimization of transiting traffic and will not show up in any of the subsequent decisions or reports.

- Transit Top N prefixes sets the number of transit prefixes that are collected each cycle and considered for optimization ([collector.flow.export.inbound_transit.topn](#))
- Match transit at egress enables or disable statistics collection for transit prefixes when packets exist the network ([collector.flow.process_transit_in_outbound](#))

4.12.6 Inbound performance optimization configuration

Refer to [Inbound performance optimization](#) for the Inbound Performance feature overview.

To adjust the Inbound performance optimization parameters, go to “Configuration→Inbound”.

The following configuration parameters can be adjusted:

- **Inbound performance (`irpinperfd.enabled`)** enables/disables inbound performance optimization. This also disables the traffic statistics collection, required for inbound performance optimization
- **Inbound Performance moderated mode (`irpinperfd.moderated`)** enables/disables manual moderation of all decisions of the inbound performance optimization algorithm
- **Probe shelf life (`irpinperfd.probing.shelf_life`)** sets the maximum age of a probe which can be reused
- **Probing failure margin (`irpinperfd.probing.failure_margin`)** sets the percentage of probing failures (maximum possible value is 25%), exceeding which is critical and leads to the inbound performance optimization algorithm restart
- **Probing timeout (`irpinperfd.probing.timeout`)** sets the time period of probing process, exceeding which is critical and leads the inbound performance optimization algorithm restart
- **Reprobing interval (`irpinperfd.probing.interval`)** sets the period of time to wait until the algorithm starts reprobing a sample and comparing it against the model
- **Top volume prefixes per inperf rule (`irpinperfd.model.topn_per_rule`)** sets the number of prefixes to be considered for the rule’s model construction. Prefixes with greater volume are taken into account first




Figure 4.12.22: Configuration editor: Inbound configuration settings

The Inbound Performance feature makes it’s decisions based on the configured rules.

i When the Moderated mode is enabled, the Inbound Performance improvements should be applied manually at the Inbound Performance Rules page.

Refer to [Inbound Performance Rules](#), [Inbound Performance Optimization](#) and [Inbound Performance](#) .

4.12.7 Explorer configuration

To adjust the Explorer-related parameters, go to “Configuration→Explorer”. The [Explorer settings](#) sections should be consulted for detailed Explorer parameters description.

Explorer parameters:

- Infrastructure IPs (`explorer.infra_ips`)
- Explorer worker threads (`explorer.maxthreads`)
- Probing algorithms and Traceroute algorithms (`explorer.probe.algorithm`, `explorer.trace.algorithms`)

- High volume task precedence ([explorer.high_vol_precedence](#))
- Process max collected IPs ([explorer.max_collector_ips](#))
- First probing packets amount ([explorer.probing.sendpkts.min](#))
- Adaptive probing packets count ([explorer.probing.sendpkts.adaptive_max](#))
- ICMP timeout ([explorer.timeout](#))
- Traceroute retry packets and packets per hop ([explorer.traceroute.sendpkts](#), [explorer.traceroute.retrypkts](#))
- Traceroute minimum and maximum ttl ([explorer.traceroute.ttl.min](#), [explorer.traceroute.ttl.max](#))

Figure 4.12.23: Configuration editor: Explorer settings

4.12.8 Providers and Peers configuration

Providers can be added and modified via “Configuration→Providers”.

See also: [Provider](#)

Figure 4.12.24: Configuration editor: Providers configuration

The provider of choice widget should be expanded to be able to edit it. Using the control buttons, a specific provider can be temporarily suspended (e.g for a short maintenance in the monitored network), shutdown (for long maintenance), or completely deleted. See also: [peer.X.shutdown](#). Commit Control can also be disabled or enabled for each provider specifically. ([peer.X.cc_disable](#)).

i The sections below describe the most common groups of provider configuration parameters. Some parameters are repeated on more than one of these groups for convenience.

4.12.8.1 Providers configuration

To add a new provider, the “Add provider” button must be used. Several parameters should be configured for the new provider:

- General settings:
 - Provider name (`peer.X.shortname`)
 - Router (`peer.X.bgp_peer`)
 - Provider description (`peer.X.description`)
 - Provider’s routing domain (`peer.X.rd`)
 - Maximum load per interface (`peer.X.limit_load`)
 - Flow agents (`peer.X.flow_agents`)
 - Circuit issues detection (`peer.X.circuit.control`) indicating how IRP should react to excessive persistent loss over a particular provider.
 - Use of BMP data (5.6)

The screenshot shows the 'ORANGE' configuration editor with the 'General' tab selected. The form contains the following fields and controls:

- Provider/link name:** Text input with 'Orange'.
- Router:** Buttons for 'RT', 'RTv6', and 'Add new...'.
- Provider description:** Text input with 'Orange Moldova'.
- Provider's routing domain:** Dropdown menu with 'RD1' selected.
- Maximum load per interface (Mbps):** Text input with '800'.
- Flow agents:** Text input with '185.34.202.233/562' and an 'Add new...' button.
- Circuit issues detection:** Dropdown menu with 'Disabled' selected.
- Use BMP data:** Dropdown menu with 'Do not use BMP data' selected.
- Routes configuration mode:** Dropdown menu with 'Autoconfiguration from BGP' selected.
- PBR check for provider:** Two buttons: 'DISABLED' (active) and 'ENABLED'.

Figure 4.12.25: Configuration editor: Adding a new provider, General settings

- IPv4 / IPv6 settings:
 - IPv4 / IPv6 diagnostic hops (`peer.X.ipv4.diag_hop`, `peer.X.ipv6.diag_hop`)
 - Probing IPv4 / IPv6 address (`peer.X.ipv4.master_probing`, `peer.X.ipv6.master_probing`)
 - Remote provider ASN (`peer.X.ipv4.next_hop_as`), used by the AS-Path restoration algorithm (`bgpd.as_path`)
 - Router next-hop address (`peer.X.ipv4.next_hop`), that sets the next-hop value for injected routes related to this provider
 - Provider’s alternative IPv4 address for PBR tests (5.14.35)

Figure 4.12.26: Configuration editor: Adding a new provider, IPv4 / IPv6 settings

- Commit Control configuration:
 - Provider cost per Mbps (`peer.X.cost`)
 - Provider 95th percentile (`peer.X.95th`)
 - Provider inbound 95th percentile (`peer.X.95th.in`)
 - Provider 95th calculation mode for inbound traffic
 - Provider billing day (`peer.X.95th.bill_day`)
 - Commit Control status for this provider (`peer.X.cc_disable`)
 - CC provider precedence - used for Commit Control and grouping (`peer.X.precedence`)
 - Centile value (`peer.X.95th.centile`)
 - Performance/Cost improvements within provider group toggle in case the peer load balancing is configured, (`peer.X.improve_in_group`, `peer.X.precedence`)

Figure 4.12.27: Configuration editor: Adding a new provider, Commit Control

- SNMP-related settings:

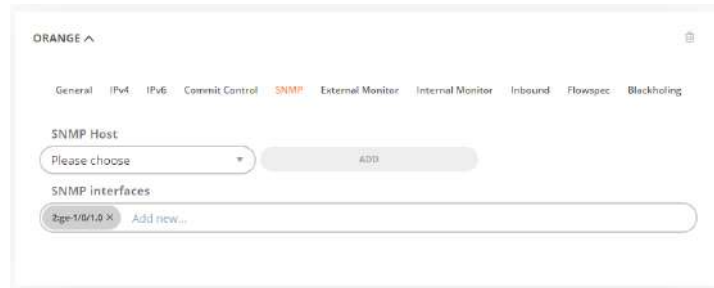


Figure 4.12.28: Configuration editor: Adding a new provider, SNMP settings

- SNMP v3 uses additional parameters depending on security services used for statistics collection:
 - SNMP security services selects if Authentication and/or Privacy services are used (`snmp.X.seclevel`)
 - SNMP authentication password if authentication is used (`snmp.X.auth_password`)
 - SNMP authentication protocol if authentication is used (`snmp.X.auth_protocol`)
 - SNMP encryption password if privacy is used (`snmp.X.priv_password`)
 - SNMP encryption protocol if privacy is used (`snmp.X.priv_protocol`)
 - SNMP Username if authentication is used (`snmp.X.auth_username`)
 - SNMP version (`snmp.X.version`)
- External Monitor settings:
 - External monitor status toggles for IPv4 / IPv6 (`peer.X.mon.ipv4.external.state`, `peer.X.mon.ipv6.external.state`)
 - ICMP/UDP ping monitored IPv4 / IPv6 addresses (`peer.X.ipv4.mon`, `peer.X.ipv6.mon`)

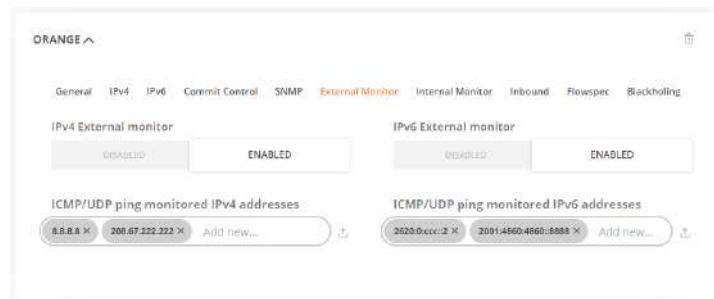


Figure 4.12.29: Configuration editor: Adding a new provider, External Monitor

- Internal Monitor settings:
 - Internal monitor status toggles for IPv4 / IPv6 (`peer.X.mon.ipv4.internal.state`, `peer.X.mon.ipv6.internal.state`)
 - BGP session monitoring IPv4 / IPv6 addresses (`peer.X.mon.ipv4.bgp_peer`, `peer.X.mon.ipv6.bgp_peer`)
 - BGP MIBs for IPv4 / IPv6 (`peer.X.mon.ipv4.internal.mode`, `peer.X.mon.ipv6.internal.mode`)
 - SNMP host for BGP Internal Monitor (`peer.X.mon.snmp`)

Figure 4.12.30: Configuration editor: Adding a new provider, Internal Monitor

- SNMP v3 uses additional parameters depending on security services used for monitoring:
- Inbound:
 - Base community for inbound improvements (5.14.27)

Inbound optimization improvements advertise larger or smaller counts of preponds to be announced to different providers. IRP instances use BGP session to communicate to routers and relies on BGP communities to communicate the improvements. Each provider is assigned a base community which represents zero preponds. Additional preponds are represented by incrementing accordingly the right-side part of the provider's community.

⚠ IRP instances use 8 prepond levels. This means that base community should be chosen to allow additional 7 values without intersecting with other community values.

Figure 4.12.31: Configuration editor: Adding a new provider, Inbound

- Flowspec:
 - IPv4 / IPv6 Redirect community (5.14.20 and 5.14.21)

Figure 4.12.32: Configuration editor: Adding a new provider, Flowspec

- Blackholing:
 - Blackholing Routers
 - Blackholing community




Figure 4.12.33: Provider Blackholing configuration details

The same parameters can be adjusted for the already configured providers.

4.12.8.2 Internet Exchanges configuration

An Exchange is configured as a special type of provider. The important things to consider when setting up an Exchange are listed below.

i It is recommended that Noction systems engineers guide you through Exchange configuration process.

- IRP needs one ACL and one PBR rule for each Peering Partner on an Exchange. Internet Exchanges can have hundreds of Peering Partners. Ensure that the number of Access Lists, Access List entries and PBR entries will not exceed router hardware limitations.

- Once the Exchange is setup via GMI, an IRP instance will retrieve the routing table on the router in order to setup Exchange peering partners. IRP will require access to the router in order to do so.

- When configuring an Exchange its Diag Hop parameter must be provided. Diag Hop for an Exchange represents the list of direct-connected networks configured on the router's interface towards the Exchange network. This parameter is labeled "IPv4 diagnostic hop" on IRP Frontend.

- Probing IPs for Exchanges no longer require one IP address per peering partner since this number might be impractically big. Instead a combination of probing IP and DSCP value is used to configure the PBRs. A single probing IP in conjunction with DSCP can cover up to 64 peering partners. A sufficient number of probing IPs will be required in order to cover all peering partners on your exchange.

⚠ As a rule of thumb it is better to list the PBR rules for transit providers before the (large number of) rules for Internet Exchange peers. If the many PBR rules assigned to peers on a large Internet Exchange are placed at the beginning of the PBR list some routers evaluate all of them before finding the terms for transit providers and this consumes valuable router CPU resources.

MDIX

ACTIVE SUSPENDED SHUTDOWN

General IPv4 IPv6 Commit Control SNMP External Monitor Internal Monitor Peering Partners Inbound Flowspec

Provider/link name: MDIX

Router: R1 X R1_v6 X Add new...

Provider description: MDIX

Provider's routing domain: RD1

Maximum load per interface (Mbps): 1000

Flow agents: 185.34.202.253/583 X Add new...

Circuit issues detection: Disabled

Use BMP data: Do not use BMP data

Figure 4.12.34: An Internet Exchange is similar to a provider and it requires configuration of the Router, the Probing IPs and the other usual attributes

- Once the Exchange is configured, the IRP instance will need to reset its BGP session towards the router interconnecting with the Exchange in order to retrieve the required routing table information about all peering partners. Once the BGP session is restarted IRP instance will start fetching Exchange routing tables.

i Allow sufficient time (~10 minutes) to fetch Exchange routing tables before continuing configuration.

- Peering partner autoconfiguration function is available. It retrieves the list of Next Hops (peering partners) on the Exchange with the corresponding list of prefixes individual peers accept.

MDIX

General IPv4 IPv6 Commit Control SNMP External Monitor Internal Monitor Peering Partners Inbound Flowspec Blackholing

Search...

193.17.78.10 X

193.17.78.11 X

193.17.78.12 X

193.17.78.13 X

193.17.78.17 X

193.17.78.20 X

193.17.78.22 X

193.17.78.33 X

193.17.78.35 X

193.17.78.4 X

193.17.78.7 X

193.17.78.8 X

193.17.78.9 X

2001:7f8:6a::1:3335:13 X

2001:7f8:6a::1:5836:4 X

2001:7f8:6a::2:5454:10 X

AUTOCONFIG PBR CHECK

Peering Partner next-hop: 193.17.78.10

Peering Partner name:

Route Server or Peering session IPv4/IPv6 address: 193.17.78.2

Probing IP address: 185.34.202.187

Probing DSCP: 0

Peering Partner status: DISABLED ENABLED

Peering Partner PBR check: DISABLED ENABLED

Figure 4.12.35: The list of peering partners for the Exchange shows details about each of them and offers access to Autoconfiguration feature and PBR ruleset generator for the router

- Use the Autoconfiguration feature to create an initial configuration for an IRP instance. Review Exchange peering partners before starting the Exchange. Consider enabling periodic auto re-configurations for selected IX in order to update periodically the value of BGP session monitoring IPv4 address from the Exchange and to add new peering partners. Refer [global.exchanges.auto_config_interval](#) and [peer.X.auto_config](#).

i Keep in mind that Autoconfiguration might tamper with all the changes you've made directly within IRP instance config files for peering partners. So, use Autoconfiguration sparingly after you've applied manual changes or try avoiding direct configuration file changes altogether.

- Once the Exchange is configured correctly, you will need to apply the PBR rules on the router(s). There is a functionality to generate PBR rules for different router vendors. You will need to review the generated ruleset and apply it on the router manually.
- It is possible that the Autoconfiguration feature on the Exchange has been run with incomplete routing tables or that new peering partners have been connected. This is especially true for very big Exchanges. In this case it is possible that some peering partners are neither in IRP nor router configurations. When such a case is detected, a warning about newly discovered peering partners gets raised. At this stage you will need to both re-autoconfigure IRP instance and extract the PBRs for the router.
- After its creation and up to its complete configuration the Exchange is kept in a Stopped state. When both IRP and the Router PBRs are configured, particular IRP instance is ready to start optimizing Exchange traffic too. Keep in mind that starting the Exchange will require a BGP session restart.

i We must mention that before applying changes to IRP instance configuration the changes are validated. If errors are detected these will be flagged and the previous good configuration will be kept intact so you will have the option to review the erroneous data and correct.

4.12.8.3 Switching a provider from one router to another

Switching a provider from one router to another is possible via manual IRP instance reconfiguration only.

Use the following steps to switch the provider from one router to another:

1. Suspend the provider using GMI Frontend "Providers and Peers" particular IRP configuration section. IRP instance withdraws all existing and stops new improvements towards shutdown provider(s)
2. You can remove traffic from the provider (by denying incoming/outgoing announces in the BGP configuration) and then physically re-cable the provider from one router to another. Then configure BGP session towards the provider on new router
3. The PBR rule for the provider should be configured on new router (move provider's PBR settings from the old router to the new one)
4. Change IRP box local PBR rules if any
5. Check if PBR works properly using traceroute tool
6. Modify (/etc/noction/irp.conf) assigned router for the provider (refer to [peer.X.bgp_peer](#))
7. Modify (/etc/noction/irp.conf) SNMP interface/IP/community for the provider (refer to [SNMP Host](#), [peer.X.snmp.interfaces](#), [peer.X.mon.snmp](#))
8. Reload Bgpd configuration (affected BGP sessions could be reset)

9. Re-activate the provider using GMI Frontend “Providers and Peers” IRP instance of choice configuration section
10. Check IRP BGP Internal and External monitor states. They must be UP.

4.12.9 SNMP hosts configuration

SNMP hosts are nodes on the network that provide or read SNMP data. Multiple SNMP hosts are configured in this section of configuration and references to them are used from elsewhere in the configuration. SNMP hosts can be configured by navigating via to “Configuration→SNMP hosts” tab.

See also: [SNMP Host](#)

Figure 4.12.36: Configuration editor: SNMP hosts configuration

- IRP instance Notifications

Notifications are messages sent to alert about some events registered by IRP instances. There are two prerequisites in order to start using notifications:

1. Configure event parameters and channels
2. Subscribe for event notifications

IRP instances can send notifications as SMS, email and SNMP Trap.

4.12.9.1 Configure notifications and events

Events configuration changes default threshold values when an event is raised. The most relevant events are presented in the following figure. For example Overloaded by (%) indicates that this event will fire when the aggregated outbound bandwidth usage for all providers exceeds by 10% the configured limits. Refer section 5.12 for details.

Figure 4.12.37: Configuration editor: Events configuration

IRP uses a local email server to send emails. Still, email sent by this server might trigger filtering policies enforced by some third parties that might block some important event notifications. It is possible to provide the details of another email server on your network that is protected for this. The following figure highlights the available email channel configuration parameters. Besides specifying the server email configuration sets the sender of email messages that will show in receiver's inbox.

i Only SMTP servers without authentication are currently supported.

EMAIL CHANNEL CONFIGURATION

Server port: 25

Email server: 127.0.0.1

From email: root@localhost

Figure 4.12.38: Configuration editor: Email configuration

SMS configuration is mandatory for sending SMS notifications.

The following SMS gateways are supported:

- [Twilio](#)
- [Plivo](#)

A valid account with a supported SMS gateway is required in order to finish configuration. Check the following figure for details:

SMS CHANNEL CONFIGURATION

SMS gateway: None

Account ID:

From phone number:

Secret:

Max message size: 150

Show advanced

Figure 4.12.39: Configuration editor: SMS configuration

The settings are:

- SMS gateway - choose one of the supported gateways
- Account ID - the public identifier assigned to your account by the SMS gateway. The following figure highlights this value for Plivo
- From phone number - the phone number that shows as sender on the receiver's mobile device. Use a phone number supported by the SMS gateway
- Secret - the secret identifier assigned to your account by the SMS gateway. The following figure highlights this value for Plivo

- Max message size - the maximum length of the SMS text. The SMS will be trimmed before sending. Subsequently the SMS gateway will split the SMS into multiple parts if required.

Refer section 5.12 for complete details about configuration parameters.

⚠ Some SMS gateways enforce the From phone number and will reject numbers that do not comply with their policy.

Figure 4.12.40: Plivo account ID and secret

There is a list of parameters that should be configured to deliver SNMP traps. Refer section 5.12 for details.

Figure 4.12.41: Configuration editor: SNMP Traps configuration

Web Hooks configuration requires only a Webhook URL to be provided. There is a series of advanced configuration parameters in case their default settings are not satisfactory. Refer section 5.12 for details.

Figure 4.12.42: Configuration editor: Web Hooks configuration

The settings are:

- Webhook URL - the unique URL the team is assigned by Web hook provider. An example for Slack.com is shown
- Bot name - optional bot name assigned to Webhook bot
- Avatar icon URL - optional parameter that designates an icon (usually PNG or JPEG image are supported) assigned to the Webhook bot in the channel
- Avatar emoji - an alternative avatar specified in the form of an emoji in “:robot-face:” format.

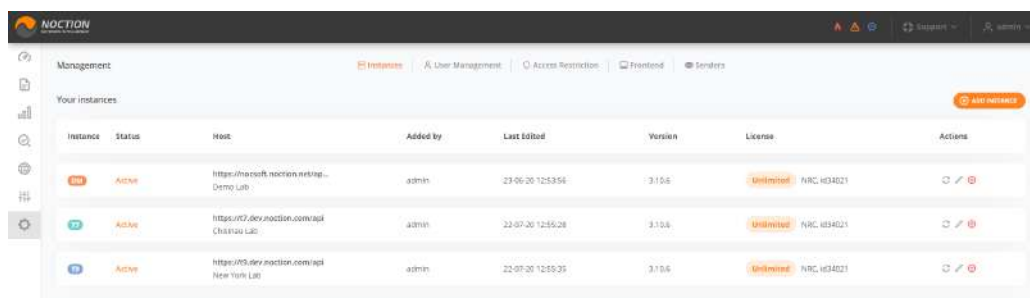
⚠ Web hooks support has been tested and verified to work for Slack.com API.

4.13 Management

4.13.1 Instances

The “Instances” tab under “Management” accumulates information about all IRP instances added to GMI. The section presents IRP shortnames, status, hosts, license, IRP versions as well as the “Last Edited” and “Added by” details.

Admins are free to add new, edit, and delete IRP instances as they like.



Instance	Status	Host	Added by	Last Edited	Version	License	Actions
001	Active	https://haassoft.noctium.net/api... Demo Lab	admin	23-06-20 12:53:56	3.1.0.6	Unlimited NRC-6034021	[Edit] [Delete]
02	Active	https://123.dev.noctium.com/api... Christou Lab	admin	22-07-20 12:55:28	3.1.0.6	Unlimited NRC-6034021	[Edit] [Delete]
03	Active	https://123.dev.noctium.com/api... New York Lab	admin	22-07-20 12:59:39	3.1.0.6	Unlimited NRC-6034021	[Edit] [Delete]

Figure 4.13.1: IRP Instances

To add a new instance, one should click the ADD INSTANCE button and fill out the corresponding fields in a popup window. The actual IRP instance access credentials should be provided for GMI to obtain the token.

Add instance

Instance host

 Host or IP including the endpoint, if applicable

Short name

 Very short instance name

Description

 Instance description

Username

 For obtaining the token from the IRP

Password

Figure 4.13.2: Adding an instance

4.13.2 User Management

The User Management tab is accessible under the Management main menu section.

Users who can access GMI are managed in User Directories. User directory is a generic name for an external user management provider for example LDAP (either POSIX or Active Directory). GMI can interface with as many user directories as needed in a specific environment.

User Directories

Order	On/Off	Name	Type	Hostname	Actions
0	<input checked="" type="checkbox"/>	LDAPS (test)	LDAP POSIX	ldap://ldap.test.nocton.com	
0	<input checked="" type="checkbox"/>	IRP Directory	LDAP POSIX	ldap://ldap.nocton.com	

All users

#	Username	Email	Enabled	User Directory	Instance	Privileges	Created by	Creation date	Actions
1	trial	trial@example.com	enabled	Internal	05	Manager	admin	04/11/2017 17:28:07	
2	user	user@user.com	enabled	Internal	06	User	mike	12/03/21 09:09:06	
3	max	max@max.com	enabled	LDAPS (test)	All instances	Admin	admin	21/07/21 11:18:08	
4	mike	mike@nocton.com	enabled	Internal	All instances	Admin	admin	26/09/2018 25:55	
5	mmandak	mmandak@nocton.com	enabled	IRP Directory	06	User	admin	21/07/21 21:04:59	
6	admin1	admin1@test.com	enabled	Internal	All instances	Admin	admin	26/09/21 09:52:21	

Figure 4.13.3: User Management

4.13.2.1 Internal user directory

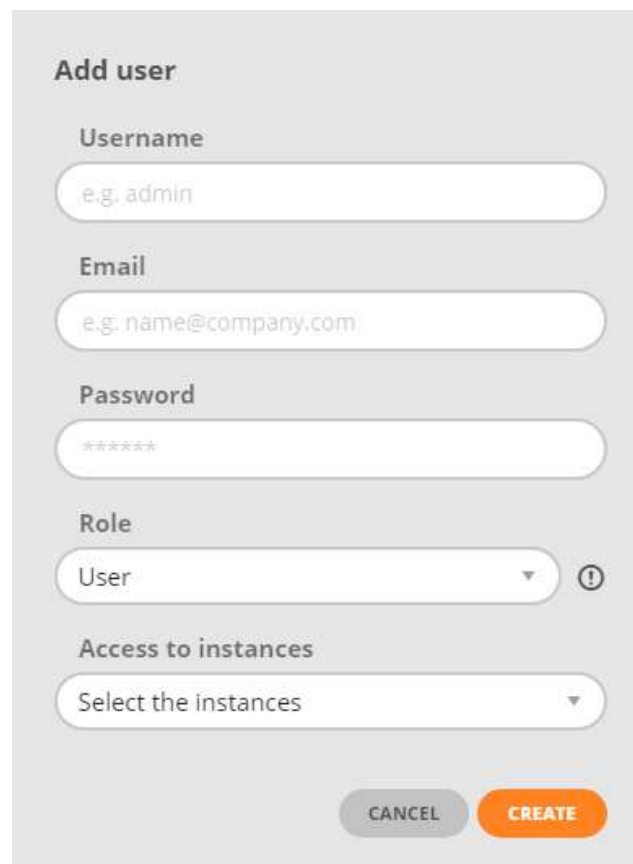
GMI includes an Internal user directory that allows multiple user accounts to be setup.

To add a new user, one should click the ADD USER button and fill out the corresponding fields in a popup window.

There are 3 user roles available in GMI, with specific limitations and privileges:

- User - cannot manage other users or instances

- Manager - cannot manage users
- Admin - full privileges



The screenshot shows a web form titled "Add user" with a light gray background. It contains several input fields: "Username" with a placeholder "e.g. admin", "Email" with a placeholder "e.g. name@company.com", and "Password" with a placeholder "*****". Below these is a "Role" dropdown menu currently set to "User" with a small information icon to its right. At the bottom of the form is an "Access to instances" dropdown menu with the text "Select the instances". At the very bottom are two buttons: a gray "CANCEL" button and an orange "CREATE" button.

Figure 4.13.4: Adding users

4.13.2.2 LDAP and Active Directory

LDAP or AD user directories can be added, updated and removed from GMI by accessing “Management→User Management” tab. Each user directory takes a series of parameters specific for the protocol.

⚠ All operations with DN's (initial bind DN, group DN's, user names) are case insensitive and also strip redundant whitespace.

Refer individual protocol documentation for how to correctly configure one or another user directory. The example below offers a generic set of parameters required to configure GMI to use Active Directory for access management.

Add new user directory

GENERAL ADVANCED BINDINGS ACCESS

User Directory Name:

Hostname: LDAP

State: ON OFF

When directory is disabled: Disable user

Port:

Order: 0 255

CANCEL SAVE

Figure 4.13.5: User Directory configuration

The general tab covers:

- User directory name - the name assigned to the directory within GMI,
- User directory hostname in the form of either IP address or domain name (LDAP/LDAPS),
- Enabling or disabling a user directory,
- The option to disable or remove users completely for a disabled directory.
- User directory port
- Order specifies when this user directory will be examined by GMI compared to other user directories,

Add new user directory

GENERAL ADVANCED BINDINGS ACCESS

Timeout: 0 120

TLS: ON OFF

Certificate verification: ON OFF

TLS CA Certificate file: No file selected CHOOSE

Initial bind DN:

Initial bind password:

CANCEL SAVE

Figure 4.13.6: User Directory Advanced Options

Advanced configuration of a user directory covers:

- Timeout before failing a connection to this user directory

- TLS use
- Certificate verification
- CA certificate used to verify server's certificate in case the Certificate verifications is turned on
- Initial binding user name that GMI uses to authenticate itself
- Initial bind password assigned to GMI.

i Usually bind passwords are not required to access directories. If a password is required and configured via GMI Frontend, it will be scrambled in GMI configuration files. If the password is specified directly in GMI configuration it is provided in clear-text with the condition that it does not start/end with scrambled password encapsulation characters.

The screenshot shows a web form titled "Add new user directory" with four tabs: "GENERAL", "ADVANCED", "BINDINGS" (selected), and "ACCESS". Under the "BINDINGS" tab, there are four input fields arranged in two rows. The first row has "Base DN" and "Username attribute" (containing the text "uid"). The second row has "Email attribute" (containing the text "email") and another "Base DN" field (empty). At the bottom right of the form are two buttons: "CANCEL" and "SAVE".

Figure 4.13.7: User Directory Bindings

Bindings maps User Directory attributes to GMI specific attributes, for example:

- Base DN specifies the root distinguished name and user subtree

i GMI recognizes BOTH short and full user identifiers. Examples below are both valid directory entries that will match user “chris” with long name “Mr. Chris Smith”:

```
cn: ops
uniqueMember: chris
```

and

```
cn: ops
uniqueMember: cn=Mr. Chris Smith,ou=employees,ou=People,dc=ops,dc=org
```

- Username and Email fields map User Directory attributes to GMI user attributes,

Edit user directory - IRP Directory

GENERAL ADVANCED BINDINGS **ACCESS**

Roles ^① Bind group Access to instances

User ou=Users Select the insta... **ADD**

Binding groups

All instances	User role: ADMIN	Bind group: cn=admin,ou=irp,ou=app-groups
E7	User role: USER	Bind group: ou=Users
E7	User role: USER	Bind group: asd
E7	User role: USER	Bind group: cn=user,ou=irp,ou=app-groups

CANCEL UPDATE

Figure 4.13.8: User Directory Access Options

Access configuration of a user directory covers:

- Roles assigned (either User, Manager or Admin),
- Bind group - the name of the attribute that uniquely identifies a given group or user.

⚠ Do NOT provide the Distinguished Name - the name that includes an object's entire path to the root of the LDAP namespace. Introduce the Relative Distinguished Name instead - an object name without a path, or a partial path (Example: "cn=support").

- Access to IRP instances available in GMI. (Users with the Admin roles have access to all IRP instances)

4.13.3 Access Restriction

This section covers restrictions on the ranges of IP addresses that can access GMI's Frontend and an option to add an SSL Certificate.

NOCTION

Management

Instances User Management **Access Restriction** Frontend Senders

FRONTEND ACCESS CONTROL LIST

Frontend **ADD**

SSL CERTIFICATE

Select your public certificate **UPLOAD**

Select your private certificate **UPLOAD**

IMPORT

Warning: options to restrict access to a web application exist. GMI assumes this is the case and accepts requests from everywhere; if access is restricted from within GMI make sure to NOT lower the default values that allow to NOT deny access to yourself.

Figure 4.13.9: Access Restriction

4.13.4 Frontend

Custom Login screen and Navigation Bar logo images can be uploaded by using the Upload buttons. The new settings are applied after clicking the “SAVE” button.

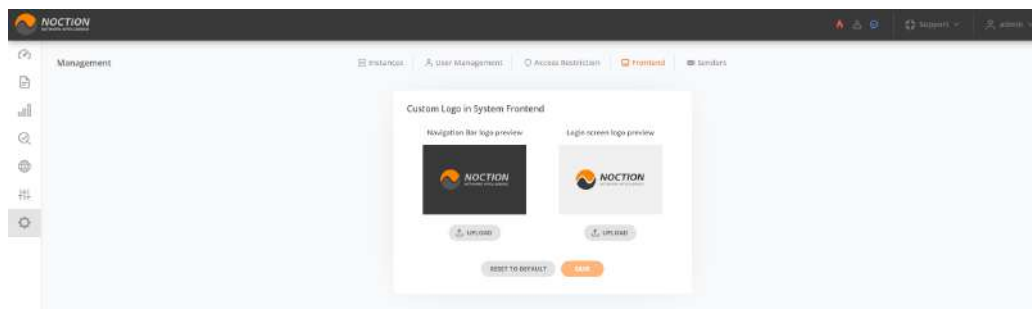


Figure 4.13.10: Frontend

4.13.5 Senders

To receive GMI reports and graphs via email or Slack, one should setup GMI senders. Click the “ADD A SENDER” button, fill out the required fields and click save.




Figure 4.13.11: Senders

4.14 Emailing subscription

Proceed to Username→Emailing and sign up for the reports and graphs emails. Edit existing subscriptions or delete them in this view as well.

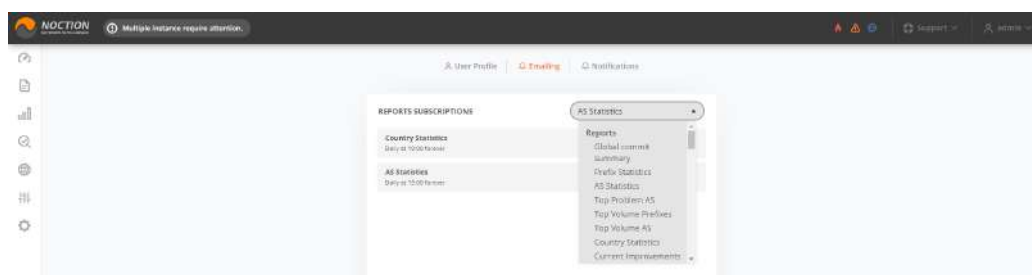


Figure 4.14.1: Reports and Graphs emailing subscription

Select the desired report or graph. Fill out the required fields and click Save.

Edit subscription

Subject

AS Statistics

Destination

info@noction.com ✕ Add destination

Send until

send forever

☒ Forever

Delivery frequency

Every day ▼

At

19:00 ▼

Comment

AS statistics report for Chisinau Lab

CANCEL

SAVE

Figure 4.14.2: Emailing subscription details

4.15 Events

The system events are diagnostic messages that were sent by the Components of IRP instances, as well as notifications caused by the BGP monitoring algorithm. The events are displayed in the status bar located at the top of the system frontend (See [Frontend structure](#)). A specific IRP instance should be selected to display the corresponding results. Several or all the events can be marked as “Read”, so they will no longer be displayed in the status bar.

NOCTION

System Events

Demo Lab

Search System Events

Last 20 items

Timestamp	Module	Type	Event
<input type="checkbox"/> 21-07-20 11:00:26	Explorer	info	Explorer started
<input type="checkbox"/> 21-07-20 11:00:10	irpapid	info	irpapid started
<input type="checkbox"/> 21-07-20 11:00:07	irpapid	info	irpapid is stopped
<input type="checkbox"/> 21-07-20 09:50:09	Explorer	info	Explorer started
<input type="checkbox"/> 21-07-20 09:50:08	irpapid	info	irpapid started
<input type="checkbox"/> 21-07-20 09:50:05	irpapid	info	irpapid is stopped
<input type="checkbox"/> 21-07-20 09:26:54	Core	info	No routing updates, liss improvement discarded due to the Static Route policy applied to the 192.16.24.0/21
<input type="checkbox"/> 21-07-20 09:12:12	Core	info	No routing updates, liss improvement discarded due to the Static Route policy applied to the 192.16.24.0/21
<input type="checkbox"/> 21-07-20 08:37:09	Core	info	No routing updates, liss improvement discarded due to the Static Route policy applied to the 192.16.24.0/21
<input type="checkbox"/> 21-07-20 08:00:35	Explorer	info	Explorer started
<input type="checkbox"/> 21-07-20 08:00:33	irpapid	info	irpapid started

Mark as Read

Figure 4.15.1: System events


4.16 Notifications and Events

IRP instances produce a huge number of various events and some of them are critical for administrator's awareness. Notifications in GMI allow administrators to subscribe to any of the available IRP instances generated events using the following channels:

- SMS
- Email
- Slack (via Webhook)
- SNMP Traps

IRP service Ippushd provides this feature. In order for Notifications to be delivered correctly the corresponding channel configuration shall be provided. By default only email notifications can be delivered since IRP uses the embedded system email service to send them.

More so, users should subscribe for specific events.

 Only events for valid subscriptions using correctly configured channels will be delivered.

Refer section [SNMP hosts configuration](#) for details about configuring, subscribing and contents of notifications.

Refer section [Notification and events](#) for details about individual configuration parameter.

4.16.1 Events

The list of events monitored by IRP that can generate notifications is provided below.

When one of the IRP components detects a transition from normal to abnormal traffic behavior or back it fires these events:

- Abnormal correction: irpflowd
- Abnormal correction: irpspand
- Inbound traffic low: SPAN
- Inbound traffic low: Flow
- Inbound traffic normal: Flow
- Inbound traffic normal: SPAN
- Outbound traffic low: SPAN
- Outbound traffic low: Flow
- Outbound traffic normal: Flow
- Outbound traffic normal: SPAN

When Commit Control limits are exceeded per provider or overall one of the following events fires. Refer section [5.12](#) for configuring the actual limits of the events.

- Commit Control overload by X Mbps
- Commit Control overload by X%
- Commit Control provider X overloaded by Y Mbps
- Commit Control provider X overloaded by Y%

When an IRP component (re)loads the configuration it validates it and depending on results fires one of the following events:

- Configuration Invalid: Bgpd
- Configuration Invalid: Core
- Configuration Invalid: Explorer
- Configuration Invalid: Irpapid
- Configuration Invalid: Irpflowd
- Configuration Invalid: Irpspand
- Configuration Ok: Bgpd
- Configuration Ok: Core
- Configuration Ok: Explorer
- Configuration Ok: Irpapid
- Configuration Ok: Irpflowd
- Configuration Ok: Irpspand

Outage detection algorithm fires one of the following events when it confirms congestion or outage problems and reroutes traffic around it:

- Congestion or Outage
- Outage: Confirmed and rerouted

Explorer periodically checks the PBRs and its expected probing performance and triggers the following events:

- Failed PBR (IPv6) check for provider
- Failed PBR (IPv4) check for provider
- Successful PBR (IPv4) check for provider
- Successful PBR (IPv6) check for provider
- Explorer performance low
- High number of VIP prefixes degrades IRP performance

IRP BGP Internal and External monitors fire the following events:

- ExternalMonitor (IPv4) Failed status for a provider. All improvements towards the provider will be withdrawn.
- ExternalMonitor (IPv4) OK status for a provider. All improvements towards the provider will be announced.
- ExternalMonitor (IPv6) Failed status for a provider. All improvements towards the provider will be withdrawn.
- ExternalMonitor (IPv6) OK status for a provider. All improvements towards the provider will be announced.
- InternalMonitor (IPv4) Failed status for a provider. All improvements towards the provider will be withdrawn.

- InternalMonitor (IPv4) OK status for a provider. All improvements towards the provider will be announced.
- InternalMonitor (IPv6) Failed status for a provider. All improvements towards the provider will be withdrawn.
- InternalMonitor (IPv6) OK status for a provider. All improvements towards the provider will be announced.

When statistics collection over SNMP is up or down IRP fires the following events:

- Provider SNMP stats down: X
- Provider SNMP stats up: X

Bgpd raises these events when BGP sessions are established/disconnected:

- IRP BGP session disconnected
- IRP BGP session established

When IRP identifies conditions to re-route traffic (make an improvement) and additionally it considers the differences to be excessive it raises these events:

- Excessive packet latency for prefix
- Excessive packet loss for prefix
- Improvements spike
- Low rate of announced IPv4 improvements
- Low rate of announced IPv6 improvements
- New improvement

Once an IRP component is started, stopped or restarted it raises the following events:

- Service started: Bgpd
- Service started: Core
- Service started: Explorer
- Service started: Irpapid
- Service started: Irpflowd
- Service started: Irpspand
- Service stopped: Bgpd
- Service stopped: Core
- Service stopped: Explorer
- Service stopped: Irpapid
- Service stopped: Irpflowd
- Service stopped: Irpspand

4.16.2 SNMP Traps

SNMP traps is a widely used mechanism to alert about and monitor a system's activity.

IRP SNMP traps not only notify about some IRP platform event but also include the list of varbinds which contain detailed information related to the thrown trap. The complete list of traps and varbinds with their descriptions can be found at `/usr/share/doc/irp/NOCTION-IRP.mib`

4.16.3 Notifications

Notifications are sent only if a valid subscription has been created. Subscriptions, similar to regular emailing of reports can be found under Username→Notifications menu.

Notifications page provides an overview of existing subscriptions with features to create, edit, delete them. Check the following figure for a preview.




Figure 4.16.1: Notifications

Use Add new subscription or Events drop-down to create a new subscription or edit, delete existing subscriptions by clicking corresponding icons. Creating or editing a subscription will bring up a pop-up like the one in the following figure.

Figure 4.16.2: Subscribe notifications

Subscription details include:

- Topic - the title given to subscription in Notification page and also a textual part of notifications

i If SMS notifications are used it is advised to keep the topic short so that there is enough space left to include details about the event

- Choice of one or more IRP instances
- Interval between notifications - sets a rate limit of how often an email and/or SMS notification is sent. Value 'No limit' for the interval depicts no rate limits and all events will trigger a notification

i SNMP Traps notification are not constrained by the interval between notifications. SNMP Traps are sent immediately as they are raised

⚠ The rate limit is enforced per subscription so it is still possible to receive multiple notifications if the subscribed events are part of different subscriptions

- Destinations - IP address of a trap receiver for SNMP Trap notifications, email addresses, phone numbers and webhook channels. Multiple destinations of same type can be provided. At least one valid destination must be provided per subscription

ℹ IRP parses and recognizes whether a provided destination is a valid IP address, email address, phone number or web hook channel.

Supported formats are:

email name@host.domain
phone +digits
webhook #channel_name
snmptrap ipv4 or ipv4:port or [ipv6] or [ipv6]:port

- Event filters - allows filtering of events by textual description or by IRP component. Remove the filters to see all the subscribed events
- Subscribed events - the full list of events supported by IRP and tick marks for the events in this subscription. At least one event is required for a valid subscription.

4.17 User Profile

One can access Profile details by navigating to <Username> → Profile in the top right menu. This section allows to see the general profile details and preferences including the user's email, date and time formats and browser notifications permission. There is also an option to change password and see the most recent active sessions details.

The screenshot displays the 'User Profile' interface with two main sections: 'GENERAL' and 'SECURITY'.

GENERAL Section:

- Email:** A text input field containing 'info@noction.com'.
- Display date format:** A dropdown menu showing 'DD-MM-YY'.
- Display time format:** A dropdown menu showing '24 hours'.
- Date format example:** '27-07-20 20:12:57'.
- Allow browser notifications:** An unchecked checkbox.

SECURITY Section:

- Password Requirements:**
 - Minimum 8 characters long
 - New and repeat password match
 - All required fields filled
 - Contains numbers (Optional)
 - Contains capital letters (Optional)
 - Contains symbols (Optional)
- Current password:** A text input field with masked characters.
- New Password:** A text input field with masked characters.
- Confirm Password:** A text input field with masked characters.

LATEST ACTIVE SESSIONS Section:

A table showing the latest active sessions, with a total count of 20. The table lists session details including the user, IP address, device, and session time.

User	IP Address	Device	Session Time
Ruppasner	19-07-20 19:01:14		X
185.34.203.23	27-07-20 20:12:29	Linux	X
Ruppasner	18-07-20 18:01:03		X
52.51.203.11	27-07-20 19:01:16	Linux	X
52.51.203.11	22-07-20 11:17:06	Windows	X

At the bottom right, there are 'RESET' and 'SAVE' buttons.

Figure 4.17.1: User Profile

4.18 Maintenance windows

Maintenance works are on everybody's agenda in current fast paced and continuously evolving networks. During maintenance network engineers are very busy and will welcome any help their systems can offer in carrying out those works with the least amount of headaches. IRP instances are clearly not in the top of network engineer's priorities and asking to suspend or shutdown providers immediately before a maintenance window starts and restart the provider back once the maintenance works end is not very helpful if not even annoying.

GMI offers the facility to plan maintenance windows in advance for providers on the connected IRP instances. Knowing when a maintenance window starts and ends, the IRP instances exclude specific provider links from either performance optimization or bandwidth control. More so, there is a capability to reshape the traffic flowing in and out of a network to anticipate any downtime on a link.

Properly configured maintenance windows allows IRP instances the time to move most of the outbound traffic and deflect most of inbound traffic away from the provider link that is scheduled for maintenance. Having only a small fraction of traffic or none at all on the maintenance link before the downtime starts avoids any (shall we say, catastrophic) spikes, possible overloads and consequently unpredictable behavior of the remaining live network equipment.

Specifically the following applies:

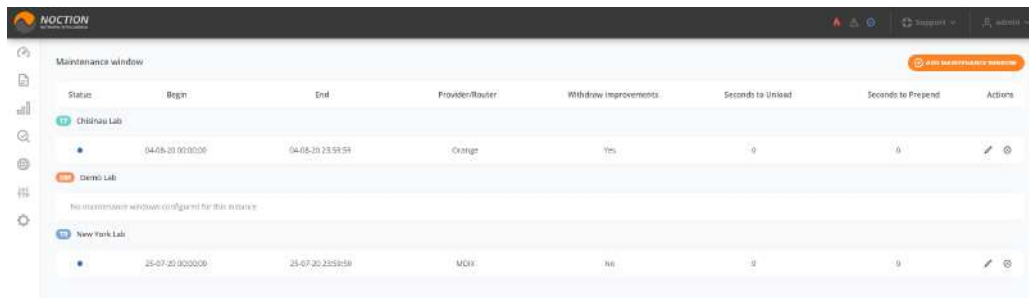
- a maintenance window is configured in advance and can be removed/revised at any time
- a maintenance window sets details for single provider. If needed multiple maintenance windows can be setup and even overlapping maintenance windows are OK
- GMI highlights maintenance windows in the Frontend, so that it is easy to spot current maintenance window status
- optionally GMI can instruct IRP instances to preserve existing improvements so that once the maintenance window ends improvements are reimplemented. It is advised that this feature is used only when the maintenance window is very short (a few minutes long)
- an unloading period can be setup. During unloading IRP instances actively re-route outbound prefixes through other available providers. While IRP instances are able to make most of the unloading improvements fast, consideration shall be given to the announcement rate limitations setup in Bgpd in order for all the improvements to reach network routers in time for maintenance window starting time
- a prepend time can be setup. This is only applicable if Inbound optimization is operational. If this time is setup then an IRP instance will prepend configured inbound prefixes with the maximum allowed number of preprends through the provider link under maintenance in order to deflect inbound traffic towards other providers.

To review Maintenance windows, navigate to the corresponding left side bar menu option.



Figure 4.18.1: Maintenance windows menu button

The list displays all the current and future maintenance windows configured for providers in various IRP instances.



Status	Begin	End	Provider/Router	Withdraw Improvements	Seconds to Unload	Seconds to Prepend	Actions
Chinhou Lab							
•	04-05-20 00:00:00	04-05-20 23:59:59	Orange	Yes	0	0	
No maintenance windows configured for this instance							
New York Lab							
•	25-07-20 00:00:00	25-07-20 23:59:59	MDN	No	0	0	

Figure 4.18.2: Maintenance windows

As shown in the screen-shot above, the list highlights:

- Current color-coded status of the maintenance window. Red and Orange depict unloading and actual maintenance phases while Blue and Green depict future planned windows and past finished windows correspondingly
- Begin and End date-times of maintenance windows
- Provider under maintenance
- Withdraw Improvement option enables or disables withdrawal of existing Outbound Improvements to that provider
- Seconds to Unload defines time interval in seconds prior to beginning of the maintenance window when a particular IRP instance starts to unload outbound traffic from provider

⚠ If Seconds to Unload/Seconds to Prepend time intervals set to zero the IRP instance does not perform corresponding action

ℹ When seconds to unload is specified it should be at least 5 minutes (300 seconds) long with larger time intervals allowing IRP more time to assess and reroute traffic flowing through provider with scheduled maintenance window. Very short time intervals might not have the desired effect since usually an IRP optimization cycle takes 2-3 minutes to finish

- Seconds to Prepend defines time interval in seconds prior to beginning of the maintenance window when IRP announces all inbound prefixes with maximum preprends to provider in order to deflect inbound traffic towards other providers.

Maintenance windows can be added, edited or deleted using the designated action buttons.

A maintenance window is added by specifying the following:

Add maintenance window

Instance

T7

Chisinau Lab

Provider

Choose a provider

Withdraw Improvements

Enable | Disable

10:47:00 21-07-20

10:46:46 22-07-20

<

July 2020

August 2020

>

		1	2	3	4	5						1	2
6	7	8	9	10	11	12	3	4	5	6	7	8	9
13	14	15	16	17	18	19	10	11	12	13	14	15	16
20	21	22	23	24	25	26	17	18	19	20	21	22	23
27	28	29	30	31			24	25	26	27	28	29	30
							31						

Your time will be converted to UTC format:

Seconds to unload

0

Seconds to prepend

0

CANCEL

CREATE

Figure 4.18.3: Maintenance window popup

Chapter 5

Configuration parameters reference

5.1 Database credentials

5.1.1 db.clickhouse.dbname

Name of the ClickHouse database that contains the IRP tables.

Default value: `irp`

5.1.2 db.clickhouse.host

Database host. ClickHouse server host name or IPv4/IPv6 address.

Default value: `localhost`

5.1.3 db.clickhouse.http_port

TCP port number used to connect to the ClickHouse server via REST API.

Default value: `8123`

5.1.4 db.clickhouse.password

The password used to connect to a ClickHouse server.

The password will be changed to a random one by the IRP installation process.

Default value: `irp`

5.1.5 db.clickhouse.port

TCP port number used to connect to the ClickHouse server via native API.

Default value: `9000`

5.1.6 db.clickhouse.username

User name used to connect to a ClickHouse server.

Default value: `irp`

5.1.7 db.dbname

Name of a MySQL database that holds the IRP tables.

Default value: `irp`

5.1.8 db.host

Database host. MySQL server host name or IPv4/IPv6 address.

Default value: `localhost`

5.1.9 db.ourhost

IRP host. IRP box host name or IPv4/IPv6 address.

Default value: `localhost`

5.1.10 db.password

The password used to connect to a MySQL server.

The password will be changed to a random one by the IRP installation process.

Default value: `irp`

5.1.11 db.port

TCP port number to use for the connection to a MySQL database.

Default value: `3306`

5.1.12 db.username

User name used to connect to a MySQL server.

Default value: `irp`

5.2 Global parameters

5.2.1 `global.agg_ipv4_max`

Defines the maximum IPv4 aggregate mask for the improved prefixes advertised by Bgpd. If `global.aggregate` is enabled, IRP will not announce any prefix with the mask that's less than the mask defined in `global.agg_ipv4_max`.

Possible values: 8–24

Default value: 16

5.2.2 `global.agg_ipv6_max`

Defines the maximum IPv6 aggregate mask for the improved prefixes advertised by Bgpd. If `global.aggregate` is enabled, IRP will not announce any prefix with the mask that's less than the mask defined in `global.agg_ipv6_max`.


Possible values: 32–48

Default value: 32

5.2.3 `global.aggregate`

If this parameter is enabled IRP analyzes entire aggregates. When one can be improved the aggregate will be announced by Bgpd (refer to `bgpd.updates.split`, `bgpd.peer.X.updates.limit.max`). Largest possible prefixes are used when aggregates exceed `global.agg_ipv4_max` & `global.agg_ipv6_max` parameters for IPv4 and IPv6 accordingly. An aggregate dictionary is maintained by IRP in order to support this capability. The aggregate dictionary is periodically populated by `refresh_asn` or Bgpd (if Bgpd has sufficient information, by means of at least one full-view BGP peering session).

Disabling this parameter instructs IRP to operate with the smallest possible disaggregated prefixes (/24 for IPv4 and /48 IPv6).

 Switching this parameter ON/OFF might leave a large number of small prefixes in the aggregate dictionary with undesirable side effects. The contents of the dictionary should be reviewed and refreshed in case it lists undesirable prefixes. To prevent issues IRP improvements should be cleared before applying this change. Clearing improvements ensures that disaggregated prefixes announced by IRP are not present on the network.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.2.4 `global.bw_overusage`

Enables or disables automatic Flowspec throttling policies for prefixes with excessive bandwidth spikes.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.2.5 `global.exchanges`

Defines the path to the Exchanges configuration file.

Default value: `/etc/noction/exchanges.conf`

Recommended value: `/etc/noction/exchanges.conf`

5.2.6 `global.exchanges.auto_config_interval`


Defines the Internet Exchanges auto re-configuration interval in seconds. Auto re-configuration is applied to providers of type Internet Exchange with enabled auto re-configuration. Refer [peer.X.auto_config](#).

Possible values: 3600–2678400

Default value: 86400

5.2.7 `global.failover`

Enables and disables failover feature.

 Enabling failover requires extensive preparatory activities. Refer [IRP Failover](#), [Failover Configuration](#), [Setup Failover wizard](#) for details.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.2.8 `global.failover.log`

Path to failover log file.

Default value: /var/log/irp/failover.log

5.2.9 `global.failover.use_communities`

This parameter allows monitoring presence of improvements from master instance by using communities specified in [bgpd.peer.X.master_communities](#) (refer [bgpd.peer.X.slave_communities](#)).

Master and slave communities in a BGP peering should be unique if the parameter is enabled.

If a network setup requires the same community to be used on both master and slave instances then the parameter should be disabled and failover shall use [bgpd.peer.X.master_localpref](#) only.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.2.10 `global.failover_identity_file`

Defines the path to failover master identity file. Set this parameter only when a SSH key file other than `~/.ssh/id_rsa` is used. Refer [global.failover](#).

Possible values: Path to identity file

5.2.11 `global.failover_role`

Defines the role of the IRP instance in a failover configuration. Known roles are Master and Slave. Used only when failover is enabled. Refer [global.failover](#).

Possible values: 0 (Master), 1 (Slave)

Default value: 0

5.2.12 `global.failover_slave.ip`

Defines the IPv4 address of the slave IRP instance used for configuration sync in a failover configuration. Used only when failover is enabled. Refer [global.failover](#).

Possible values: Valid IPv4 address

Example value: 10.10.0.2

5.2.13 `global.failover_slave.ipv6`

Defines the IPv6 address of the slave IRP instance used for configuration sync in a failover configuration. Used only when failover is enabled. Refer [global.failover](#).

Possible values: Valid IPv6 address

Example value: 2001:db8::ff00:42:8329

5.2.14 `global.failover_slave.port`

Defines the SSH port of the slave node in a failover configuration. Used only when failover is enabled. See also [global.failover](#).

Possible values: 1-65535 (Valid port number)

Default value: 22

5.2.15 `global.failover_timer_fail`

Defines the period of time in seconds during which master is considered alive before the slave becomes active in a failover configuration. Used only when failover is enabled.

See also [global.failover](#).

Possible values: 30-3600

Default value: 300

5.2.16 `global.failover_timer_failback`

Defines the period of time in seconds for slave to switch back to standby after it detects master became alive in a failover configuration. Used only when failover is enabled.

See also [global.failover](#).

Possible values: 30-3600

Default value: 300

5.2.17 `global.flowspec`

Enables/disables Flowspec capability globally.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.2.18 `global.flowspec.pbr`

Enables/disables use of Flowspec policies instead of Policy Based Routing. This is available only when Flowspec is enabled. Refer [global.flowspec](#).

Possible values: 0 (Disabled), 1 (Enabled in intrusive), 2 (Enabled always)

Default value: 0

5.2.19 `global.frontend_acl`

Defines whether access to the IRP Frontend must be restricted and controlled by an ACL. See also [global.frontend_acl_ips](#).

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.2.20 `global.frontend_acl_ips`

Defines the IPs or networks that should be allowed to access the IRP Frontend.

Possible values: Valid IPv4, IPv6 address or subnet definition in CIDR format

Default value: 0.0.0.0/0 ::/0

5.2.21 `global.ignored.asn`

Defines the list of ASNs to be ignored by IRP.

Format:

1. Space-separated list of AS numbers that must be ignored by IRP.
2. Absolute path to a newline-separated file containing a list of AS numbers that must be ignored by IRP.

This parameter should list all AS numbers that must be ignored by Irpspand, Irpflowd, Explorer and the Core. No improvements will be performed by the Core for prefixes that are announced by ASNs listed within this parameter. No data will be gathered by Irpspand/Irpflowd for any source or destination IPv4/IPv6 address that are announced by ASNs that are listed within this parameter. No probes will be sent by Explorer to any destination IPv4/IPv6 address that are announced by ASNs listed within this parameter.

Refer also [global.ignored_communities](#), [global.ignorednets](#).

Possible values: 1 - 4294967295.

5.2.22 `global.ignored.unannounced`

Prefixes which aren't present in a BGP routing table wouldn't be analyzed nor optimized by Outbound optimization algorithms.

As example, spoofed IP addresses may not belong to advertized prefixes but the traffic itself may have significant volume and be optimized by IRP.

Therefore, in such a case the recommendation is to enable the parameter to do not optimize traffic from unknown origins.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.2.23 `global.ignored_communities`

Defines the list of BGP Community attributes the network uses to mark prefixes that IRP must ignore. IRP monitors routes marked with at least one of these BGP Community attributes and dynamically updates a list of prefixes to ignore. The decision what prefixes to mark with one of these attributes is applied on the network and network operators do not need to be explicitly set in IRP too.

This (dynamic) list of prefixes will be ignored by Irpspand, Irpflowd, Explorer and the Core. No improvements will be performed by the Core for such prefixes. No data will be gathered by Irpspand/Irpflowd for any source or destination IPv4/IPv6 address within such prefixes. No probes will be sent by Explorer to any destination IPv4/IPv6 address within such prefixes.

Refer also [global.ignored.asn](#), [global.ignorednets](#).

Possible values: valid BGP community attribute list.


5.2.24 `global.ignorednets`

Defines the list of networks to be ignored by IRP.

Format:

1. Space-separated list of local IPv4/IPv6 prefixes that should be ignored by IRP.
2. Absolute path to a newline-separated file containing a list of IPv4/IPv6 prefixes that should be ignored by IRP.

If netmask is not clearly specified, the system assumes /32 for IPv4 addresses, and /128 for IPv6 addresses.

-  224.0.0.0/3 is always ignored.
- IPv6 probing is performed only to 2000::/3 address range (see: [IPv6 Address Space](#))
- All IPv4/IPv6 addresses assigned to IRP server are automatically added to ignored networks


This parameter lists all IPv4/IPv6 addresses that should be ignored by Irpspand, Irpflowd, Explorer and the Core. No improvements will be performed by the Core for /24 subnets listed within this parameter as /24 or a less specific network. No data will be gathered by Irpspand/Irpflowd for any source or destination IPv4/IPv6 address listed within this parameter. No probes will be sent by the Explorer to any destination IPv4/IPv6 address listed within this parameter.

Refer also [global.ignored.asn](#), [global.ignored_communities](#).

Possible values: See above.

Default value: 127.0.0.0/8 10.0.0.0/8 169.254.0.0/16 172.16.0.0/12 192.168.0.0/16
100.64.0.0/10 203.0.113.0/24 198.18.0.0/15 192.0.0.0/24 192.0.2.0/24
2001:db8::/32

Recommended value: See default value.


-  Recommended networks to ignore are:
- networks from: [Section 3 of RFC5735](#), as listed in the default value above
- networks from: [Section 2 of RFC3849](#)

5.2.25 `global.improve_mode`

Defines the IRP operating mode (see [IRP Optimization modes](#)).

This parameter adjusts the priorities and rules for prefix improvements. Prefixes can be improved in three different ways:

1. **Performance optimization:** Decrease loss, then decrease latency;
2. **Cost optimization:** Decrease loss, then decrease cost while keeping the latency within the preconfigured level;

 Loss, cost and latency are improved in strict order, depending on the selected operating mode.

Possible values: 1 (Performance), 2 (Cost) (see above).

Default value: 1

See also: [Commit Control](#), `core.commit_control`, `peer.X.cc_disable`

5.2.26 `global.inbound.injection`

Defines how IRP is used to inject Inbound improvements.

In “Pull” mode an external script is used to pull Inbound improvements from IRP API and re-configure routers’ access lists.

Example script `/usr/bin/irpTransitPull.pl` could be used to work in the Pull mode.

Possible values: 0 (Pull), 1 (BGP)

Default value: 1

5.2.27 `global.inbound_conf`

Defines path to file with configured inbound prefixes.

Possible values: path to file

Default value: `/etc/noction/inbound.conf`

Recommended value: `/etc/noction/inbound.conf`

5.2.28 `global.inbound_transit`


Enables or disables inbound optimization of transiting traffic. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.2.29 `global.ipv6_enabled`

Defines whether IPv6 is enabled in the system. Currently used for the Frontend only. Even if IPv6 is enabled via this parameter, other components configuration must be adjusted for IPv6 as well.

 IRP prior to 3.3-2 allows configuration of both IPv4 and IPv6 sessions on a single router, while IRP 3.3-2 requires definition of two separate BGP sessions. It is recommended to split BGP sessions in the form of two different routers before upgrading to 3.3-2, otherwise IRP configuration will not be valid. Make sure the newly added router is linked to corresponding providers to ensure IPv6 optimization works properly. InternalMon for IPv6 session monitoring requires correct configuration of BGP MIB (IPv6) mode (see `peer.X.mon.ipv6.internal.mode` parameter). Currently support for Brocade, Cisco and Juniper MIBs is available.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.2.30 `global.master_management_interface`


Defines the management network interface. In most cases it is the same as the probing interface. When failover is enabled (`global.failover`) slave's management interface (`global.slave_management_interface`) must be configured too.

Possible values: any valid system network interface name

Default value: `eth0`

5.2.31 `global.master_probing_interface`

Defines the probing network interface. In most cases it is the same as the management interface. When failover is enabled (`global.failover`) slave's probing interface (`global.slave_probing_interface`) must be configured too.


 This parameter is used only for displaying operating system interface(s) status in Frontend. It does not actually configure Explorer behavior, which depends on [Provider](#).

Possible values: any valid system network interface name

Default value: `eth0`

5.2.32 `global.master_rd`

Specifies the Routing Domain that hosts the master node of IRP in a failover configuration. By default RD=1 hosts IRP nodes.

 It is recommended that master node of IRP is hosted in RD=1 at all times.

Possible values: 1-100

Default value: 1

Recommended value: 1

5.2.33 `global.nonintrusive_bgp`

Instructs the system to run in a non-intrusive BGP mode (see [IRP Operating modes](#)).

All improvements made in a non-intrusive mode, will not be automatically injected into the routers.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

Recommended value: 1 at first start, 0 after manual tests are performed and the system is ready to go into intrusive mode

5.2.34 `global.offpeak_hour`


Defines customer's network usual off-peak hour of the day.

Possible values: 0 - 23

Default value: 3

5.2.35 `global.outbound`

Enables or disables outbound optimization.

 Outbound optimization is disabled only for standalone Inbound optimization is used.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

Recommended value: 1

5.2.36 `global.png.datadir`

Defines the file-system directory path for storing image files (Graphs).

Default value: /usr/share/irp/web/RRD

Recommended value: /usr/share/irp/web/RRD

5.2.37 `global.policies`

Defines the path to the Routing Policies (1.2.9) configuration file.

Default value: /etc/noction/policies.conf

Recommended value: /etc/noction/policies.conf

5.2.38 `global.policy.unpack_max`

Limits the amount of unpacked prefixes per policy.

The parameter protects against specifying filters that are too broad.

Possible values: 1-100000

Default value: 5000

5.2.39 `global.private_key`

Public/Private key pair used to interoperate between IRP instances for purposes such as Global Commit Control.

Default value: generated once during installation/upgrade

5.2.40 `global.public_key`


Public/Private key pair used to interoperate between IRP instances for purposes such as Global Commit Control.

Default value: generated once during installation/upgrade

5.2.41 `global.rd_rtt`

Defines the latency distances between routing domains in the format `rda:rdb:rtt` where `rda` is the id assigned to one routing domain, `rdb` is the id assigned to the second routing domain and `rtt` represents the round trip time in milliseconds between them. `rda` and `rdb` must be different and assigned to providers. IRP assumes that distance from `rda` to `rdb` is equal to distance from `rdb` to `rda`.

The parameter takes a collection of such triplets that will define all the available inter-datacenter links between routing domains.

 It is important that the RTT value between routing domains is accurate. In case the value differs significantly from the correct value IRP will make improvement decision based on incorrect information and it will make unnecessary global improvements that will reroute more traffic via inter-datacenter links.

Refer [peer.X.rd](#), [peer.X.flow_agents](#), [bgpd.rd_local_mark](#)

5.2.42 `global.rrd.age_max`

Defines the maximum trusted interface load data age (seconds)

Data older than this interval will not be trusted by IRP. If interface rate for each provider link has not been updated for a specified amount of time, then IRP behavior will be changed as follows:

- If provider's `limit_load` is set, no further Cost/Performance improvements will be performed to that provider
- Commit Control will not perform further in/out improvements for this provider

Possible values: 120-240

Default value: 120

Recommended value: 120. Should be increased only if frequent SNMP timeouts occur.

5.2.43 `global.rrd.datadir`

Defines the file system directory path for storing RRD database files.

Possible values: valid directory

Default value: `/var/spool/irp`

Recommended value: `/var/spool/irp`

5.2.44 `global.slave_management_interface`

Defines the management network interface for slave node in a failover configuration. In most cases it is the same as the probing interface. When failover is disabled ([global.failover](#)) the parameter is not used. See also [global.master_management_interface](#).

Possible values: any valid system network interface name

Default value: `eth0`

5.2.45 `global.slave_probing_interface`

Defines the probing network interface for the slave node in a failover configuration. In most cases it is the same as the management interface. When failover is disabled ([global.failover](#)) the parameter is not used.

See also [global.master_probing_interface](#).

Possible values: any valid system network interface name

Default value: `eth0`

5.2.46 `global.slave_rd`

Specifies the Routing Domain that hosts the slave node of IRP in a failover configuration. By default RD=1 hosts IRP nodes.

Possible values: 1-100

Default value: 1

Recommended value: 1

5.2.47 `global.user_directories_conf`

Defines the path to the User Directories configuration file.


Default value: `/etc/noction/user_directories.conf`

Recommended value: `/etc/noction/user_directories.conf`

5.3 API daemon settings

5.3.1 `apid.allowed_ips`

This parameter defines the list of IPv4/IPv6 prefixes or addresses that are allowed to make requests to Irpapid.

 By default, only server-local access to Irpapid is allowed. If there is a remote GMI installation, the IP address(es) of the GMI instance(s) should be added to this parameter.

Possible values: list of IPv4/IPv6 prefixes

Default value: `127.0.0.1 ::1`

5.3.2 `apid.grpc.port`

Specifies the TCP port on which the Irpapid is listening for interprocess communication.

Possible values: 1-65535

Default value: 7602

5.3.3 `apid.listen.master_ip`

Defines the IPv4/IPv6 address of the API. Allows IRP API calls to target another node on the network. If no value provided then localhost is used. When failover is enabled ([global.failover](#)) slave's listen IP ([apid.listen.slave_ip](#)) must be configured too.

Possible values: list of IPv4/IPv6 addresses

Default value: `::`

5.3.4 apid.listen.port

Defines the TCP port on which the irpapid is listening. If no value provided port 10443 is used.

Possible values: 1–65535

Default value: 10443

5.3.5 apid.listen.slave_ip

Defines the IPv4/IPv6 address of the API of the slave node in a failover configuration. Allows IRP API calls to target another node on the network. If no value provided then localhost is used. When failover is disabled ([global.failover](#)) slave's listen IP is not used.

See also [apid.listen.master_ip](#).

Possible values: list of IPv4/IPv6 addresses

Default value: ::

5.3.6 apid.log

Defines the file-system path to the irpapid log file.

Default value: /var/log/irp/irpapid.log

5.3.7 apid.log.level

Defines the logging level for the irpapid service.

Possible values: emerg, fatal, alert, crit, error, warn, notice, info, debug

Default value: info

Recommended value: info

5.3.8 apid.maxthreads

Defines the number of threads that irpapid is allowed to run. If no value provided 50 threads are used.

Possible values: 1–300


Default value: 50

5.3.9 apid.path.mtr

System path to MTR utility.

Possible values: /valid/path

Default value: /usr/sbin/mtr


 Default value varies depending on platform

5.3.10 apid.path.ping

System path to ping utility.

Possible values: /valid/path

Default value: /usr/bin/ping


 Default value varies depending on platform

5.3.11 apid.path.ping6

System path to ping for IPv6 utility.

Possible values: /valid/path

Default value: /usr/sbin/ping6


 Default value varies depending on platform

5.3.12 apid.path.traceroute

System path to traceroute utility.

Possible values: /valid/path


Default value: /usr/bin/traceroute

 Default value varies depending on platform

5.4 Bgpd settings


5.4.1 bgpd.as_path

Defines the way Bgpd handles the AS-PATH attribute in the outgoing announcements. The selected options are evaluated in the configured order and the first valid AS-PATH will be used.

 Note that IRP will refuse to announce improvements when all configured options fail to produce a valid AS-PATH. Option 0 allows announcements with empty AS-PATH but this is undesirable for other reasons noted below.

Keeping a correct AS-PATH can be a requirement for some NetFlow/sFlow processing since the provider AS or the destination AS for the corresponding flow record is taken from the BGP routing table.

Some installations use outgoing filters that allow empty AS-Path while redistributing iBGP routes to upstreams. In such cases, Bgpd must be configured not to advertise the improvements with an empty AS-PATH to prevent further redistribution of the improvements to upstream routers.

 The reconstructed AS-path does not always correspond to the actual BGP AS-path.

”Use AS-Path from BMP” option is tied to the `bgpd.improvements.remove.bmp_exact_aggregate` parameter. Must be enabled and placed first in the list of AS path restore priority options for the improvements to be removed when an exact size prefix gets withdrawn from BMP.

Refer also to `peer.X.ipv4.next_hop_as`, `peer.X.ipv6.next_hop_as`, `peer.X.aspath_for_ix`, `bgpd.as_path_borrowing`, `bgpd.improvements.remove.bmp_exact_aggregate`.


Examples:

`bgpd.as_path=2 3` - this will instruct Bgpd to take first path from the database (reconstructed AS-Path). If that path is empty, then an AS-path with the provider’s AS number and the prefix AS number should be composed.

`bgpd.as_path=3 0` - this will instruct Bgpd to compose an AS-path with the provider’s AS number and the prefix AS number. If AS-Path remains empty, the improvements with an empty AS-Path are announced.

Possible values: Space separated list of options in the order of preference


- 0 - Allow empty AS-PATH
- 2 - Use non-empty reconstructed AS-PATH (Announce AS-path reconstructed from traceroute)
- 3 - Reconstruct AS path with provider ASN and prefix origin ASN
- 4 - Use AS-Path from BMP
- 5 - Use AS-Path from BGP Alternative paths (RFC 7911)

 Second algorithm produces meaningful AS-PATH only when `explorer.trace.all` is enabled

Default value: 5 4 2 3

5.4.2 `bgpd.as_path_borrowing`

Allows BMP to borrow AS path from other Provider with the same autonomous system number.

 The parameter modifies only 4th algo of `bgpd.as_path` parameter.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.4.3 `bgpd.db.timeout.withdraw`

Defines the time period (in seconds) before prefixes are withdrawn from the routing tables, after a database failure. This allows BGP daemon to function independently for a period of time after the IRP database becomes inaccessible due to a failure or manual intervention.

Possible values: 600 – 21600(6 hours)


Default value: 14400

Recommended values: 3600–14400

5.4.4 bgpd.full_control

Sets default behavior of IRP in regards to inbound prefixes control and specifically:

- only announce improvements or
- only announce improvements and include all allowed providers or
- fully control inbound prefixes by always announcing to allowed providers.

 This system wide default behavior can be overridden at inbound prefix level by specifying desired parameter value for `inbound.rule.X.full_control`.

Refer to [inbound.rule.X.full_control](#) for details.

Possible values: 0 (Improvements), 1 (If improved), 2 (All)

Default value: 0

5.4.5 bgpd.improvements.remove.bmp_exact_aggregate

Removes an improvement when the exact prefix for such improvement is no longer seen in BMP and the AS-Path from BMP is in use.

See also [bgpd.improvements.remove.withdrawn](#), [bgpd.as_path](#).

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.4.6 bgpd.improvements.remove.hold_time

Defines the time interval (in seconds) for deleting (from the IRP database) the improvements affected by "Remove on next-hop eq" or "Remove on aggregate withdraw" conditions (see [bgpd.improvements.remove.next_hop_eq](#) and [bgpd.improvements.remove.withdrawn](#)).

This reduces the effects of route flapping to improvement cleanup.

Verification is performed on each BGP Scan, so that the values that are lower than the value of the `bgpd.scaninterval` parameter are not taken into account. The higher the values - the longer the time interval needed for improvements, which are still valid after aggregate route is withdrawn or on equal next-hop.

Bgpd does improvements check against the routers received and sent via iBGP in periodic time intervals. The process is referred to as the BGP Scan process.

Time interval between BGP Scannings can be configured in [bgpd.scaninterval](#).


Possible values: 1 - 1000 sec

Default value: 60

Recommended values: 30 - 120

5.4.7 bgpd.improvements.remove.next_hop_eq

Instructs Bgpd to remove a prefix from improvements when aggregate route's next hop has changed and points to the same next hop as the improvement.

 This parameter shouldn't be enabled when `bgpd.updates.split` is disabled in a multi-routing domain configuration.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

Recommended value: 0

5.4.8 `bgpd.improvements.remove.withdrawn`

Instructs Bgpd to remove prefix from improvements when aggregate is being withdrawn from the router. See also [bgpd.improvements.remove.bmp_exact_aggregate](#).

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

Recommended value: 1

5.4.9 `bgpd.improvements.strip_non_irp_communities`

Instructs Bgpd to exclude from Updates of IRP improvements other communities except those configured in IRP.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.4.10 `bgpd.log`

Defines the complete path to the Bgpd log file

Possible values: full path to log file

Default value: /var/log/irp/bgpd.log

5.4.11 `bgpd.log.level`

Defines the logging level for the Bgpd service.

Possible values: emerg, fatal, alert, crit, error, warn, notice, info, debug

Default value: info

Recommended value: info

5.4.12 `bgpd.mon.guardtime`

Defines the Bgpd Monitoring guard time. All the controlled IP addresses must respond within the specified amount of time, for the provider to be restored from the FAIL state. In FAIL state, all improvements are withdrawn from that provider.

It is recommended that [bgpd.mon.guardtime](#) is set as a double value of the [bgpd.mon.holdtime](#).

Possible values: 10-600

Default value: 60

Recommended value: 60

5.4.13 bgpd.mon.holdtime

Defines the Bgpd Monitoring holdtime.

If any controlled IP address does not respond to all the requests during the holdtime, then corresponding provider enters the FAIL state. In FAIL state, all the improvements are withdrawn from that provider.

Possible values: 10-60

Default value: 30

Recommended value: 30

5.4.14 bgpd.mon.internal.flap_guardtime

Defines time interval to protect from BGP session flapping.

Internal monitor will keep FAIL state until BGP session stays established within the configured period of time.

Possible values: 0-86400

Default value: 0

Recommended value: 600

5.4.15 bgpd.mon.keepalive

Defines the Bgpd Monitoring keepalive interval (in seconds) between consequent ICMP Echo Requests to single controlled IP address.

Possible values: 1-10

Default value: 10

Recommended value: 5

5.4.16 bgpd.mon.longholdtime

Defines the Bgpd Monitoring long holdtime.

If any controlled IP address does not respond to all the requests during the long holdtime, then corresponding provider enters the FAIL state. In FAIL state, all the improvements are withdrawn from that provider.

Possible values: 60-3600


Default value: 1800

Recommended value: 1800

5.4.17 bgpd.monitor.type

Defines how IRP monitors Improvements to Internet Exchanges:

- by using coarse-grained internal monitors that merely validate an IX peer is live or
- by using fine-grained prefix monitors for each IX improvement that validate that the IX peer still advertises the improved prefix.

 Prefix monitors might consume significant router CPU resources when relying on SNMP to determine if an IX peer advertises the improved prefix and the number of IX improvements is large.

Refer to [bgpd.prefix.monitor.interval](#) for details.

Possible values: 0 (Use internal monitor), 1 (Use prefix monitor)

Default value: 0

5.4.18 bgpd.no_export


Controls how Bgpd appends NO_EXPORT or NO_ADVERTISE communities to outbound improvements.

A router is responsible (as defined in [RFC 1997](#)) for preventing the utter route redistribution (NO_ADVERTISE) or the route advertisement outside of an autonomous system (NO_EXPORT).

It's strictly recommended to have this option turned on as an additional measure in preventing Outbound Improvement leakage outside of an optimized network.

One may consider disabling the feature in the following conditions:

- Different autonomous system numbers are used inside of an optimized network.
The feature may prevent redistribution of a Global improvement in MRD configuration.
- Downstream BGP peering.
The feature may prevent redistribution of routes to downstream, making the improved routes not visible to it.

 Disable with caution

Possible values: 0 Disabled

- 1 Append NO_EXPORT community to outbound improvements
- 2 Append NO_ADVERTISE community to outbound improvements

Default value: 1

5.4.19 bgpd.policy.cascade.amount

Defines the maximum number of downstream AS for cascading policies. If IRP identifies more downstream elements in AS-PATH from the designated AS the policy for those AS will not be enforced.


Possible values: 1-1000000

Default value: 1000

Recommended value: 1000

5.4.20 bgpd.prefix.monitor.interval

Prefix monitor tracks at given interval in seconds if a peering partner on an Exchange is still announcing the prefix that IRP improved through it.


 This is intended to avoid cases when after IRP makes an Improvement through a peer on an IX the peer stops announcing/servicing the route.

Possible values: 1-10

Default value: 10

5.4.21 bgpd.prefix.monitor.search_interval

Defines the interval between retries of prefix monitor failed initialization attempts in case a prefix isn't advertised by an IX peering partner or if a SNMP error occurs.


 During this time IRP will not be able to make improvements through the affected IX peers.

Possible values: 300–3600

Default value: 300

5.4.22 bgpd.prefixlist.asn

Specifies a collection of AS numbers that are analyzed for inbound optimization of transiting traffic. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

 IRP ignores prefixes /25 and shorter for IPv4 and /65 and shorter for IPv6 being present in network's routing table. This means that traffic belonging to these small prefixes are accounted under the immediately larger prefix that fits the above criteria.

Possible values: list of valid AS numbers

5.4.23 bgpd.prefixlist.prefixes

Specifies a collection of IPv4 or/and IPv6 prefixes in CIDR notation that are analyzed for inbound optimization of transiting traffic. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: list of valid CIDR prefixes

5.4.24 bgpd.rd_local_mark

Specifies a marker to distinguish local improvements from global improvements in the case of multiple routing domains optimization. Parameter represents a valid value for BGP community attribute of the form X:Y. Value in bgpd.rd_local_mark is APPENDED to communities attribute. Refer [global.rd_rtt](#), [peer.X.rd](#), [peer.X.flow_agents](#), [rd.X.community.local](#).


Possible values: X:Y

Default value: 65535:1

5.4.25 bgpd.retry_probing.new.bmp_path_change

Defines on what new provider AS Path changes to re-probe an already improved prefix. The possible options are:

- 0: Disabled
- 1: On major AS-Path change
- 2: On any AS-Path change

 Major AS Path changes are those traversing a different set of Autonomous Systems. AS Path changes such as prepended ASN are ignored when only major AS Path changes are monitored.

Refer also to [bgpd.retry_probing.old.bmp_path_change](#), [peer.X.bmp](#).


Possible values: 0 – 2

Default value: 0

5.4.26 `bgpd.retry_probing.old.bmp_path_change`

Defines on what old provider AS Path changes to re-probe an already improved prefix. The possible options are:

- 0: Disabled
- 1: On major AS-Path change
- 2: On any AS-Path change

 Major AS Path changes are those traversing a different set of Autonomous Systems. AS Path changes such as prepended ASN are ignored when only major AS Path changes are monitored.

Refer also to `bgpd.retry_probing.new.bmp_path_change`, `peer.X.bmp`.

Possible values: 0 - 2

Default value: 0

5.4.27 `bgpd.scaninterval`

The interval in seconds between the execution of the BGP Scan process.

BGP scan is used for:

- checking the improvements belonging to aggregated routes
- checking the improvements for **aggregate withdrawn(5.4.8)** and for **aggregate next-hop equal to improvements next-hop(5.4.7)** conditions
- checking the improvements for changed BGP attributes
- checking for changes for the improvements to be announced to/withdrawn from iBGP neighbors

Possible values: 10-100

Default value: 20

Recommended value: 60

5.4.28 `bgpd.snmp.concurrent_requests`

Specifies the number of allowed concurrent (in-flight) SNMP requests sent to a router.

Possible values: 1-2000

Default value: 10

5.4.29 `bgpd.snmp.packets_interval`

Time interval in milliseconds between transit improvement's monitor SNMP packets. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 1-100

Default value: 10

5.4.30 bgpd.snmp.simultaneous

The number of the simultaneous PDUs that will be contained in a SNMP request.

Possible values: 1–300

Default value: 10

Recommended value: 10

5.4.31 bgpd.transit.monitor.election_interval

Transit monitor election interval as a factor of reconfirm intervals. Refer [bgpd.transit.monitor.fast_reconfirm_interval](#).

Possible values: 1–10

Default value: 4

5.4.32 bgpd.transit.monitor.fast_reconfirm_interval

Transit monitor fast reconfirm interval in seconds. The reconfirm interval sets the periodicity by which transit monitors verify presence of alternative routes from other providers for transit improvements. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 1–60

Default value: 15

5.4.33 bgpd.transit.monitor.retries

Number of SNMP retries before a timeout of transit improvement's monitor. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 1–2000

Default value: 3

5.4.34 bgpd.transit.monitor.timeout

Timeout in milliseconds of individual SNMP requests used to monitor transit improvements. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 1–2000


Default value: 1000

5.4.35 bgpd.updates.split

Instructs Bgpd to split advertised prefixes in two equal parts (e.g /24 is split into two /25 prefixes).

This parameter should be enabled in order to preserve the original BGP UPDATE attributes received from the corresponding aggregates.

Refer to [global.aggregate](#), [global.agg_ipv4_max](#), [global.agg_ipv6_max](#) parameters.

 If this option is enabled, the number of announced prefixes will be twice the [core.improvements.max](#).


If the `bgpd.peer.X.updates.limit.max` parameter value is established, then the limitation is set on the total amount of announced prefixes, AFTER split. For example, if the value of `core.improvements.max` is set to 10000 and `bgpd.peer.X.updates.limit.max` is set to 5000, then the amount of the improvements towards this particular provider is no more than 2500, split onto 5000.


Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.5 BGP sessions settings

5.5.1 `bgpd.peer.X.as`

 Mandatory for each iBGP session definition.

 Parameter changes cause reset of BGP session.

Defines the Autonomous System Number for the iBGP session.

Possible values: 1-4294967295

5.5.2 `bgpd.peer.X.blackholing.ipv4.next_hop`

Defines IPv4 address which will be used as `next_hop` in BGP UPDATE for Blackholing routes sent to this router.

The next-hop address should be known by the router.

Possible values: IPv4 address

5.5.3 `bgpd.peer.X.blackholing.ipv6.next_hop`


Defines IPv6 address which will be used as `next_hop` in BGP UPDATE for Blackholing routes sent to this router.

The next-hop address should be known by the router.

Possible values: IPv6 address

5.5.4 `bgpd.peer.X.blackholing.localpref`


Defines localpref value used by IRP in BGP UPDATE for Blackholing routes sent to this router.

 Avoid collisions of localpref values assigned to IRP both within its configuration and/or on customer's network.

Possible values: 0 - 4294967295

Default value: 100

5.5.5 `bgpd.peer.X.cap_4byte_as`

 Parameter changes cause reset of BGP session.

Defines usage of 16 or 32-bit autonomous system numbers. Capability can be negotiated during session setup or forced to either 16 or 32 bits.

- 1 - Negotiated on OPEN
- 2 - Always 16-bit AS path
- 3 - Always 32-bit AS path

When the router is operating in legacy mode and does not negotiate capabilities but still sends 32-bit AS Path then Bgpd considers this a malformed AS_PATH attribute and disconnects the session. To avoid this the parameter must be set to force use of 32bit AS numbers.

For example: A BGP session with IRP with “disable-capability-negotiation” option configured on Vyatta router.

As a result, the BGP session is established and then teared down with log messages:

Listing 5.1: Error log

```
Jan 29 10:20:40.777965 WARN: BGP session RTR/IPv4 (10.0.0.1 AS 65530)
    Incoming UPDATE error: Invalid elements ignored. Malformed AS_PATH
Jan 29 10:20:40.778021 ERROR: BGP session RTR/IPv4 (10.0.0.1 AS 65530)
    Incoming UPDATE error: malformed AS_PATH xxxxxxxx
```

The solution is to set `bgpd.peer.RTR.cap_4byte_as = 3`, then the BGP session succeeds.

Possible values: 1 (Negotiate), 2 (Force 16bit), 3 (Force 32bit)

Default value: 1

5.5.6 bgpd.peer.X.flowspec

Enables/disables FlowSpec capability for BGP session.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.5.7 bgpd.peer.X.flowspec.redirect_type

The parameter defines the protocol dialect used to specify FlowSpec redirect.

Typically, Huawei routers use the IETF dialect, while Cisco routers use the Simpson dialect.

Nokia dialect has been implemented as described in the [Unicast Routing Protocols Guide Release 22.2.R1](#).

Possible values: 0 (IETF), 1 (Simpson), 2 (Nokia)

Default value: 1

5.5.8 bgpd.peer.X.inbound.ipv4.next_hop

Defines IPv4 address which will be used as next_hop in BGP UPDATE for inbound improvements/announcements towards this router.

The next-hop address should be known by the router (refer [Routers configuration for Inbound](#)).

An inbound rule can be configured with a rule-specific next-hop. Refer to [inbound.rule.X.next_hop](#).

Possible values: IPv4 address

5.5.9 bgpd.peer.X.inbound.ipv6.next_hop

Defines IPv6 address which will be used as next_hop in BGP UPDATE for inbound improvements/announcements towards this router.

The next-hop address should be known by the router (refer [Routers configuration for Inbound](#)).


An inbound rule can be configured with a rule-specific next-hop. Refer to [inbound.rule.X.next_hop](#).

Possible values: IPv6 address

5.5.10 bgpd.peer.X.inbound.master_localpref

Defines localpref value used by IRP for inbound improvements for this router.

Assigned localpref value should allow Inbound Improvement to become best route.

 Avoid collisions of localpref values assigned to IRP both within its configuration and/or on customer's network.


Possible values: 0 - 4294967295

Default value: 102

5.5.11 bgpd.peer.X.inbound.slave_localpref

Defines localpref value used by IRP for inbound improvements for this router.

Assigned localpref value should allow Inbound Improvement to become best route.

 Avoid collisions of localpref values assigned to IRP both within its configuration and/or on customer's network.

Possible values: 0 - 4294967295

Default value: 101

5.5.12 bgpd.peer.X.keepalive

Defines the session keepalive time (sec). See [RFC1771](#) for details. Hold time is calculated as keepalive * 3.

Possible values: 1-100

Default value: 60

Recommended value: 1/3 holdtime

5.5.13 bgpd.peer.X.listen

Instructs the BGP daemon to listen for incoming sessions. It must be set to 1 to be [RFC1771](#) compliant. It can be set to 0 to resolve specific issues.


Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

Recommended value: 1

5.5.14 bgpd.peer.X.master_communities

Defines the BGP Community that will be appended by Bgpd to all advertised prefixes. The format is: “X:Y”.


 Avoid colisions of communities values assigned to IRP both within its configuration and/or on customer’s network.

When failover is enabled ([global.failover](#)) slave’s BGP Community ([bgpd.peer.X.slave_communities](#)) must be configured too.


In case Bgpd receives the full or partial RIB (Routing Information Base), values for: **MED**, **Origin**, **LocalPref** and **Communities** are taken from a less-specific Aggregate route. If values for **MED**, **Origin**, **LocalPref** are set in config, it will override any value from Aggregate route. If value for **Communities** is set in config, it will be appended to communities from Aggregate route.

5.5.15 bgpd.peer.X.master_localpref

Defines the local-preference value for prefixes (improvements) announced by Bgpd.

 Avoid colisions of localpref values assigned to IRP both within its configuration and/or on customer’s network.

When failover is enabled ([global.failover](#)) slave’s BGP LocalPref ([bgpd.peer.X.slave_localpref](#)) must be configured too.


 If failover is enabled, master’s LocalPref value must be greater than slave’s LocalPref value.


In case Bgpd received the full or partial RIB (Routing Information Base), values for **MED**, **Origin**, **LocalPref** and **Communities** will be taken from less-specific Aggregate route. If values for **MED**, **Origin**, **LocalPref** are set in config, it will override any value from Aggregate route. If value for **Communities** is set in config, it will be appended to communities from Aggregate route.

Possible values: 0 – 4294967295

Default value:

5.5.16 bgpd.peer.X.master_our_ip


 Mandatory for IPv4 BGP session.


 Parameter changes cause reset of BGP session.

Defines IPv4 address of IRP server end of this iBGP session. When failover is enabled ([bgpd.peer.X.slave_our_ip](#)) slave’s IP address ([bgpd.peer.X.slave_localpref](#)) must be configured too.

Possible values: IPv4 address

5.5.17 bgpd.peer.X.master_our_ipv6

 Mandatory for IPv6 BGP session.

 Parameter changes cause reset of BGP session.

Defines IPv6 address of IRP server end of this iBGP session. When failover is enabled (`bgpd.peer.X.slave_our_ip`) slave's IPv6 address (`bgpd.peer.X.slave_our_ipv6`) must be configured too.


Possible values: IPv6 address


5.5.18 `bgpd.peer.X.master_password`

Defines iBGP session's password. When failover is enabled (`bgpd.peer.X.slave_our_ip`) slave's local IPv6 address (`bgpd.peer.X.slave_password`) must be configured too.

Possible values: up to 80 alphanumeric characters

5.5.19 `bgpd.peer.X.master_peer_ip`


 Mandatory for IPv4 BGP session.


 Parameter changes cause reset of BGP session.

Defines iBGP session's router's IPv4 address.

Possible values: IPv4 address

5.5.20 `bgpd.peer.X.master_peer_ipv6`


 Mandatory for IPv6 BGP session.


 Parameter changes cause reset of BGP session.

Defines iBGP session's router's IPv6 address.

Possible values: IPv6 address

5.5.21 `bgpd.peer.X.master_router_id`

 Mandatory for IPv6 BGP session either as standalone master or when failover is enabled and the value should be different from `bgpd.peer.X.slave_router_id`.

 Parameter changes cause reset of BGP session.

Defines IRP server's router ID. The BGP router ID is used in the BGP algorithm for determining the best path to a destination where the preference is given to the BGP router with the lowest router ID.

Possible values: 4-byte value in the IPv4 address format. Any valid IPv4 address can be used.

5.5.22 bgpd.peer.X.med

Defines the Multi-Exit Discriminator (MED) value for prefixes (improvements) announced by Bgpd.

In case Bgpd received the full or partial RIB (Routing Information Base), values for **MED**, **Origin**, **LocalPref** and **Communities** will be taken from less-specific Aggregate route. If values for **MED**, **Origin**, **LocalPref** are set in config, it will override any value from Aggregate route. If value for **Communities** is set in config, it will be appended to communities from Aggregate route.


Possible values: 0 - 4294967295

5.5.23 bgpd.peer.X.origin

Defines the Origin value for prefixes announced by Bgpd.

Possible values: 0 (IGP), 1 (EGP), 2 (INCOMPLETE)

5.5.24 bgpd.peer.X.shutdown

 Parameter changes cause reset of BGP session.

Defines whether the corresponding iBGP session is active or shutdown.

Possible values: 0 (Active), 1 (Shutdown)

Default value: 0

5.5.25 bgpd.peer.X.slave__communities

Defines the BGP Community that will be appended by Bgpd on the slave node of a failover configuration to all advertised prefixes. The format is: "X:Y".

See also [bgpd.peer.X.master__communities](#).

5.5.26 bgpd.peer.X.slave__localpref


Defines the local-preference value for prefixes announced by Bgpd.


See also [bgpd.peer.X.master__localpref](#).

Possible values: 0 - 4294967295

Default value:

5.5.27 bgpd.peer.X.slave__our_ip

 Mandatory for IPv4 BGP session.


 Parameter changes cause reset of BGP session. Affects only BGP session of slave node in a failover configuration.


Defines IPv4 address of IRP server end of this iBGP session of slave node in a failover configuration.

See also [bgpd.peer.X.master__our_ip](#).

Possible values: IPv4 address

5.5.28 `bgpd.peer.X.slave_our_ipv6`

 Mandatory for IPv6 BGP session in failover configuration.

 Parameter changes cause reset of BGP session. Affects only BGP session of slave node in a failover configuration.

Defines IPv6 address of IRP server end of this iBGP session of slave node in a failover configuration.

See also `bgpd.peer.X.master_our_ipv6`.

Possible values: IPv6 address

5.5.29 `bgpd.peer.X.slave_password`


Defines iBGP session's password for slave node in a failover configuration.


See also `bgpd.peer.X.master_password`.

Possible values: up to 80 alphanumeric characters

5.5.30 `bgpd.peer.X.slave_peer_ip`

Optionally distinct router IPv4 address may be specified in the parameter to be used in order to establish the BGP session from slave instance.

 Mandatory for IPv4 BGP session.

 Parameter changes cause reset of BGP session.


Defines iBGP session's router's IPv4 address.


See also `bgpd.peer.X.master_peer_ip`.

Possible values: IPv4 address

5.5.31 `bgpd.peer.X.slave_peer_ipv6`

Optionally distinct router IPv6 address may be specified in the parameter to be used in order to establish the BGP session from slave instance.

 Mandatory for IPv6 BGP session.


 Parameter changes cause reset of BGP session.


Defines iBGP session's router's IPv6 address.

See also `bgpd.peer.X.master_peer_ipv6`.

Possible values: IPv6 address

5.5.32 `bgpd.peer.X.slave_router_id`


 Mandatory for IPv6 BGP session either as standalone slave or when failover is enabled and the value should be different from `bgpd.peer.X.master_router_id`.

 Parameter changes cause reset of BGP session. Affects only BGP session of slave node in a failover configuration.

Defines IRP server's router ID. The BGP router ID is used in the BGP algorithm for determining the best path to a destination where the preference is given to the BGP router with the lowest router ID.

Possible values: 4-byte value in the IPv4 address format. Any valid IPv4 address can be used.

5.5.33 `bgpd.peer.X.transit.mib`


 Parameter changes cause reset of BGP session.

Defines the MIB used by transit improvements monitors. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

- 0 - Generic (BGP4-MIB)

Default value: 0


5.5.34 bgpd.peer.X.transit.snmp

 Parameter changes cause reset of BGP session.

Points to the SNMP host (and its parameters) used for transit improvements monitor on this session. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: valid SNMP host identifier

5.5.35 bgpd.peer.X.transit.status

 Parameter changes cause reset of BGP session.

Enables or disables transit improvements through this BGP session (router). Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.5.36 bgpd.peer.X.updates.limit.max

Represents a maximum number of prefixes that can be announced simultaneously in one session.

If [bgpd.updates.split](#) is ON, the number of announced prefixes is twice the number of the improvements. If the BGP neighbor (typically - the edge router) has any hardware/software limitation for the number of routes in active routing table, then Bgpd can be instructed not to announce more than a specified amount of prefixes. Value 0 means no limit for current iBGP session.

Values less than maximum allowed improvements, can cause not all the improved prefixes to be injected into such peer.

Higher values can be incompatible with the router (please consult the router's vendor regarding the maximum amount of entries in routing table as well as the BGP table).

Possible values: 0-100000

Default value: 0

Recommended value: 0

5.5.37 bgpd.peer.X.updates.limit.ps

Defines the maximum number of updates per second that will be sent to the current BGP neighbor. Low values will slow down the improvements injection. High values can cause router to drop the improvement without installing it into the routing database.

Possible values: 1-1000000

Default value: 500

Recommended values: 100-1000

5.6 BMP monitoring station settings

5.6.1 irpbmpd.grpc.port

Specifies the TCP port on which the Irpbmpd is listening for interprocess communication.

Possible values: 1-65535

Default value: 7603

5.6.2 irpbmpd.log

Defines the file-system path to the BMP monitoring station (irpbmpd) log file.

Default value: `/var/log/irp/irpbmpd.log`

5.6.3 irpbmpd.log.level

Defines the logging level for the BMP monitoring station (irpbmpd) service.

Possible values: `emerg, fatal, alert, crit, error, warn, notice, info, debug`

Default value: `info`

Recommended value: `info`

5.6.4 irpbmpd.port

Defines the TCP port where BMP monitoring station (irpbmpd) listens for monitoring routers to establish BMP sessions.

i When BMP information is available for all providers also consider setting BMP as a source of AS_PATH information under `bgpd.as_path`.

Default value: `7854`

Recommended value: `1-65535`

5.6.5 irpbmpd.sources

Defines IP addresses if the valid BMP source. BMP connections coming in from other IP addresses that are not listed in this parameter will be ignored.

It is recommended to specify only trusted IP addresses.

Possible values: list of valid IPv4/IPv6 addresses

Default value: `0.0.0.0/0 ::/0`

Recommended value: Trusted IPv4/IPv6 flow exporter addresses

5.7 Collector settings

5.7.1 collector.detect.explorer_ips

Instructs the collector to ignore the Explorer-generated traffic, which can initiate false IRP reaction to the network events (Excessive loss, blackout).

i This feature is used only by the SPAN collector, instructing it to ignore the IPv4 traffic when the network address translation is used to masquerade IP addresses.

Possible values: `0 (Disabled), 1 (Enabled)`

Default value: `0`

Recommended value: `0`

5.7.2 collector.export.archive_inbound

Enable Inbound traffic distribution history.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.7.3 collector.export.archive_transit

Enable Transit traffic distribution history.

Warning: may significantly increase disk consumption.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.7.4 collector.export.top_volume_ips

Defines the maximum number of top hosts per prefix for which usage values are stored. Collector persists up to this number of IP address, usage tuples together with other prefix statistics being collected. When the value of the parameter is zero no such statistics will be collected.

Possible values: 0-10

Default value: 5

5.7.5 collector.export.ttl

Defines the collector-gathered data lifetime (in seconds). Larger values lead to excessive database size. Aggregated data is kept for one year, disregarding this parameter.

Possible values: 1-'unlimited'

Default value: 86400

Recommended value: 1-30days

5.7.6 collector.export.volume.high.top_n

Defines the number of top volume prefixes.

Top N prefixes will be marked for priority probing in descending order of volume. Lower values lead to small number of events for priority probing of high volume prefixes. Higher values lead to overflow of probing queue with jobs for unimportant prefixes.

Possible values: 0-'unlimited'

Default value: 50

Recommended value: 10-50

5.7.7 collector.export.volume.min

Defines the minimum collected prefix volume (bytes). The prefix will not be exported into the IRP database if its traffic volume is less than the value of the current parameter ([collector.export.volume.min](#)) and the percentage of the traffic volume less the [collector.export.volume.min_pct](#) parameter.

Lower values will lead to a higher number of prefixes exported into the IRP database.

Possible values: 1-2000000000

Default value: 100000

Recommended value: 10000-1000000

5.7.8 collector.export.volume.min_pct

Defines the minimum prefix volume (%) for a prefix to be exported into the IRP database. The prefix will not be exported into the IRP database if its percentage of the traffic volume is less than the value of the current parameter ([collector.export.volume.min_pct](#)) and the traffic volume less than [collector.export.volume.min](#). The percentage of the traffic volume is calculated according to the export interval (60 seconds). Lower values will lead to a higher number of prefixes exported into the IRP database.

Possible values: 0.0001-100

Default value: 0.01

Recommended value: 0.01-0.1

5.7.9 collector.flow.all_outbound

The parameter allows Irpflowd to process all outbound traffic even not listed in [collector.ournets](#).

Enabling the feature requires [peer.X.flow_agents](#) to be properly configured for all providers to distinguish outbound direction.

Possible values: 0 (Ournets only), 1 (All traffic)

Default value: 0

5.7.10 collector.flow.buffer.size

Defines the buffer size (in packets) for the Flow collector (irpflowd). Higher values lead to extra memory used by the collector. Slightly increase this value if the network has traffic spikes that cause buffer overflows and packet drop in the collector.

Possible values: 10000-1000000

Default value: 50000

Recommended value: 50000-500000

5.7.11 collector.flow.enabled

Enables the Flow collector. If the parameter is enabled, Flow Collector will be used to gather network prefix data, but it is still possible to use SPAN Collector to acquire network events.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

Recommended value: 0 - for SPAN collector (see [collector.span.enabled](#)), 1 - for Flow collector.

5.7.12 collector.flow.export.inbound_transit.topn

Specifies the number of largest volume transit prefixes that are exported each collector cycle for inbound optimization of transiting traffic. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 1-10000

Default value: 100

5.7.13 collector.flow.listen.nf

Defines the UDP port the Flow collector will listen to, for NetFlow traffic.

Possible values: 1–65535

Default value: 2055

Recommended value: 2055

5.7.14 collector.flow.listen.sf

Defines the UDP port the Flow collector will listen to, for sFlow traffic.

Possible values: 1–65535

Default value: 6343

Recommended value: 6343

5.7.15 collector.flow.log

Defines the file-system path to the irpflowd log file.

Default value: /var/log/irp/irpflowd.log

5.7.16 collector.flow.log.level

Defines the logging level for the irpflowd service.


Possible values: emerg, fatal, alert, crit, error, warn, notice, info, debug

Default value: info

Recommended value: info

5.7.17 collector.flow.process_transit_in_outbound

Enables or disables matching of prefix traffic at network egress for inbound optimization of transiting traffic. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

 This should be enabled only in complex network topologies where ingress prefix statistics can be incomplete.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.7.18 collector.flow.sources

Defines the valid NetFlow/sFlow/jFlow exporters IP addresses. Flow data coming in from other IP addresses that are not listed in this parameter will be ignored.

It is recommended to specify only trusted Flow exporters IP addresses.

Possible values: list of valid IPv4/IPv6 addresses

Default value: 0.0.0.0/0 ::/0

Recommended value: Trusted IPv4/IPv6 flow exporter addresses

5.7.19 collector.ournets

 Mandatory if `collector.span.enabled` is enabled

Defines the list of networks to be analyzed. Typically this is the list of prefixes advertised by your AS.
Format:

1. Space-separated list of local IPv4/IPv6 prefixes that should be analyzed and optimized by the IRP.
2. Absolute path to a newline-separated file containing a list of local IPv4/IPv6 prefixes that should be analyzed and optimized by the IRP.

IPv4 prefixes are treated as /24 subnets and IPv6 prefixes as /48 subnets, if netmask is not clearly specified.

 The downstream clients' networks can be indicated as well.

Possible values: IPv4 or IPv6 prefixes

5.7.20 collector.sessions.max

Defines the maximum number of TCP sessions inside the collector process. It can be used to minimize memory usage.

 Both Flow and Span collectors use this parameter.

Estimated memory usage can be calculated as follows: $\text{usage} = 80\text{MB} + (\text{collector.sessions.max} * 1464)$

Possible values: 10000–10000000

Default value: 2000000

Recommended value: Depends on available server memory and the maximum number of simultaneous sessions during peak hours.

5.7.21 collector.snmp.enhanced.sec

Defines during what second of a minute the SNMP enhanced algorithm should run.

Possible values: 30–58

Default value: 45

Recommended value: 45

5.7.22 collector.span.buffer.size

Defines the buffer size (in packets) for the Span collector (irpspand). Higher values lead to extra memory used by the collector. Slightly increase this value if the network has traffic spikes that cause buffer overflows and packet drop in the collector.

Possible values: 10000–5000000

Default value: 100000

Recommended value: 100000–5000000

5.7.23 collector.span.enabled

Enables the Span collector. Span collector will be used to gather network prefix data, if the parameter is enabled.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

Recommended value: 0 - if no traffic analysis should be performed by the Span collector, 1 - if the Span collector should analyze mirrored traffic for network events and/or prefix statistics.

See also: [collector.flow.enabled](#)

5.7.24 collector.span.interfaces

 Mandatory if [collector.span.enabled](#) is enabled

Defines a space-separated network interfaces list for passive packet analysis by the Span collector.

Example:

```
collector.span.interfaces = eth0 eth1 eth2
```

5.7.25 collector.span.log

Defines the file-system path to the irpspand log file.

Default value: /var/log/irp/irpspand.log

5.7.26 collector.span.log.level

Defines the logging level for the irpspand service.

Possible values: emerg, fatal, alert, crit, error, warn, notice, info, debug

Default value: info

Recommended value: info

5.7.27 collector.span.min_delay

Enables fast probing tasks queuing for prefixes with one of the following issues:

- Blackouts
- Congestion
- Excessive packet delays.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

Recommended value: 0

See also: [collector.span.min_delay.probing_queue_size](#)

5.7.28 `collector.span.min_delay.probing_queue_size`

Defines the number of slots in the probing queue that can be used by the `min_delay` algorithm (see: [collector.span.min_delay](#))

Possible values: 1–200


Default value: 50

Recommended value: 30–200

See also: [collector.span.min_delay](#)

5.7.29 `collector.span.size_from_ip_header`

When parameter is enabled packet size is determined from header of the IPv4/IPv6 packet. Otherwise, packet size is determined from link layer information (original packet size from PCAP/SNF).

 Enable this when source packets are stripped before entering Noction IRP's appliance SPAN interface.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.7.30 `collector.span.threshold.blackout`

Defines the percentage of retransmitted packets in regards to the total packets number. If retransmit is higher than this value, this prefix is considered to have a blackout.

Possible values: 1–100

Default value: 70

Recommended value: 70–90

5.7.31 `collector.span.threshold.congestion`

Defines the percentage of retransmitted packets in regards to the total packets number. If retransmit is higher than this value, this prefix is considered to have a congestion.

Possible values: 1–100

Default value: 50

Recommended value: 30–70

5.7.32 `collector.span.threshold.delay`

Percentage of excessive delayed packets by the total packets number (see [collector.span.threshold.excessive](#)). If this percentage is higher than this value, we consider this prefix to have a delay.

Possible values: 1–100

Default value: 20

Recommended value: 10–30

5.7.33 collector.span.threshold.excessive

Defines the excessive RTT (%). The value is compared to the average round trip time. If RTT is higher by `collector.span.threshold.excessive` in % than the average RTT, then the counter of excessive delay packets is incremented.

Possible values: > 0

Default value: 200

Recommended value: 100-500

5.7.34 collector.speaking_ips

Defines the number of speaking IPs to be stored in the database.

Possible values: 0-1000


Default value: 100

Recommended value: 100

5.8 Core settings

5.8.1 `core.circuit.high_loss_diff`

Defines the upper threshold of consistent packet loss over a provider when compared to other providers on the network. A provider exceeding this threshold is marked for shutdown. An event of this type will be raised and available to subscribers to act upon. Threshold loss difference is determined over a configured past time horizon and compared with all other providers on the network over the same interval.


 While the provider is marked for shutdown IRP cannot do this itself. Instead network engineers and/or other network capabilities are expected to be triggered in order to divert as much traffic as possible away from provider with circuit issues.

Possible values: 2–50

Default value: 15

5.8.2 `core.circuit.hist_interval`

Defines time interval in minutes used to determine average packet loss over a provider when compared to other providers on the network for circuit issues detection algorithm.

 Keep in mind that shorter time intervals might cause false positives where the averages are high simply because a few large and random packet loss probes are able to push the numbers above thresholds. At the same time longer time intervals will take longer to spot issues and thus will extend the time period where a circuit with issues is being used while alternatives were available.

Possible values: 1–240

Default value: 5

5.8.3 `core.circuit.inbound`

Defines if and when IRP announces Max prepends for inbound prefixes through provider with circuit issues. The options are as follows:

- No changes - any prepends through provider are not changed
- Prepend on warn - IRP announces Max Prepends once Warning level for circuit issues is reached
- Prepend on shutdown - IRP announces Max Prepends only when shutdown level for circuit issues is exceeded.

Refer also [core.circuit.transit](#).

Possible values: 0 (No changes), 1 (Prepend on warn), 2 (Prepend on shutdown)

Default value: 0

5.8.4 `core.circuit.recover_hold_time`

Defines the time interval in seconds before IRP attempts to restore a circuit with issues.

Possible values: 60–3600

Default value: 600

5.8.5 `core.circuit.recover_loss_diff`

Defines the normal threshold of packet loss when a provider with circuit issues can be considered to be ok. At this stage IRP will restore the provider to its full capacity status and will retract all other measures taken in the past in order for the network traffic to avoid the circuit with issues.

❌ This parameter is only used if the ways to react to a circuit issue include restoring it and this only can happen within the given time interval. Otherwise the circuit should be restored by manual intervention of network engineers after they have verified the issue is no longer a problem.

ℹ Note that this parameter should be both smaller than `core.circuit.high_loss_diff` and `core.circuit.warn_loss_diff`

Possible values: 0-48

Default value: 5

5.8.6 `core.circuit.recover_monitored_intervals`

Defines the time interval in minutes during which IRP will continuously evaluate average loss for provider(s) with circuit issues in order to determine if it is back to normal.

❌ This parameter is only used if the ways to react to a circuit issue include restoring it and this only can happen within the given time interval. Otherwise the circuit should be restored by manual intervention of network engineers after the circuit issue is no longer a problem.

Possible values: 1-30

Default value: 5

5.8.7 `core.circuit.transit`

Defines if and when IRP announces Max preponds for transit prefixes through provider with circuit issues. The options are as follows:

- No changes - any preponds through provider are not changed
- Prepend on warn - IRP announces Max Preponds once Warning level for circuit issues is reached
- Prepend on shutdown - IRP announces Max Preponds only when shutdown level for circuit issues is exceeded.


Refer also `core.circuit.inbound`.

Possible values: 0 (No changes), 1 (Prepend on warn), 2 (Prepend on shutdown)

Default value: 0

5.8.8 `core.circuit.warn_loss_diff`

Defines the lower threshold of packet loss when a provider seems to be having circuit issues. At this stage IRP will start raising alerts that network engineers and/or network management systems can subscribe in order to act on them. At this level IRP starts taking other preventive measures such as re-probing outbound improvements made to provider with circuit issues.

 Note that this parameter should be both smaller than `core.circuit.high_loss_diff` and larger than `core.circuit.recover_loss_diff`.

Possible values: 1-49

Default value: 10

5.8.9 `core.circuit.withdraw_on_warn`


Enables or disables withdrawing outbound improvements to a provider with circuit issues at or above warning level.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.8.10 `core.commit_control`

Enables or disables Commit Control (see the [Commit Control](#) section)

 The following parameters must be set in the configuration file to ensure proper functionality of the Commit Control feature:

1. SNMP parameters must be set for each provider:
 - (a) `SNMP Host`
 - (b) `peer.X.snmp.interfaces`
2. Each provider needs to have `peer.X.95th` set to the desired Commit level
3. `peer.X.precedence` must be set for each provider

When Commit Control is disabled, all existing Commit Control improvements are removed from current improvements. The improvements are preserved temporarily until `core.commit_control.probe_ttl` after Core service restart. If Commit Control is re-enabled before expiration these improvements will be restored into current improvements.

See also:

[Commit Control](#)

`core.commit_control.rate.group`

`core.commit_control.rate.high`


`core.commit_control.rate.low`

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.8.11 `core.commit_control.agg_bw_min`

Defines the minimum bandwidth for a single prefix. Prefixes, whose bandwidth is less than the specified value, are ignored by Commit Control and are not being used in the Commit Control algorithm.

 It is not recommended to set high values for this parameter since it limits the number of prefixes that can be rerouted by the Commit Control mechanism.

i This parameter is considered only during initial improvement. During retry probing of existing Commit Control improvements IRP might detect that the current bandwidth usage for some prefixes is below this limit. IRP will preserve these improvements as relevant if there are no other reasons to withdraw them.

Its recommended to decrease value if summary throughput is lower than 200-500 Mbps.

Possible values: 0.1-5000

Default value: 1

Recommended value: 1

5.8.12 `core.commit_control.del_irrelevant`

If enabled, Commit Control algorithm deletes improvements that have traffic volume less than value in the [5.8.11](#) parameter.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

Recommended value: 1

5.8.13 `core.commit_control.group_balance.check_monitor`

Defines how monitors affect group load balancing.

Group load balance will include providers in a group only when a provider has no failed monitors.

0 Ignores monitors

1 Excludes providers with failed IPv4 monitor

2 Excludes providers with any failed monitor

Possible values: 0 - 2

Default value: 0

5.8.14 `core.commit_control.inbound.enabled`

Enables or disables Inbound bandwidth control.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.8.15 `core.commit_control.inbound.improvement.delay`

The parameter specifies minimal time delay before next Inbound/Transit optimization cycle.

Time delay should be sufficient to cover route propagation time and time required to collect and process changes in bandwidth distribution.

Possible values: 60 - 1800

Default value: 300

5.8.16 `core.commit_control.inbound.moderated`

Enables or disables review and moderate feature of inbound bandwidth control.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.8.17 `core.commit_control.inbound.rate.high`

Defines provider's high load rate limit (%). Refer [core.commit_control.rate.high](#) for details.

This limit is used when inbound and outbound commit control operate independently. Otherwise [core.commit_control.rate.high](#) value overrides this. Refer also [peer.X.95th.mode](#).

Possible values: 50 - 99

Default value: 90

5.8.18 `core.commit_control.inbound.rate.low`

Defines provider's low load rate limit (%). Refer [core.commit_control.rate.low](#) for details.

This limit is used when inbound and outbound commit control operate independently. Otherwise [core.commit_control.rate.low](#) value overrides this. Refer also [peer.X.95th.mode](#).

Possible values: 30 - 99

Default value: 70

5.8.19 `core.commit_control.inbound.volume_estimation`

Defines method of estimating inbound bandwidth. The available options are:

- 0: Last minute data
- 1: Largest of last minute and current hour average
- 2: Largest of last minute and daily average

Possible values: 0 - 2

Default value: 1

5.8.20 `core.commit_control.loss_override`

Defines when IRP allows Commit Control improvements to adjust a provider's bandwidth considering current and new provider's loss measurements. The available options are to allow Commit Control improvements when there is:

- 0: Better or equal loss
- 1: Better or irrelevant loss difference
- 2: Any loss difference

Refer [core.performance.loss_pct](#) for details.

Possible values: 0 - 2

Default value: 0

5.8.21 `core.commit_control.probe_ttl`

Defines the TTL (time-to-live) for a specific probe. If the probe results are older than the value specified here, the system will disregard it and the Commit Control algorithm will schedule it for expedited re-probing.

Possible values: 600–86400

Default value: 7200

Recommended value: 7200

5.8.22 `core.commit_control.probing_queue_size`

Defines the number of slots in the probing queue that can be used by the Commit Control algorithm. This sets the upper limit on the number of prefixes that can be scheduled for probing by Commit Control. Note that this queue has higher priority than ordinary and retry probing. At the same time Commit Control relies heavily on existing probing results and most of the time this queue will be empty.

Possible values: 1–500

Default value: 100

Recommended value: 10–100

5.8.23 `core.commit_control.rate.group`

If using provider load balancing in group, this parameter defines allowed deviation of a provider's current bandwidth(in %) compared with other providers in the same group.

Example 1. There are three grouped providers with 95th set to 1 Gbps, 2 Gbps and 3 Gbps with a total current bandwidth of 600 Mbps. In this case, load balancing expects that each provider's bandwidth usage will be 100 Mbps, 200 Mbps and 300 Mbps accordingly. These values are proportional to individual provider's 95th settings in the group. Let's assume `core.commit_control.rate.group` is set to 5%. If a provider's bandwidth exceeds by more than the 5% limit the expected value (105 Mbps, 210 Mbps and 315 Mbps accordingly and the total for the group did not change), IRP will start active rerouting of excessive bandwidth from that provider to providers within or outside this group.

 Load balancing in a provider group is performed even when the 95th is not exceeded.

Possible values: 1–30

Default value: 5

Recommended value: 5

See also: [Provider load balancing](#)

5.8.24 `core.commit_control.rate.high`

Defines provider's high load rate limit (%).

IRP will stop Latency/Cost improvement if provider bandwidth is over `core.commit_control.rate.high` % of load (percents of 95th). The improvements will start happening again after provider's bandwidth drops below `core.commit_control.rate.low` % of load.

These parameters are used for passive 95th overload prevention as well as an additional method for Commit Control to be less aggressive.

 Provider's Latency/Cost improvement ability does not change when current bandwidth rate varies between `core.commit_control.rate.low` and `core.commit_control.rate.high`.

Example 2. if provider's current load is 1100 Mbps, 95th is set to 1000 Mbps. Commit Control will try to unload 200Mbps (to reduce current load to 90% of the 95th level) in order to prevent 95th overload as well as allow provider's bandwidth natural growth during peak hours.

See also: [Provider load balancing](#)

Possible values: 50–99, but larger or equal to `core.commit_control.rate.low`

Default value: 90


Recommended value: 90

5.8.25 `core.commit_control.rate.low`

Defines provider's low load rate limit (%).

IRP will start Latency/Cost improvements again after provider's bandwidth drops below `core.commit_control.rate.low` % of load. Latency/Cost improvements will be stopped if provider bandwidth is over `core.commit_control.rate.high` % of load (percents of 95th).

These parameters are used for passive 95th overload prevention as well as an additional method for Commit Control to be less aggressive.

 Provider's Latency/Cost improvement ability does not change when current bandwidth rate varies between `core.commit_control.rate.low` and `core.commit_control.rate.high`.

See also: [Provider load balancing](#)


Possible values: 30–99, but lower or equal to `core.commit_control.rate.high`

Default value: 80

Recommended value: 80

5.8.26 `core.commit_control.react_on_collector`

Defines if Commit Control algorithm should react immediately on collector overload values.

 Enabling this feature represents a tradeoff. Take into consideration the benefits of a faster react time vs reliance on older probes and statistics when making Commit Control improvements.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

Recommended value: 1

5.8.27 `core.commit_control.worst_loss`

Defines amount of loss worsening allowed when IRP performs Commit Control improvements and `core.commit_control.loss_override` is set to "Any loss difference".

Refer `core.performance.loss_pct` for details.

Possible values: 1 – 99


Default value: 2

5.8.28 `core.cost.worst_ms`

Latency worsening.

Defines the allowed latency worsening for an improved prefix while running in [Cost optimization](#) mode (see [global.improve_mode](#)) and the [Commit Control](#) feature is enabled. While operating in these modes IRP will consider as alternatives candidates with latency degradation not exceeding this limit. Of course, the candidate with least latency degradation or even with better latency will be chosen as an improvement.

When IRP detects that an existing Commit Control improvement has alternatives with latencies better than specified value it will replace it with a new performance (latency) improvement. Over-usage of the alternative route will be avoided.

 This parameter is applicable for local improvements within a Routing Domain. Global improvements will also take into account [core.global.worst_ms](#) values.

Possible values: 1-unlimited

Default value: 10

Recommended value: 10

5.8.29 `core.eventqueuelimit`

Defines the exploring events queue size for IPv4 prefixes.

Lower values may result in IRP missing some important network issues.

Possible values: 1-10000

Default value: 1000

Recommended value: 500-2000

5.8.30 `core.eventqueuelimit.retry_probe_pct`

Defines a percentage of the total exploring queue length (see [core.eventqueuelimit](#)) to be used for retry probing.

Possible values: 1-100

Default value: 40

Recommended value: 20-60

5.8.31 `core.eventqueuelimit_ipv6`

Defines the exploring events queue size for IPv6 prefixes.

Lower values may result in IRP missing some important network issues.


Possible values: 1-10000

Default value: 1000

Recommended value: 500-2000

5.8.32 `core.flowspec.max`

Defines the maximum number of IPv4 Flowspec rules that IRP is allowed to advertise.


 Routers have a low limit of FlowSpec entries. If the limit is exceeded, the router behavior may be unpredictable.

Possible values: 1-20000

Default value: 100

5.8.33 `core.flowspec.max_ipv6`

Defines the maximum number of IPv6 Flowspec rules that IRP is allowed to advertise.

 Routers have a low limit of FlowSpec entries. If the limit is exceeded, the router behavior may be unpredictable.

Possible values: 1-20000

Default value: 100

5.8.34 `core.global.allow_commit`

Enables or disables global commit control Improvements in a multiple routing domain configuration. By default these Improvements are enabled.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.8.35 `core.global.allow_latency_cost`

Enables or disables latency and cost global Improvements in a multiple routing domain configuration. By default these Improvements are enabled.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.8.36 `core.global.worst_ms`

Defines acceptable latency worsening for cost and commit for global improvements. Refer [core.cost.worst_ms](#)

Possible values: 1-unlimited

Default value: 30

Recommended value: 10

5.8.37 `core.improvements.clearslots`

Specifies the number of outdated improvements that will be discarded when slots are needed for new improvements. Outdated are improvements as defined by [core.improvements.clearslots.days_max](#). In case there are no outdated improvements usual improvement replacement mechanisms are employed relying on [Improvements weight](#).

Possible values: 0 - 100

Default value: 10

Recommended value: 10

5.8.38 `core.improvements.clearslots.days_max`

Sets the number of days after which an improvement is considered outdated. Outdated improvements have not been re-confirmed for a long time. The oldest of them will be discarded by the slot clearing process for outdated improvements when new slots are needed for new improvements. Refer [core.improvements.clearslots](#).


Possible values: 1 – 60

Default value: 7

Recommended value: 7

5.8.39 `core.improvements.inbound_transit.max`

Specifies the maximum number of transit improvements. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

 Increasing this limit should be done with care. Each transit improvement is continuously monitored if alternative routes from providers are still present on the router(s) and this consumes router CPU resources.


Possible values: 1–1000

Default value: 100

5.8.40 `core.improvements.inbound_transit.ttl.clean`

Specifies how IRP should treat old transit improvements. Old transit improvements are those that exceed [core.improvements.inbound_transit.ttl.max](#). The options are to either

- gradually decrease the number of preprends and withdraw the transit improvement when no preprends are left or
- withdraw the inbound transit improvement immediately when it exceeds [core.improvements.inbound_transit.ttl.max](#)

 Transit improvements cleanup performed once a day at configured off-peak hour ([global.offpeak_hour](#)).

Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 0 (Decrease preprends), 1 (Remove immediately)

Default value: 0

5.8.41 `core.improvements.inbound_transit.ttl.max`

Specifies the maximum period of time to preserve a transit improvement. If IRP does not have any other reasons to adjust a transit improvement after this limitation is exceeded its preprends will be decreased or the improvement will be withdrawn. Refer [core.improvements.inbound_transit.ttl.clean](#), [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 901 – 345600

Default value: 86400

5.8.42 `core.improvements.inbound__transit.ttl.min`


Specifies the minimal period of time to preserve a transit improvement. In order to give sufficient time to the entire Internet to adjust to a transit improvement and to avoid route instability IRP blocks making changes to a transit improvement for this duration of time. IRP can still offer route changes to the same transit prefixes routed through other providers. Refer [Optimization of transiting traffic](#), [Optimization of transiting traffic](#).

Possible values: 600 – 86400

Default value: 14400

5.8.43 `core.improvements.max`

Defines the maximum number of active IPv4 improvements.

 The number of announced prefixes can be twice this value if `bgpd.updates.split` is enabled.

Possible values: 1–100000

Default value: 10000

Recommended value: 10000

5.8.44 `core.improvements.max__ipv6`

Defines the maximum number of active IPv6 improvements.

Possible values: 1–100000

Default value: 2000

Recommended value: 2000

5.8.45 `core.improvements.safe__removal`

Specifies when IRP removes/withdraws outbound improvements. Outbound improvements become irrelevant if old provider performance metrics are no longer worse than the improved provider metrics. Still, original routes for some improvements change and providers that will service those prefixes once IRP improvements are withdrawn might have worse performance metrics. Safe improvement removal accounts for this possibility and removes improvements only when ALL providers have performance metrics that are not worse than the (fresh) improvement provider metrics.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.8.46 `core.improvements.ttl.retry__probe`

Defines the retry probing interval (in seconds). Once an improvement is older than this value, it will be sent for re-probing.

Possible values: 600–86400

Default value: 14400

Recommended value: 7200–14400

5.8.47 `core.log`

Defines the file-system path to the core log file.

Default value: `/var/log/irp/core.log`

5.8.48 `core.log.level`

Defines the logging level for the core service.

Possible values: `emerg`, `fatal`, `alert`, `crit`, `error`, `warn`, `notice`, `info`, `debug`

Default value: `info`

Recommended value: `info`

5.8.49 `core.outage_detection`

Enables the Outage Detection algorithm.

Possible values: `0` (Disabled), `1` (Enabled)

Default value: `0`

Recommended value: `1`

5.8.50 `core.outage_detection.limit_pct`

Defines the problem prefixes threshold (in percent) over which the system confirms an outage (if `core.outage_detection` is enabled).

See also: [Outage detection](#)

Possible values: `1-100`

Default value: `60`

Recommended value: `60`

5.8.51 `core.outage_detection.volume_top_n`

Top-N relevant volume prefixes

The outage detection mechanism considers only those prefixes in the "Top N by volume." Separately for IPv4 and IPv6.

Possible values: `1-20000`

Default value: `5000`

5.8.52 `core.overusage.check_interval`

Defines the interval to check overusage for Flowspec policies prefixes in order to apply automatic throttling Flowspec policies.

Possible values: `60-600`

Default value: `60`

Recommended value: `60`

5.8.53 `core.overusage.hold_timer`

Defines the retention time for throttling rules after prefix average bandwidth returns to normal for automatic throttling Flowspec policies.

Possible values: 60–600

Default value: 60

Recommended value: 300

5.8.54 `core.overusage.out.average.period`

Defines the duration of the time interval in hours used to determine average prefix usage for automatic throttling Flowspec policies.

Possible values: 1–24

Default value: 1

Recommended value: 1

5.8.55 `core.overusage.out.average.relevant_min`

Defines minimum prefix bandwidth average (in Mbps) that is relevant for automatic throttling Flowspec policies.

Possible values: 1–100000

Default value: 50

Recommended value: 50

5.8.56 `core.overusage.out.threshold.throttle`

Defines the multiplier of prefix average bandwidth applied on outbound traffic for automatic throttling Flowspec policies.

Possible values: 2–10000

Default value: 2

Recommended value: 2

5.8.57 `core.overusage.out.threshold.trigger`

Defines the multiplier of prefix average bandwidth that sets the overusage threshold of a prefix. An automatic throttling Flowspec policy is created when current prefix bandwidth exceeds the average multiplied by this multiplier.

Possible values: 2–10000

Default value: 10

Recommended value: 10

5.8.58 `core.performance.loss_pct`

Defines the relevant packet loss difference (in percent) after which a loss improvement is made.

For current route loss > 50%: Do not perform rerouting, if packet loss difference between routes is less than 15%.

For current route loss <= 50%: Do not perform rerouting, if packet loss difference between routes is less than `core.performance.loss_pct`.

Possible values: 1-15

Default value: 3

Recommended value: 3-5

5.8.59 `core.performance.rtt.diff_ms`

Defines the relevant RTT difference (ms) after which a latency improvement is made. If the RTT difference between the current route and another provider is less than `core.performance.rtt.diff_ms`, then the system ignores this prefix as an improvement candidate.

 `core.performance.rtt.diff_ms` and `core.performance.rtt.diff_pct` are grouped together at decision making, using a logical AND condition.


Possible values: 1-'unlimited'

Default value: 10

Recommended value: 2-20

5.8.60 `core.performance.rtt.diff_pct`

Defines the relevant RTT difference (in percent) after which a latency improvement is made. If the RTT difference between the current route and another provider is less than `core.performance.rtt.diff_pct`, then the system ignores this prefix as an improvement candidate.

 `core.performance.rtt.diff_ms` and `core.performance.rtt.diff_pct` are grouped together at decision making, using a logical AND condition.


Possible values: 1-100


Default value: 10

Recommended value: 5-20

5.8.61 `core.performance.rtt.ix_diff_ms`

Defines the relevant RTT difference (ms) to make latency improvements across Internet Exchanges and transit providers. A latency improvement from an IX to a transit provider is allowed only when the latency is improved by at least this threshold. Inversely, a latency improvement from a transit provider to an IX is allowed even if latency degradation is less than this threshold.

 A default value of zero for this parameter disables this feature. A value smaller than or equal to `core.performance.rtt.diff_ms` isn't allowed. When the feature is disabled then the `core.performance.rtt.diff_ms` and `core.performance.rtt.diff_pct` parameters are considered to determine relevancy of latency improvements.

 `core.performance.rtt.ix_diff_ms` and `core.performance.rtt.ix_diff_pct` are grouped together at decision making, using a logical AND condition.


Possible values: 1-1000000

Default value: 0

Recommended value: 20-50

5.8.62 `core.performance.rtt.ix_diff_pct`

Defines the relevant RTT difference (in percent) to make latency improvements across Internet Exchanges and transit providers. Refer `core.performance.rtt.ix_diff_ms` for details.

 `core.performance.rtt.ix_diff_ms` and `core.performance.rtt.ix_diff_pct` are grouped together at decision making, using a logical AND condition.

Possible values: 1-100

Default value: 10

Recommended value: 5-20

5.8.63 `core.probes.ttl.failed`

Defines the failed probe lifetime (in seconds).

Regular and Retry probing will not be performed until the age of the failed probe is less than `core.probes.ttl.failed`

Possible values: 120-14400

Default value: 300

Recommended value: 300

5.8.64 `core.probes.ttl.max`

Defines the probe lifetime (in seconds).

Probed prefixes data is kept in the IRP database for the specified amount of time. Automatic cleanup of outdated data is performed periodically by Dbcron.

See also: [Administrative Components](#)

High values will lead to database size growth.

Possible values: 86400-604800

Default value: 86400

Recommended value: 86400

5.8.65 `core.probes.ttl.min`

Defines the relevant probe lifetime (in seconds)

Ordinary probing will not be performed for a specific prefix if its probe age is less than `core.probes.ttl.min`.

Possible values: 300-43200

Default value: 7200

Recommended value: 1800-14400

5.8.66 `core.problem.outage_timeout`

Outage detection timeout (in seconds).

The IRP will disregard a possible outage, if it was not confirmed during last `core.problem.outage_timeout`.

Possible values: 0–3600

Default value: 600

Recommended value: 600

See also: `core.outage_detection`

5.8.67 `core.problem.rtt.diff_pct`

Defines the RTT difference (in percent) between the current route and the new route, for Outage detection algorithm. IRP will choose a new route for problematic prefixes only if the RTT difference (in percent) is greater than this parameter.

Possible values: 1–100

Default value: 50

Recommended value: 30–60

5.8.68 `core.vip.interval.probe`

Defines the VIP prefixes probing interval, in seconds. All networks and ASNs specified in `/etc/nocton/policies.conf` and will be queued for priority probing every `core.vip.interval.probe` seconds.

i For probing intervals lower than 5 minutes (300 seconds) IRP uses fast probing algorithms avoiding for example long lasting tracing.

See also: [VIP Improvements](#)

Possible values: 60–14400

Default value: 3600

5.9 Explorer settings

5.9.1 `explorer.aipi`

If enabled, AIPI (Adaptive Inter-Packet Interval) algorithm is executed using inter-packet intervals configured for each provider.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

Recommended value: 0

See also: `peer.X.diag_hop.interval_min`, `peer.X.diag_hop.interval_max`

5.9.2 explorer.algorithms

Enables or disables the use of new scanning, probing and tracing Explorer algorithms.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.9.3 explorer.high_vol_precedence

High volume tasks have priority over the tasks with variable cardinality.

The parameter establishes the balance between the tasks allocated for processing high volume prefixes (high priority for the Customer network) and the network events tasks (such as blackout, congestion, etc).

Possible values: 10–90

Default value: 50

See also: [explorer.maxthreads](#), [collector.export.volume.high.top_n](#)

5.9.4 explorer.infra_ips

Defines the IPv4/IPv6 addressees, used in internal infrastructure to determine the current route for a specific prefix. If netmask is not specified, then /32 is assumed for IPv4 address and /128 for IPv6 address (host addresses).

Format:

1. Space-separated list of infrastructure IPv4/IPv6 addresses.
2. Absolute path to a newline-separated file containing a list of local IPv4/IPv6 prefixes that should be analyzed and optimized by the IRP.

Possible values: See above.

Recommended value: all IPv4/IPv6 addresses used in the internal infrastructure.

5.9.5 explorer.interval.infra

Defines the time interval (in milliseconds) between the packets sent to infrastructure hops (5.9.4). By decreasing this value the Explorer performance enhances.

Possible values: 1–200

Default value: 5

Recommended value: 5

5.9.6 explorer.interval.other

Defines the time interval (in milliseconds) between the packets sent to destination hops (5.9.4).

Possible values: 1–1000

Default value: 200

Recommended value: 200

5.9.7 explorer.interval.other.trace

Defines the time interval (in milliseconds) between traceroute packets sent to non-infrastructure hops.

Possible values: 1–1000

Default value: 20

Recommended value: 20

5.9.8 explorer.ipv4_test

Defines public IPv4 address that will be used for ensuring that PBR is operational.

Possible values: IPv4 address

Default value: 8.8.8.8

Recommended value: 8.8.8.8

5.9.9 explorer.ipv6_test

Defines public IPv6 address that will be used to verify that PBR is operational.

Possible values: IPv6 address

Default value: 2620:0:ccc::2

Recommended value: 2620:0:ccc::2

5.9.10 explorer.log

Defines the file-system path to the explorer log file.

Default value: /var/log/irp/explorer.log

5.9.11 explorer.log.level

Defines the logging level for the explorer service.

Possible values: emerg, fatal, alert, crit, error, warn, notice, info, debug

Default value: info

Recommended value: info

5.9.12 explorer.max_collector_ips

Process max collected IPs. Maximum number of IPs to probe for a specific prefix.


Possible values: 1–255

Default value: 10

Recommended value: 10

5.9.13 explorer.maxthreads

Defines the number of simultaneously processed exploring tasks.

 If this value is excessively large Explorer tasks can take very long to finish. Hardware/router diagnostic packet rate limitations kick in causing all threads to wait for exploring tasks to finish.

Possible values: 10–2000

Default value: 200

Recommended value: 200–300

5.9.14 `explorer.probe.algorithm`

Defines the probing algorithm(s) used by explorer, in the preferred order.

The following algorithms can be used:

- UDP - udp requests (destination port: 33434)
- ICMP - regular ICMP ping sweep
- TCP_SYN - tcp syn scan (destination port: 33434)

Possible values: ICMP UDP TCP_SYN

Default value: ICMP UDP TCP_SYN

Recommended value: ICMP UDP TCP_SYN

5.9.15 `explorer.probe.indirect__priority`

Defines the execution priority between direct and indirect probing algorithms. Parameter value of 1 means that indirect probing will be executed prior to direct probing algorithm. Parameter value of 2 disables indirect probing.

 This parameter conflicts with `explorer.trace.all` configuration option.

Possible values: 0 (Direct then Indirect), 1 (Indirect then Direct), 2 (Direct only)

Default value: 0

Recommended value: 0

5.9.16 `explorer.probing.sendpkts.adaptive__max`

Defines the maximum number of ping packets for adaptive prefix probing.

Possible values: positive Integer value, higher or equal to `explorer.probing.sendpkts.min`

Default value: 100

Recommended value: 100

5.9.17 `explorer.probing.sendpkts.min`

Defines the default ping packets count for each probe. If any loss is detected, additional packets (up to `explorer.probing.sendpkts.adaptive__max`) are sent, for a more accurate packet loss detection.

Possible values: positive Integer value

Default value: 5

Recommended value: 5

5.9.18 explorer.probing.simultaneous

Defines the number of scanned IP addresses.

Possible values: 1-100

Default value: 50

Recommended value: 50

5.9.19 explorer.scanning.confirm_ips

Defines the number of scanned speaking IP addresses for a provider with loss.

Possible values: 1-10

Default value: 5

Recommended value: 5

5.9.20 explorer.scanning.replypkts.min

Defines the number of response packets from a scanned speaking IP address to qualify as a candidate.

Possible values: 1-10

Default value: 5

Recommended value: 5

5.9.21 explorer.scanning.rtt.dispersion_ms

Defines RTT maximum dispersion in milliseconds to qualify a speaking IP as candidate.

Possible values: 1-1000

Default value: 50

Recommended value: 50

5.9.22 explorer.scanning.sendpkts.factor

Defines the number of attempts to send packets to all selected speaking IPs.

Possible values: 1-10

Default value: 5

Recommended value: 5

5.9.23 explorer.timeout

Defines the probes ICMP timeout (in ms)

Default value: 2000

Recommended value: 2000

5.9.24 `explorer.timeout.infra`

Defines the waiting time (in milliseconds) for infrastructure hops' (5.9.4) responses expected during trace execution.

Possible values: 50–20000

Default value: 10000

Recommended value: 10000

5.9.25 `explorer.trace.algorithms`

Defines the traceroute algorithm(s) to be used, arranged by priority. See `explorer.probe.algorithm` for algorithms description.

Possible values: UDP ICMP TCP_SYN

Default value: UDP ICMP

Recommended value: UDP ICMP

5.9.26 `explorer.trace.all`

Defines tracing behaviour. If traces are forced, each probed prefix unconditionally runs traces through all configured providers. Regular tracing is needed for Outage Detection and for AS Path reconstruction. Traces can be disabled altogether by setting `explorer.trace.all` to 2.



Third algorithm should be enabled in `bgpd.as_path` when traces are fully disabled.



Forcing ALL traces (`explorer.trace.all = 1`) can significantly slow down probing for networks. Review probing times before and after changing this parameter and if probing times become unacceptable revert to previous setting.



When all traces are disabled (`explorer.trace.all = 2`) IRP is missing essential trace data and is not able to take action on some types of problems. This effectively cuts off those features.

Possible values: 0 (Regular), 1 (Force all), 2 (Disable all)

Default value: 0

Recommended value: 0

5.9.27 `explorer.trace.diag_hop_unique`

Enforces the uniqueness of per-provider diagnostic hops `peer.X.ipv4.diag_hop`, `peer.X.ipv6.diag_hop`.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.9.28 `explorer.traceroute.retrypkts`

Defines the number of additional traceroute packets to be sent to the intermediate hop in case it has not replied and might be an infrastructure boundary hop.

Possible values: 0–50

Default value: 10

Recommended value: 10

5.9.29 `explorer.traceroute.sendpkts`

Defines the number of traceroute packets per hop to be sent for each probe.

Possible values: 3–5

Default value: 3

Recommended value: 3

5.9.30 `explorer.traceroute.simultaneous`

Defines the number of simultaneously probed hops during trace execution initiated from the hop defined by (5.9.33).

Possible values: 1–30

Default value: 5

Recommended value: 5

5.9.31 `explorer.traceroute.simultaneous.infra`

Defines the number of simultaneously probed infrastructure hops during trace execution.

Possible values: 1–30

Default value: 3

Recommended value: 3

5.9.32 `explorer.traceroute.ttl.max`

Defines the maximum traceroute TTL. Essentially this defines the last hop at which explorer stops the trace.

Possible values: 1–255

Default value: 30

Recommended value: 30

5.9.33 `explorer.traceroute.ttl.min`

Defines the minimum traceroute TTL. Basically this defines the first hop after which explorer analyzes the trace.

Possible values: 1–255

Default value: 2

Recommended value: 2–5

5.10 Inbound Performance

Feature overview: [Inbound performance optimization](#)

How to configure: [Inbound performance optimization configuration](#)

5.10.1 `irpinperfd.enabled`

Enables or disables the inbound performance feature.

Refer to: [1.2.18](#)

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.10.2 `irpinperfd.log`

Defines the path to the log file in the server file-system.

Default value: `/var/log/irp/irpinperfd.log`

5.10.3 `irpinperfd.log.level`

Defines the logging level for the Irpinperfd service.

Possible values: `emerg`, `fatal`, `alert`, `crit`, `error`, `warn`, `notice`, `info`, `debug`

Default value: `info`

5.10.4 `irpinperfd.model.debug`

Enables additional logging of Inbound Performance statistical model.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.10.5 `irpinperfd.model.topn_per_rule`

The number of top volume prefixes to be analyzed per an inbound performance rule.

Increasing this value may raise the accuracy of improvements at the expense of the probing slots.

Possible values: 1 - 1000

Default value: 200

5.10.6 `irpinperfd.moderated`

Enables or disables the moderated mode of handling inbound improvements.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.10.7 `irpinperfd.probing.confirmation_interval`

The time interval (seconds) between two probings of the sample prefixes for active rules.

Possible values: 3600 - 86400

Default value: 18000

5.10.8 `irpinperfd.probing.failure_margin`

The value indicates the minimal acceptable percentage of the failed probes.

Possible values: 0 - 25

Default value: 10

5.10.9 `irpinperfd.probing.interval`

Time interval (seconds) between two probings of the sample prefixes.

Possible values: 300 - 7200

Default value: 600

5.10.10 `irpinperfd.probing.shelf_life`

The value indicates how much time (seconds) the probe will be actual for inbound performance improvement.

Possible values: 30 - 3600

Default value: 300

5.10.11 `irpinperfd.probing.timeout`

Defines the waiting time (seconds) of the probing results. Probing result absence will lead to the algorithm restarts if probing is not finished within the configured time period.

Possible values: 120 - 3600

Default value: 300

5.11 Threat mitigation

5.11.1 `irpdetectd.bgp.reaction`

Threat Mitigation BGP blackholing reaction that gets used by default when no custom rule is defined.

Possible values: 0 - Drop; 1 - Redirect

Default value: 0

5.11.2 `irpdetectd.bgp.redirect.bgp_peers`

The list of BGP communities for BGP redirect default reaction.

The list of BGP router(s) that receive BGP redirect announcements.

 The parameter is also used by the Bgpd component.

Possible values: list of BGP routers

5.11.3 `irpdetectd.bgp.redirect.communities`

 The parameter is also used by the Bgpd component.

Possible values: list of BGP communities

5.11.4 `irpdetectd.blackhole.threshold.kpps`

Blackhole threshold kpps.

Default kilo packets per second limit that triggers a blackholing event.

A value of 0 disables the feature functionality.

The default rate can be overridden by a custom rule.

Possible values: 0 - 1000000

Default value: 0

5.11.5 `irpdetectd.blackhole.threshold.mbps`

Blackhole threshold mbps.

Default megabits per second limit that triggers a blackholing event.

A value of 0 disables the feature functionality.

The default rate can be overridden by a custom rule.

Possible values: 0 - 1000000

Default value: 0

5.11.6 `irpdetectd.flowspec.ipv4.redirect`

The default IPv4 address used by Threat Mitigation FlowSpec redirect reaction.

 The parameter is also used by the Bgpd component.

Possible values: IPv4 address

5.11.7 `irpdetectd.flowspec.ipv6.redirect`

The default IPv6 address used by Threat Mitigation FlowSpec redirect reaction.

 The parameter is also used by the Bgpd component.

Possible values: IPv6 address

5.11.8 `irpdetectd.flowspec.reaction`

Threat Mitigation FlowSpec reaction that gets used by default when no custom rule is defined.

Possible values: 0 - Drop; 1 - Redirect

Default value: 0

5.11.9 `irpdetectd.flowspec.threshold.kpps`

Flowspec threshold kpps.

Default kilo packets per second limit that triggers a Flowspec event.

A value of 0 disables the feature functionality.

The default rate can be overridden by a custom rule.

Possible values: 0 - 1000000

Default value: 0

5.11.10 `irpdetectd.flowspec.threshold.mbps`

Flowspec threshold mbps.

Default megabits per second limit that triggers a Flowspec event.

A value of 0 disables the feature functionality.

The default rate can be overridden by a custom rule.

Possible values: 0 - 1000000

Default value: 0

5.11.11 `irpdetectd.ip_batch_size`

IP batch size of heuristic detection.

Size of IP batch that will be processed by heuristic DDoS detection in one DB request, 0 will disable DDoS detection

Possible values: 0 - 5000

Default value: 0

5.11.12 `irpdetectd.ipv4.prefix_size`

The default size of an IPv4 prefix which gets blocked by the BGP/FlowSpec threat mitigation action.

Possible values: 16 - 32

Default value: 32

5.11.13 `irpdetectd.ipv6.prefix_size`

The default size of an IPv6 prefix which gets blocked by the BGP/FlowSpec threat mitigation action.

Possible values: 32 - 128

Default value: 128

5.11.14 `irpdetectd.mode`

DDoS Mode.

Threat Mitigation modes:

Automatic - threat mitigation actions are performed automatically when an attack gets detected;

Moderated - users need to confirm the threat mitigation action manually;

Disabled - turns off Threat Mitigation altogether.

Possible values: 0 - Disabled; 1 - Manual; 2 - Moderated; 3 - Automated

Default value: 1

5.11.15 `irpdetectd.protected_addresses`

Protected addresses.

Possible values: 0 - Protect analyzed prefixes; 1 - Protect all prefixes

Default value: 0

5.11.16 `irpdetectd.time.keep`

Time keep.

The amount of time (minutes) to keep an approved/automatic flowspec/blackholing event active.

Possible values: 5 - 4320

Default value: 1800

5.11.17 `irpdetectd.time.monitor`

Time monitor.

Amount of time (minutes) between the DDoS attack detection and the automatic activation of the defense mechanism.

Possible values: 1 - 60

Default value: 3

5.11.18 `irpdetectd.whitelist`

Whitelist.

Prefixes that should not be considered for blocking by the DDoS detection mechanism.

Possible values: IPv4 or IPv6 prefixes

5.12 Notification and events

5.12.1 `pushd.email.from`

Defines the From email address used for sending email notifications.

Possible values: valid email address

Default value: root@localhost

5.12.2 `pushd.email.host`

Defines the IPv4, IPv6 address or email server host name used to send emails.

Possible values: Hostname or IPv4/IPv6 address

Default value: 127.0.0.1

5.12.3 `pushd.email.port`

Defines the TCP port used for sending emails.

Possible values: 1 - 65535

Default value: 25

Recommended value: 25

5.12.4 pushd.listen.port

Defines the TCP listen port of irppushd service.

Possible values: 1-65535

Default value: 10499

5.12.5 pushd.log

Defines the complete path to the irppushd log file

Possible values: full path to log file

Default value: /var/log/irp/irppushd.log

5.12.6 pushd.log.level

Defines the logging level for the irppushd service.

Possible values: emerg, fatal, alert, crit, error, warn, notice, info, debug

Default value: info

Recommended value: info

5.12.7 pushd.sms.account_sid

Defines the public account identifier issued to user by the SMS gateway. Usually this is a random string.

Possible values: random string

5.12.8 pushd.sms.auth_token

Defines the secret authentication token issued by the SMS gateway to the account holder. Usually this is a longer random string.

Possible values: random string

5.12.9 pushd.sms.gateway


Defines the SMS gateway preferred by the user. A valid account with this SMS gateway is required to send SMS.

Possible values: twilio, plivo

Default value: none

5.12.10 pushd.sms.message_size

Defines the maximum size of SMS texts. Texts that are longer than this limit will be trimmed and a ... will be added. The SMS gateway will split the SMS text into multiple messages if the request includes a longer than supported SMS message.

 The SMS gateway enforces its own text size limits. In case the notification exceeds SMS gateway limits the SMS message can be either trimmed or rejected.

Possible values: 150-6000

Default value: 150

5.12.11 `pushd.sms.phone_number`

Defines the From phone number used for sending SMS notifications.

❗ SMS gateways might have policies regarding the valid phone numbers that they will accept. Check with your SMS gateways if there are any restrictions and if yes what phone numbers can be provided in this configuration parameter.

Possible values: valid phone number

5.12.12 `pushd.sms.uri.plivo`

Defines the URL of the Plivo SMS API.

❗ Do not change this parameter unless Plivo SMS API URL has changed.

Possible values: valid URL

Default value: `https://api.plivo.com/v1/Account/%1%/Message/`

Recommended value: `https://api.plivo.com/v1/Account/%1%/Message/`

5.12.13 `pushd.sms.uri.twilio`

Defines the URL of the Twilio SMS API.

❗ Do not change this parameter unless Twilio SMS API URL has changed.

Possible values: valid URL

Default value: `https://api.twilio.com/2010-04-01/Accounts/%1%/Messages.json`

Recommended value: `https://api.twilio.com/2010-04-01/Accounts/%1%/Messages.json`

5.12.14 `pushd.templates.datadir`

Defines the directory for notification templates used by irppushd service to format email and sms messages.

Possible values: full path to notification templates

Default value: `/etc/noction/templates`

5.12.15 `pushd.webhook.avatar_emoji`

Defines an emoji to be used as the IRP bot's avatar image. The emoji is expected to be in the "robot-face:" notation.


ℹ The avatar emoji is used if an avatar icon in `pushd.webhook.avatar_urlis` not specified.

Possible values: valid emoji in `:colon:` notation

Default value: `:robot-face:`

5.12.16 pushd.webhook.avatar_url

Defines the URL of the IRP bot's avatar image. URL should be a publicly accessible PNG or JPEG image that the webhook provider is able to retrieve and use.

 This is the default avatar image. The avatar icon will be used instead of an emoji if both are specified. Refer to [pushd.webhook.avatar_emoji](#).

Possible values: valid URL

Default value: `http://www.noction.com/round-logo.png`

Recommended value: `http://www.noction.com/round-logo.png`

5.12.17 pushd.webhook.botname

Defines the name assigned to IRP bot.

Possible values: text


Default value: IRP

Recommended value: IRP

5.12.18 pushd.webhook.url

Defines the URL assigned to Web Hooks user/team. Web Hooks work by POST-ing JSON content to this designated URL for example

POST `https://hooks.slack.com/services/T00000000/B00000000/XXXXXXXXXXXXXXXXXXXXXXXXXXXX`

 This parameter is required in order for any subscription to web hooks channels to work. Currently only the Slack.com web hooks API has been tested and confirmed to work.

Possible values: valid URL

5.12.19 trap.bgpd_announced_rate_low.limit_pct

Defines the announcements rate limit percentage.

Possible values: 1 - 100

Default value: 60

Recommended value: 60

5.12.20 trap.core_cc_improvements_spike.diff_pct

Defines the size (in percent) of the spike in the number of improvement compared to preceeding period average defined in [trap.core_cc_improvements_spike.period_sec](#).

Possible values: 1 - 100

Default value: 50

5.12.21 `trap.core_cc_improvements_spike.period_sec`

Defines the time interval preceeding a spike. The number of improvements is averaged over this time and a spike event is triggered when the subsequent measurement exceeds the size (in percent) defined by `trap.core_cc_improvements_spike.diff_pct`.

Possible values: 30 - 86400

Default value: 300

5.12.22 `trap.core_cc_overload.inbound_limit_mbps`

Defines the overall overload limit in Mbps for the Inbound Commit Control overload by X Mbps event. The event is raised if the overall inbound traffic of the network exceeds the configured value (sum of `peer.X.95th` for all providers) by this limit.

Possible values: 1 - 10000

Default value: 100

Recommended value: 100

5.12.23 `trap.core_cc_overload.inbound_limit_pct`

Defines the overall overload limit in percent for the Inbound Commit Control overload by X% event. The event is raised if the overall inbound traffic of the network exceeds the configured value (sum of `peer.X.95th` for all providers) by this limit.

Possible values: 1 - 1000

Default value: 10

Recommended value: 10

5.12.24 `trap.core_cc_overload.limit_mbps`

Defines the overall overload limit in Mbps for the Outbound Commit Control overload by X Mbps event. The event is raised if the overall outbound traffic of the network exceeds the configured value (sum of `peer.X.95th` for all providers) by this limit.

Possible values: 1 - 10000

Default value: 100

Recommended value: 100

5.12.25 `trap.core_cc_overload.limit_pct`

Defines the overall overload limit in percent for the Outbound Commit Control overload by X% event. The event is raised if the overall outbound traffic of the network exceeds the configured value (sum of `peer.X.95th` for all providers) by this limit.

Possible values: 1 - 1000

Default value: 10

Recommended value: 10

5.12.26 trap.core_cc_provider_overload.inbound_limit_mbps

Defines the per provider overload limit in Mbps for the Inbound Commit Control provider X overloaded by Y Mbps event. The event is raised if the inbound traffic for a provider exceeds the configured value (refer [peer.X.95th](#)) by more than this limit.

Possible values: 1 - 10000

Default value: 100

Recommended value: 100

5.12.27 trap.core_cc_provider_overload.inbound_limit_pct

Defines the per provider overload limit in percent for the Inbound Commit Control provider X overloaded by Y% event. The event is raised if the inbound traffic for a provider exceeds the configured value (refer [peer.X.95th](#)) by more than this limit.

Possible values: 1 - 1000

Default value: 10

Recommended value: 10

5.12.28 trap.core_cc_provider_overload.limit_mbps

Defines the per provider overload limit in Mbps for the Outbound Commit Control provider X overloaded by Y Mbps event. The event is raised if the outbound traffic for a provider exceeds the configured value (refer [peer.X.95th](#)) by more than this limit.

Possible values: 1 - 10000

Default value: 100

Recommended value: 100

5.12.29 trap.core_cc_provider_overload.limit_pct

Defines the per provider overload limit in percent for the Outbound Commit Control provider X overloaded by Y% event. The event is raised if the outbound traffic for a provider exceeds the configured value (refer [peer.X.95th](#)) by more than this limit.

Possible values: 1 - 1000

Default value: 10

Recommended value: 10

5.12.30 trap.core_excessive.prefixes

Defines list of prefixes to restrict events prefixExcessiveLatency & prefixExcessiveLoss.

Events would be generated if the list is empty or aggregate from the list contains a prefix from event.

Possible values: list of valid IPv4/IPv6 prefixes

Default value: 300

Recommended value: 300

5.12.31 `trap.core_excessive_latency.limit`

Defines the packet latency limit in milliseconds.

Possible values: 1 - 10000

Default value: 300

Recommended value: 300

5.12.32 `trap.core_excessive_loss.limit`

Defines the packet loss percentage limit.

Possible values: 1 - 100

Default value: 50

Recommended value: 50

5.12.33 `trap.destination.auth_password`

SNMP user's password for IRP generated traps. Applicable only when SNMP v3 with authentication is used. Refer to `trap.destination.version`, `trap.destination.seclvl`, `trap.destination.auth_username`.

Possible values: text

5.12.34 `trap.destination.auth_protocol`

SNMP authentication protocol for IRP generated traps. Applicable only when SNMP v3 with authentication is used. Refer to `trap.destination.version`, `trap.destination.seclvl`.

- 0 - MD5
- 1 - SHA

Possible values: 0, 1

Default value: 1

5.12.35 `trap.destination.auth_username`

SNMP user name for IRP generated traps. Applicable only when SNMP v3 with authentication is used. Refer to `trap.destination.version`, `trap.destination.seclvl`.

Possible values: text

5.12.36 `trap.destination.community`

Defines the SNMP community used for sending SNMP traps.

Possible values: textual SNMP v2c community name

Default value: public

5.12.37 `trap.destination.port`

Defines the UDP port used for sending SNMP traps.

Possible values: 1 - 65535

Default value: 162

5.12.38 trap.destination.priv_password

SNMP password for IRP generated traps. Applicable only when SNMP v3 with privacy is used. Refer to [trap.destination.version](#), [trap.destination.seclvl](#).

Possible values: text

5.12.39 trap.destination.priv_protocol

SNMP encryption protocol for IRP generated traps. Applicable only when SNMP v3 with privacy is used. Refer to [trap.destination.version](#), [trap.destination.seclvl](#).

 DES isn't supported in RedHat Enterprise Linux 9

- 0 - DES
- 1 - AES

Possible values: 0, 1

Default value: 1

5.12.40 trap.destination.seclvl

SNMP security services for IRP generated traps. IRP supports either No security, Authentication only or both Authentication and Privacy. Applicable only when SNMP v3 is used ([trap.destination.version](#)). Depending on security services used further parameters should be configured, for example: [trap.destination.auth_username](#), [trap.destination.auth_password](#), [trap.destination.auth_protocol](#), [trap.destination.priv_password](#), [trap.destination.priv_protocol](#).

- 0 - Neither Authentication nor Privacy
- 1 - Authentication only
- 2 - Authentication and Privacy

Possible values: 0, 1, 2

Default value: 0

5.12.41 trap.destination.version

SNMP version for IRP generated traps.

- 2 - SNMP v2c
- 3 - SNMP v3

Possible values: 2, 3

Default value: 2

5.13 Administrative settings

5.13.1 dbcron.api.log

Defines file-system path to the dbcron API log file.

Default value: /var/log/irp/api.mlog

5.13.2 dbcron.api.socket

Defines dbcron API listen socket.

Default value: /var/spool/irp/api

5.13.3 dbcron.log

Defines the file-system path to the dbcron log file.

Default value: /var/log/irp/dbcron.log

5.13.4 dbcron.log.level

Defines the logging level for the dbcron service.

Possible values: emerg, fatal, alert, crit, error, warn, notice, info, debug

Default value: info

Recommended value: info

5.13.5 globalcc.key_service.port

Specifies the TCP port used by the Globalcc component for the purpose of public key exchange.

Possible values: 1-65535

Default value: 7600

5.13.6 globalcc.log

Defines the file-system path to the Globalcc log file.

Default value: /var/log/irp/globalcc.log

5.13.7 globalcc.log.level

Defines the logging level for the Globalcc service.

Possible values: emerg, fatal, alert, crit, error, warn, notice, info, debug

Default value: info

Recommended value: info

5.13.8 globalcc.service.port

Specifies the TCP port on which the Irpmng is listening for interprocess communication.

Possible values: 1-65535

Default value: 7601

5.13.9 irpdetected.log

Defines the file-system path to the Iripdetected log file.

Default value: /var/log/irp/irpdetected.log

5.13.10 `irpdetected.log.level`

Defines the logging level for the `Irpdetected` service.

Possible values: `emerg`, `fatal`, `alert`, `crit`, `error`, `warn`, `notice`, `info`, `debug`

Default value: `info`

Recommended value: `info`

5.13.11 `irpmng.grpc.port`

Specifies the TCP port on which the `Irpmng` is listening for interprocess communication.

Possible values: 1-65535

Default value: 7604

5.14 Provider

5.14.1 `peer.X.95th`

 Mandatory if `core.commit_control` is enabled.

Defines current provider's desired 95th value (also known as Committed Interface Rate).

The value of `peer.X.limit_load`(if set to positive value) must be greater or equal to value of `peer.X.95th`. Please refer to `core.commit_control` and `Commit Control` for Commit Control description.

Depending on the `peer.X.95th.bill_day` parameter value, the behavior of the 95th level calculation has two different states.

If `peer.X.95th.bill_day` is set to -1, then this parameter is treated as "committed interface rate limit" and actual 95th calculation is not performed. In this case, IRP will actively reroute traffic from an overloaded provider to keep it below the `peer.X.95th` level.

If `peer.X.95th.bill_day` is set to a positive value, then the system calculates the actual 95th value based on the historical traffic information from `peer.X.95th.bill_day` to the current date of this month. Then, the result is compared to the `peer.X.95th` value, and the maximum of these two is chosen as the target 95th. Thus, if the 95th for the current month has already exceeded the `peer.X.95th` value, IRP will keep the traffic usage from increasing above current 95th.

Possible values: 1-9999999

5.14.2 `peer.X.95th.bill_day`

Defines the first billing period day. If current month has less days than the specified value, then last day of the month will be taken as the start of the new billing period. The parameter is used for 95th percentile calculation.

If -1 is specified, then `peer.X.95th` is treated as "Committed Interface Rate".

Possible values: -1, 1-31

Default value: 1

Recommended value: First day of the billing period.

5.14.3 `peer.X.95th.centile`

Defines the actual percentage for the 95th percentile, as some providers may have non-standard traffic billing policies.

Possible values: 1–99

Default value: 95

Recommended value: please consult your agreements with this specific provider

5.14.4 `peer.X.95th.in`


Defines the 95th level (in Mbps) for inbound when inbound and outbound commit control operate independently. Refer [peer.X.95th](#), [peer.X.95th.mode](#).

Possible values: 1–1000000

5.14.5 `peer.X.95th.mode`

Defines the way how 95th value is determined. The following modes are supported:

- 0: Separate 95th for in/out
- 2: 95th from greater of in or out
- 3: Largest of the two 95th for in/out

 Only mode with separate 95th for each inbound and outbound traffic keeps inbound and outbound commit control independent of each other.

Refer also [peer.X.95th](#), [peer.X.95th.in](#).

Possible values: 0–3

Default value: 2

5.14.6 `peer.X.aspath_for_ix`

Modifies the way IX AS number is treated for outbound improvements towards an Internet Exchange provider for 2nd and 3rd options of [bgpd.as_path](#). The supported options are:

- 0: Strip IX ASN if present
- 1: Preserve IX ASN if present

Refer also [bgpd.as_path](#).

Possible values: 0–1

Default value: 0


5.14.7 `peer.X.auto_config`

Enables or disables periodic execution of Internet Exchanges auto re-configuration process once per [5.2.6](#) interval.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.14.8 `peer.X.bgp_peer`

 **Mandatory**

Defines the iBGP session(s) over which an improvement made for this provider is advertised. This parameter should be set to a valid BGP neighbor name as defined in [BGP sessions settings](#). Typically, this is the session with the Edge router, which this provider is connected to. Multiple sessions should be set only if there are some additional (backup) links to this provider on multiple Edge routers.

Possible values: One or a list of valid BGP neighbor names, as defined in [BGP sessions settings](#).

5.14.9 `peer.X.blackholing.bgp_peer`

Specifies a BGP session(s) to be used to send blackholing announcements to.

Possible values: One or a list of valid BGP neighbor names, as defined in [BGP sessions settings](#).

5.14.10 `peer.X.blackholing.community`

Specifies a BGP community mark to be used in order to distinguish routes that are to be sent towards the provider's router responsible for blackhole routes.

5.14.11 `peer.X.bmp`

Defines BMP data usage for this provider.


- 0: Do not use BMP data
- 1: Use BMP data if available
- 2: Use BMP data only

Possible values: 0 – 2

Default value: 0

5.14.12 `peer.X.bmp.check_routes`

Allows using BMP to check route availability from a Provider configured as a partial routing or IX.

 IRP may unexpectedly use routes received from a Provider, but filtered by an edge router, because BGP filters aren't applied to BMP data.
The parameter must be enabled before switching `peer.X.bmp` to "Use BMP data only" mode.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.14.13 `peer.X.cc_disable`

Instructs IRP to exclude this provider from the Commit Control algorithm.

 Has effect only if `core.commit_control` is enabled.


Possible values: 0, 1

Default value: 0

5.14.14 `peer.X.circuit.control`

Defines whether provider will be monitored for excessive loss issues and how to react to them. The possible values are:

- 0: Disabled
- 1: Warn only
- 2: Warn and disconnect
- 3: Disconnect and restore


 IRP tries to induce a disconnect of BGP session via FlowSpec rules that drop any further packets on the BGP session with provider's router. It is thus recommended that the keepalive timer is set to 20 seconds for BGP sessions with providers that are expected to be disconnected if circuit issues are detected on them. This will minimize the time a circuit with excessive loss is kept operational and causes harm.

Possible values: 0-3

Default value: 0

5.14.15 `peer.X.cost`

Defines the cost per Mbps for the current provider. All `peer.X.cost` parameters should be specified in the same currency.

 Parameter's value should be the same for each provider in a provider's group (refer to 5.14.56) and cost mode is enabled (5.2.25)

Possible values: 1-'unlimited'

Default value: 0

5.14.16 `peer.X.description`

Defines the provider's description (full provider name)

Possible values: text

5.14.17 `peer.X.diag_hop.interval_max`

Defines maximum value for Adaptive Inter-Packet Interval. Adaptive Inter-Packet Interval is used while sending packets to a specific Provider's router (packets with a specific TTL). Inter-packet interval is raised up to the maximum value when the router does not respond to the packets and is decreased to the minimum one in case the router responds. The Adaptive Inter-Packet Interval is changed gradually to obtain better performance.

A value of the 5.9.5 parameter is used in case zero value is specified in this parameter.

Possible values: 0 - 50000000 (50 ms)

Default value: 0

5.14.18 `peer.X.diag_hop.interval_min`

Defines minimum value for Adaptive Inter-Packet Interval. Adaptive Inter-Packet Interval is used while sending packets to a specific Provider's router (packets with a specific TTL). Inter-packet interval is raised up to the maximum value when the router does not respond to the packets and is decreased to the minimum one in case the router responds. The Adaptive Inter-Packet Interval is changed gradually to obtain better performance.

Possible values: 10000 (0.01 ms) – 10000000 (10 ms)

Default value: 1000000

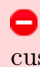
5.14.19 `peer.X.flow_agents`

Specifies a collection of Flow agents in the form IPv4/interfaceIdentifier. Flow agents are used to assign specific Flow statistics to a designated provider. Refer [Flow agents, Optimization for multiple Routing Domains](#) for details.

Possible values: forward slash separated IP address and numeric identifier in 1..2147483647 range


5.14.20 `peer.X.flowspec.ipv4.redirect_community`

Defines the BGP Community that will be appended by Bgpd to advertised Flowspec redirect policies for IPv4 sessions. The format is: “X:Y”.

 Avoid collisions of communities values assigned to IRP both within its configuration and/or on customer's network.

5.14.21 `peer.X.flowspec.ipv6.redirect_community`

Defines the BGP Community that will be appended by Bgpd to advertised Flowspec redirect policies for IPv6 sessions. The format is: “X:Y”.

 Avoid collisions of communities values assigned to IRP both within its configuration and/or on customer's network.

5.14.22 `peer.X.flowspec.pbr.enabled`

Enables or disables Flowspec PBR for a provider.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.14.23 `peer.X.flowspec.pbr.use_bgp_peer`

Specifies whether FlowSpec PBR entries should be advertised only to assigned router(s) or to all routers with FlowSpec PBR enabled.

Possible values: 0 (All routers), 1 (Assigned routers)

Default value: 0

5.14.24 `peer.X.global_group`

Includes the provider into a existing Global Group.

See also: [Global Group](#)

Possible values: 1-100

5.14.25 `peer.X.group_loadbalance`

Defines whether commit control algorithm load balances the group of providers or optimizes it on aggregated bandwidth usage. For details refer [Provider load balancing](#) and [Commit control of aggregated groups](#).

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.14.26 `peer.X.improve_in_group`

Enables or disables improvements within a provider group. This parameter must be equal for all providers in a provider group.

If disabled, IRP will not make any Cost or Performance improvements inside the group.


Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

Recommended value: depends on network infrastructure and policies


5.14.27 `peer.X.inbound.community_base`

Defines the base community of consecutive range of eight community values which IRP uses to announce inbound improvements for this provider over BGP. Base community instructs an announcement of a route without any additional prepend via a particular provider. Next communities in the range instructs for additional prepends from one up to seven.

 For example: when `peer.1.inbound.community_base` is set to 65530:50100, the value 65530:50102 represents “prepend 2 times while advertizing to provider 1”.


Possible values: valid BGP community attribute

5.14.28 `peer.X.ipv4.diag_hop`

 Mandatory

Defines the diagnostic hop or subnet in CIDR format for the current provider. Usually, it is the IP address of the first hop of the specific provider. The parameter is used to make sure that a route actually passes through this provider.


It is also used for Internet Exchanges autoconfiguration process. The Peering Partners next-hops will be gathered based on the provided subnet(s) in CIDR format.

 Any parameter changes (via Frontend) will caused the Bgpd component restart if the provider type (`peer.X.type`) is Internet Exchanges

See also: [explorer.trace.diag_hop_unique](#), [peer.X.ipv6.diag_hop](#)

Possible values: List of valid IPv4 addresses or prefixes (usually equal to `peer.X.ipv4.next_hop`)

5.14.29 `peer.X.ipv4.master_probing`

 Mandatory

Defines the local IPv4 address, assigned for probing via the current provider. When failover is enabled (`global.failover`) slave's probing IP (`peer.X.ipv4.slave_probing`) must be configured too.

See also: [Specific PBR configuration scenarios](#)

Possible values: valid local IPv4 address

5.14.30 `peer.X.ipv4.mon`


Defines the List of IPv4 addresses to be monitored by Bgpd to ensure that the immediate upstream path through this provider is operational. Each specified IP address is probed every `bgpd.mon.keepalive` seconds to verify accessibility from Bgpd (ICMP/UDP pings are used). Highly available IP addresses that reliably answer to ICMP/UDP pings should be used. It is recommended to setup at least two IP addresses belonging to different networks.

If all monitored IP addresses do not respond for `bgpd.mon.holdtime` seconds, then any improvements designated for this provider will be withdrawn from edge router(s). Improvements will be re-announced after all IP addresses respond within `bgpd.mon.guardtime` seconds.

Possible values: space-separated list of valid IPv4 addresses

Default value: 208.67.222.222 8.8.8.8

5.14.31 `peer.X.ipv4.next_hop`

 Mandatory

Defines the next-hop IPv4 address for BGP route injection. Usually, it is the IPv4 address of the BGP partner from the provider.

Possible values: valid IPv4 address

5.14.32 `peer.X.ipv4.next_hop_as`

Defines the provider autonomous system number for IPv4. This parameter must be set in order for the 3rd algorithm of Bgpd AS-PATH to work properly (refer to `bgpd.as_path`).

If this parameter is set, then its value will be used as part of AS-PATH for outgoing improvements if the 3rd algorithm is enabled in `bgpd.as_path`.

In case of Internet Exchanges, the AS-Path begins with peering partner's AS number, instead of AS number of route server.

The first AS number will be stripped from AS-Path when advertising improvement towards Exchange in case the first AS number is equal to value set in this parameter.

See also: `provider.X.rule.Y.next_hop_as`

Possible values: 0-4294967295

Default value: 0

5.14.33 peer.X.ipv4.route_server

Defines the Internet Exchange Route Server(s) IP address(es) used to properly auto-configure `provider.X.rule.Y.bgp_peer` parameter.

See also `peer.X.mon.ipv4.internal.mode`, `peer.X.mon.ipv4.internal.state`

Possible values: valid IPv4 address(es)

5.14.34 peer.X.ipv4.slave_probing

Defines the local IPv4 address, assigned for probing via the current provider for the slave node in a failover configuration.

See also: [Specific PBR configuration scenarios](#)


5.14.35 peer.X.ipv4_pbr_check

Defines per-provider alternative IPv4 address to be used for PBR tests. This overrides `explorer.ipv4_test`.

See also: `peer.X.ipv6_pbr_check`

Possible values: valid IPv4 address

5.14.36 peer.X.ipv6.diag_hop


 Mandatory

Defines the IPv6 diagnostic hop for the current provider. Usually, it is the IP address of the first hop of the specific provider. The parameter is used to verify that a route actually passes through this provider.

See also: `explorer.trace.diag_hop_unique`, `peer.X.ipv4.diag_hop`

Possible values: List of valid IPv6 addresses or prefixes, usually equal to `peer.X.ipv6.next_hop`

5.14.37 peer.X.ipv6.master_probing

 Mandatory for IPv6

Defines the local IPv6 address assigned for probing via the current provider. When failover is enabled (`global.failover`) slave's probing IPv6 address (`peer.X.ipv4.slave_probing`) must be configured too.

See also: [Specific PBR configuration scenarios](#)

Possible values: valid local IPv6 address

5.14.38 peer.X.ipv6.mon


Defines the List of IPv6 addresses to be monitored by Bgpd to ensure that the immediate upstream path through this provider is operational. Each specified IP address is probed every `bgpd.mon.keepalive` seconds to verify accessibility from Bgpd (ICMP/UDP pings are used). Highly available IP addresses that reliably answer to ICMP/UDP pings should be used. It is recommended to setup at least two IP addresses belonging to different networks.

If all IP addresses do not respond for `bgpd.mon.holdtime` seconds, then any improvements designated for this provider will be withdrawn from the edge router(s). Improvements will be announced again after all IP addresses respond within `bgpd.mon.guardtime` seconds.

Possible values: space-separated list of valid IPv6 addresses

Default value: 2620:0:ccc::2 2001:4860:4860::8888

5.14.39 `peer.X.ipv6.next_hop`

 Mandatory for IPv6

Defines the next-hop IPv6 address for BGP route injection. Usually, it is the IPv6 address of the BGP partner from the provider.

Possible values: valid IPv6 address

5.14.40 `peer.X.ipv6.next_hop_as`

Defines the provider autonomous system number for IPv6. This parameter must be set in order for the 3rd algorithm of Bgpd AS-PATH to work properly (refer to [bgpd.as_path](#)).

If this parameter is set, then it's value will be used as part of an AS-PATH for outgoing improvements if the 3rd algorithm is enabled in [bgpd.as_path](#).

See also: [provider.X.rule.Y.next_hop_as](#)

Possible values: 0-4294967295

Default value: 0

5.14.41 `peer.X.ipv6.route_server`

Defines the Internet Exchange Route Server(s) IP address(es) used to properly auto-configure [provider.X.rule.Y.bgp_peer](#) parameter.

See also [peer.X.mon.ipv6.internal.mode](#), [peer.X.mon.ipv6.internal.state](#)

Possible values: valid IPv6 address(es)

5.14.42 `peer.X.ipv6.slave_probing`

Defines the local IPv6 address assigned for probing via the current provider for the slave node in a failover configuration.

See also [peer.X.ipv6.master_probing](#).

Possible values: valid local IPv6 address

5.14.43 `peer.X.ipv6_pbr_check`

Defines per-provider alternative IPv6 address to be used for PBR tests. This overrides [explorer.ipv6_test](#).

See also: [peer.X.ipv4_pbr_check](#).

Possible values: valid IPv6 address

5.14.44 `peer.X.limit_load`

Defines the load limit for current provider. IRP will improve routes to this provider as long as its load is less than the specified value in megabits per second. SNMP must be properly configured in order for this feature to work properly. If this limit is configured, but it is impossible to retrieve interface data, no new improvements will take place.

The value of [peer.X.limit_load](#)(if set to positive value) must be greater or equal to value of [peer.X.95th](#).

Possible values: -1 (unlimited), 1-1000000

Default value: -1

Recommended value: it is recommended to set it to not more than ~60-80% of the physical interface rate

See also: [peer.X.snmp.interfaces](#)

5.14.45 `peer.X.mon.ipv4.bgp_peer`

Defines the current provider's IPv4 address that must be monitored by Bgpd, usually equal to `peer.X.ipv4.next_hop`

Possible values: valid IPv4 address

See also: `peer.X.mon.snmp`

5.14.46 `peer.X.mon.ipv4.external.state`

Enables or disables external monitors for IPv4.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.14.47 `peer.X.mon.ipv4.internal.mode`

Specifies router supported BGP MIB to monitor this provider IPv4 BGP sessions.

- 0 - Generic (BGP4-MIB)
- 1 - Cisco (CISCO-BGP4-MIB)
- 2 - Juniper (BGP4-V2-MIB-JUNIPER)
- 3 - Brocade (draft-ietf-idr-bgp4-mibv2-12)

Possible values: 0, 1, 2, 3

Default value: 0

5.14.48 `peer.X.mon.ipv4.internal.state`

Enables or disables internal monitors for IPv4.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.14.49 `peer.X.mon.ipv6.bgp_peer`

Defines the current provider's IPv6 address that must be monitored by Bgpd, usually equal to `peer.X.ipv6.next_hop`

Possible values: valid IPv6 address

See also: `peer.X.mon.snmp`

5.14.50 `peer.X.mon.ipv6.external.state`

Enables or disables external monitors for IPv6.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.14.51 `peer.X.mon.ipv6.internal.mode`

Specifies router supported BGP MIB to monitor this provider IPv6 BGP sessions.

- 1 - Cisco (CISCO-BGP4-MIB)
- 2 - Juniper (BGP4-V2-MIB-JUNIPER)
- 3 - Brocade (draft-ietf-idr-bgp4-mibv2-12)

Possible values: 1, 2, 3

5.14.52 `peer.X.mon.ipv6.internal.state`

Enables or disables internal monitors for IPv6.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.14.53 `peer.X.mon.snmp`


Identifier of SNMP host to use for monitoring this provider. Refer to [SNMP Host](#), `snmp.X.name`.

Possible values: text

5.14.54 `peer.X.mon.vpn_instance_id`

The Provider routing instance index.

The parameter is useful when a provider belongs to a non-default routing instance, and the Internal Monitor should be addressed to monitor a BGP session inside that specific routing instance.

 Applicable only to Huawei and Arista modes of Internal Monitor.

Default value depends on the mode: Huawei routers have the main routing instance ID equal to zero while Arista routers have this ID equal to 1.

Refer to `peer.X.mon.ipv4.internal.mode`, `peer.X.mon.ipv6.internal.mode`, [Arista BGPv2 MIB](#), [Huawei BGP VPN MIB](#).

Possible values: 0-4294967295

5.14.55 `peer.X.pbr_check`


Enables/disables PBR check for a specific provider.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1


Recommended value: 1

5.14.56 `peer.X.precedence`

 Precedence value of “0” has a special meaning. IRP excludes providers from grouping if a provider’s precedence value is set to 0.

Defines this provider’s precedence, used for provider grouping and defining commit priorities.

See `core.commit_control` and [Commit Control](#) for Commit Control functionality.

 All providers belonging to a single group must have the same value for the following parameters:

- `peer.X.95th.bill_day`
- `peer.X.95th.mode`
- `peer.X.95th.centile`
- `peer.X.cc_disable`
- `peer.X.cost`
- `peer.X.global_group`
- `peer.X.group_loadbalance`
- `peer.X.improve_in_group`

If two different providers have the same `peer.X.precedence` value (except 0), then these providers are considered to be a group, and traffic is balanced between them (refer to `peer.X.group_loadbalance`).

The provider or a group with the lowest precedence value are also used by Commit Control as last resort destination (if all the providers are exceeding their 95th, then traffic is rerouted to the upstream with the smallest precedence - usually this upstream has either the highest throughput, or the lowest cost).

For example:

```
peer.1.precedence = 20
peer.2.precedence = 10
peer.3.precedence = 30
```

If the `peer.X.95th` is exceeded for all configured providers, then the excessive traffic from providers 1 and 3 will be rerouted to the 2nd provider.

If provider groups are used:

```
peer.1.precedence = 50
peer.2.precedence = 40
peer.3.precedence = 40
peer.4.precedence = 40
peer.5.precedence = 30
```

Traffic between providers 2, 3 and 4 is evenly distributed. If all providers are already overloaded, then the excessive traffic will be rerouted to the 5th provider (since it has the lowest precedence value).

Possible values: 0–100

Default value: 0

5.14.57 `peer.X.rd`

Assigns a routing domain identifier to the provider. Providers in the same routing domains should be assigned the same identifier.

Providers with identifier 1 (one) are assumed to be in the same routing domain that hosts IRP instance. Refer `global.rd_rtt`, `peer.X.flow_agents`, `bgpd.rd_local_mark`.

Possible values: 1-100

Default value: 1


5.14.58 `peer.X.routes_config`

Defines whether Internet Exchanges auto configuration is enabled or not. In case it is enabled, IRP Bgpd gathers Peering Partners (next-hops) with their announced routes and stores the data in IRP database. Otherwise, manual configuration is required.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.14.59 `peer.X.shortname`

 Mandatory

Defines the provider's short, abbreviated name (3-20 characters). This parameter's value is used for reports and graphs.

Possible values: text

5.14.60 `peer.X.shutdown`

Defines whether provider is Active (0), Suspended (1) or Shutdown (2).

After the provider suspend action, all the improvements will be stored for short period of time (1h). In case of short maintenance windows, use provider suspend.

After the provider shutdown action, all the improvements will be removed. In case of long maintenance windows, use provider shutdown.

After the provider reactivation (after a previous suspend), all the improvements will be sent to retry probing.

Possible values: 0, 1, 2

Default value: 0

5.14.61 `peer.X.snmp.enhanced`

Enables or disables heuristically enhanced SNMP collection for provider.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.14.62 `peer.X.snmp.interfaces`

Defines a collection of interfaces used for current provider. Individual values are in the form `SNMPHostID:Interface` where

- `SNMPHostID` is the identifier of an SNMP host configured under [5.20](#)
- Interface could be specified either by SNMP index (`ifIndex`) or by name (`ifDescr`).

Possible values: collection of `SNMPHostID:Interface` pairs

5.14.63 `peer.X.type`

Defines the type of a provider. The following types are available:

- 0 - Transit provider
- 1 - Partial routes provider
- 2 - Exchanges provider

Possible values: 0, 1, 2

Default value: 0

5.15 Inbound rule

5.15.1 `inbound.rule.X.bgp_peer`

Defines the router(s) where IRP announces improvements of this inbound prefix.

Possible values: a router identifier(s) separated by space

5.15.2 `inbound.rule.X.communities`

The list of additional BGP communities that will be set for the Inbound Prefix.

5.15.3 `inbound.rule.X.enabled`

Enables or disables Inbound Optimization for this prefix.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.15.4 `inbound.rule.X.full_control`

Specifies individual inbound prefix control level by IRP. Default behavior can be set at system level under [bgpd.full_control](#). The settings specify whether default behavior is inherited from the system level or set individually for this prefix to announce either Improvements only or fully announce the prefix to allowed providers.

Refer to [for details](#).


Possible values: -1 (Use default), 0 (Improvement), 1 (If improved), 2 (All)

Default value: -1

5.15.5 inbound.rule.X.next_hop

Defines a next_hop specific for the inbound rule. When the next_hop is unset, then a value from `bgpd.peer.X.inbound.ipv4.next_hop/bgpd.peer.X.inbound.ipv6.next_hop` is taken.


Next-hop should point to a router that routes/terminates traffic. Otherwise null route should be configured for the Next-Hop.


 next_hop value should be of the same IP address family as the inbound prefix it is assigned to. Refer `inbound.rule.X.prefix`.

Possible values: IP address

5.15.6 inbound.rule.X.prefix

Defines an inbound prefix. Inbound prefixes should belong to your network.

 Inbound Optimization of transit networks will be supported in future releases.

 Each inbound prefix logically extends the list of ournets when it is processed by IRP Collector. Refer `collector.ournets`.


Possible values: prefix in CIDR format

5.15.7 inbound.rule.X.providers

Defines the list of providers where this inbound prefix can be advertised. In case this parameter value is empty then it is treated as prefix will be advertised to ALL providers.

Possible values: providerIDs collection

5.16 Routing Policy

 IRP used a legacy format for policies.conf where a single policy parameters immediately followed the prefix/ASN that defined them. IRP updated policies.conf format and assigned each policy a unique identifier of the form “policy.X.” used to prefix each attribute. Notice the format changes in the samples below.

Legacy routing policy example:

```
asn=65530
vip=0
policy=deny
providers=1 3 5
enabled=1
notes=AS65530 deny via Level3, Cogent and nLayer
```

Normalized routing policy example:

```
policy.1.asn=65530
policy.1.vip=0
policy.1.policy=deny
policy.1.providers=1 3 5
policy.1.enabled=1
policy.1.notes=AS65530 deny via Level3, Cogent and nLayer
```

5.16.1 `policy.X.asn`

Defines the Autonomous System Number which is used to match prefixes from the BGP routing table, originated from the specified ASN.

 Only one matching parameter should be specified per policy

Possible values: 1-4294967295

5.16.2 `policy.X.cascade`

Defines the ASN policies that cascade to downstream AS from designated AS. The parameter applies to ASN policies (and not to prefixes).

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.16.3 `policy.X.community`

A community to mark improvements belonging to a policy.

Possible values: X:Y

5.16.4 `policy.X.country`

Defines the two-character country code which is used to match prefixes from the internal geodata database.

 Only one matching parameter should be specified per policy

Possible values: Two-character ISO country Code

5.16.5 `policy.X.enabled`

Defines whether the policy is enabled or not

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.16.6 `policy.X.forcelocal`

If the parameter is enabled then global improvements aren't performed.

Refer: [policy.X.policy](#)

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.16.7 `policy.X.notes`

Defines a description of the routing policy

Possible values: text

5.16.8 `policy.X.policy`

Defines the type of policy.

allow IRP ensures that traffic flows only via providers listed in the `policy.X.providers` parameter

deny IRP ensures that traffic does not flow via providers listed in `policy.X.providers` parameter

static IRP selects prefixes matched by a corresponding parameter from the BGP routing table and makes static improvements via provider(s) specified in the `policy.X.providers` parameter.

static_exact This option can be enabled when the `policy.X.prefix` parameter gets used.

Possible values: `allow`, `deny`, `static`, `static_exact`

Default value: `allow`

5.16.9 `policy.X.prefix`

Defines the aggregate prefix which is used to match sub-prefixes from the BGP routing table.

If the `policy.X.policy` parameter is set to “static_exact” then the prefix is used “as-is” to make a route via specified provider(s).

If the `policy.X.policy` parameter is set to “static” or “static_exact” and the `policy.X.forcelocal` parameter is enabled, then multiple providers (one per routing domain) can be specified in the `policy.X.providers` parameter.

Refer: `policy.X.forcelocal`, `policy.X.policy`

 Only one matching parameter should be specified per policy

Possible values: Valid IPv4 or IPv6 prefix

5.16.10 `policy.X.priority`

Defines what policy to prioritize in case of overlapping prefixes as might be the case of a large aggregate prefix policy extending over an ASN policy. The same specific prefixes might be covered by both and priority sets which one to actually enforce in favor of the others. Policies with higher priority are prioritized.

Possible values: 0–50

Default value: 0

5.16.11 `policy.X.providers`

Defines the list of providers (IDs) affected by the policy. If the ‘providers’ property is not specified, then all the providers will be included in the policy by default.

Possible values: 1–unlimited

5.16.12 `policy.X.vip`

Defines whether VIP probing is enabled for the specified policy.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.17 Exchange

5.17.1 provider.X.rule.Y.bgp_peer

Defines the BGP Router-ID used for BGP session state monitoring, which is performed by the Internal BGP Monitor ([BGP Monitoring](#)).

Depending on whether a Route Server or Peering sessions are used, the parameter value should be set to Route Server Router-ID or Peering Partner Router-ID accordingly.

Possible values: IPv4 address

5.17.2 provider.X.rule.Y.enabled

Defines whether the rule is enabled or not.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.17.3 provider.X.rule.Y.next_hop

Defines the Peering Partner's IPv4/IPv6 next-hop address.

Possible values: IPv4/IPv6 address

5.17.4 provider.X.rule.Y.next_hop_as

Defines the peering partner's autonomous system number. This parameter must be set in order for the 3rd algorithm of Bgpd AS-PATH to work properly (refer to [bgpd.as_path](#)).

If this parameter is set, then its value will be used as part of AS-PATH for outgoing improvements if the 3rd algorithm is enabled in [bgpd.as_path](#).

In case of Internet Exchanges, the AS-Path begins with peering partner's AS number, instead of AS number of route server.

The first AS number will be stripped from AS-Path when advertising improvement towards Exchange, in case the first AS number is equal to the value set in this parameter.

See also: [peer.X.ipv4.next_hop_as](#), [peer.X.ipv6.next_hop_as](#)

Possible values: 0-4294967295

Default value: 0

5.17.5 provider.X.rule.Y.pbr_check

Enables/disables PBR check for a configured peering partner.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.17.6 provider.X.rule.Y.probing_dscp

Defines the DSCP (Differentiated services code point) used with [provider.X.rule.Y.probing_ip](#) for a specific Peering Partner.

Possible values: 0-63

5.17.7 provider.X.rule.Y.probing_ip

Defines the probing IPv4/IPv6 address used for a specific Peering Partner.

Possible values: IPv4/IPv6 address

5.17.8 provider.X.rule.Y.shortname

Defines a short name for a configured Peering Partner.

Possible values: text

5.17.9 provider.X.rule.Y.transit

The option could be turned on if a Peering Partner provides full Internet access.

IRP will check for route availability, assuming a route is always available.


Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.18 Global Group

5.18.1 globalgroup.X.members

Defines a list of IP addresses or host names of all IRP instances (including Failover slave nodes) to form the Global Group.

 Value of the parameter must be the same on all instances of the Global Group.

5.18.2 globalgroup.X.outbound.95th

Amount of bandwidth provided by this IRP instance to the Global Group.

If the value -1 is set, then IRP automatically uses total amount of outbound bandwidth set in the [peer.X.95th](#) parameter for all the providers in the Global Group.

Possible values: -1 (follow Providers), 1 - 1000000

5.18.3 globalgroup.X.outbound.95th.reserve


Amount of bandwidth reserved to operation of this IRP instance.

Possible values: 0 (not used) - 1000000

Default value: 0

5.18.4 globalgroup.X.outbound.rate_high

Defines Global Group's high load rate limit (%).

 Value of the parameter must be the same on all instances of the Global Group.


See also: [core.commit_control.rate.high](#)

Possible values: 30-99, but greater than or equal to [globalgroup.X.outbound.rate_low](#)

Default value: 95

5.18.5 `globalgroup.X.outbound.rate_low`

Defines Global Group's low load rate limit (%).

 Value of the parameter must be the same on all instances of the Global Group.


See also: [core.commit_control.rate.low](#)

Possible values: 30–99, but lower than or equal to [globalgroup.X.outbound.rate_high](#)

Default value: 85

5.18.6 `globalgroup.X.shortname`

Defines short name of the Global Group.

 Value of the parameter must be the same on all instances of the Global Group.

5.19 User Directory

5.19.1 `directory.X.admin_role_groups`

Defines the specific objects/groups that list IRP users with Admin privileges. Also refer [directory.X.user_role_attr](#).

Possible values: text

5.19.2 `directory.X.base_dn`

Defines the base DN (distinguished name) of users in the directory. Base DN is used in conjunction with [directory.X.username](#) and [directory.X.user_dn](#) to uniquely identify and query the directory during authentication.

Possible values: text

5.19.3 `directory.X.email`

Defines the directory attribute that stores user's email address.

Possible values: text

5.19.4 `directory.X.fullname`

Defines the directory attribute that stores user's full name.

Possible values: text

5.19.5 `directory.X.hostname`

Defines the hostname or the IP address of the User Directory.

Possible values: domain name or IP address

5.19.6 `directory.X.initial_bind_dn`

Defines User DN assigned to IRP to bind to directory. This user's password is configured in `directory.X.initial_bind_password`.

Possible values: `text`

5.19.7 `directory.X.initial_bind_password`

Defines the password assigned to IRP to bind to directory. Refer `directory.X.initial_bind_dn`.

Possible values: `text`


5.19.8 `directory.X.name`

Assigns a name to user directory within IRP.

Possible values: `text`

5.19.9 `directory.X.order`

Defines the sequence this user directory is queried when multiple user directories are configured. Smaller values take precedence.

 The value must be unique across configured user directories.

Possible values: 0-255

Default value: 255

5.19.10 `directory.X.port`

Defines the port used by the directory, for example 389 for LDAP.

Possible values: 1-65535

5.19.11 `directory.X.secret`

Defines the shared secret for TACACS+ host and IRP to secure communications.

Possible values: `text`

5.19.12 `directory.X.state`

Enables or disables use of this directory within IRP.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 1

5.19.13 `directory.X.timeout`

Defines the timeout (in seconds) after which IRP stops trying to communicate with the directory.

Possible values: 0-300

Default value: 60

5.19.14 `directory.X.tls`

Enables or disables use of TLS for this directory.

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.19.15 `directory.X.tls_cacertfile`

Defines the path to a CA certificate file used when server certificate validation is enabled and CA certificate cannot be included in trusted server certificates store, for example a single purpose self-signed CA certificate is used. Refer [directory.X.verify_cert](#).

Possible values: text (path)

5.19.16 `directory.X.type`

Defines user directory type for example LDAP, Active Directory, Internal etc.

Possible values: internal, ldap_posix, active_directory, tacacs+

Default value: internal

5.19.17 `directory.X.user_dn`

Defines the DN of users. User DN is used in conjunction with [directory.X.username](#) and [directory.X.base_dn](#) to uniquely identify and query the directory during authentication. Only users matching [directory.X.user_filter](#) are granted access to the system.

Possible values: text

5.19.18 `directory.X.user_filter`

Defines a filter used to match users that will be granted access. The filter uses standard LDAP filter syntax with the exception of outermost parenthesis that are expected to be dropped from configuration.

Possible values: text

5.19.19 `directory.X.user_role_attr`

Defines the attribute that stores the list of users in a role/group. Refer [directory.X.admin_role_groups](#).

Possible values: text

5.19.20 `directory.X.username`

Defines the directory attribute that represents the username of users in the directory. This attribute together with [directory.X.user_dn](#) and [directory.X.base_dn](#) are used to uniquely identify and query the directory during authentication.

Possible values: text

5.19.21 `directory.X.verify_cert`

Enables or disables verification of directory server certificate when TLS is enabled. When verification is enabled a path to CA certificate file must be provided too. Refer [directory.X.tls](#), [directory.X.tls_cacertfile](#).

Possible values: 0 (Disabled), 1 (Enabled)

Default value: 0

5.20 SNMP Host

5.20.1 snmp.X.auth_password

User's password used to authenticate SNMP communications. Applicable only when SNMP v3 with authentication is used. Refer to [snmp.X.version](#), [snmp.X.seclevel](#), [snmp.X.auth_username](#).

Possible values: text

5.20.2 snmp.X.auth_protocol

SNMP authentication protocol. Applicable only when SNMP v3 with authentication is used. Refer to [snmp.X.version](#), [snmp.X.seclevel](#).

- 0 - MD5
- 1 - SHA


Possible values: 0, 1

5.20.3 snmp.X.auth_username

SNMP user name. Applicable only when SNMP v3 with authentication is used. Refer to [snmp.X.version](#), [snmp.X.seclevel](#).

Possible values: text

5.20.4 snmp.X.community

 Mandatory parameter

Defines the SNMP community.

Possible values: textual community name


5.20.5 snmp.X.context

Defines the SNMPv3 context name.

Possible values: textual context name

5.20.6 snmp.X.ip

Defines IPv4/IPv6 address of the SNMP host.

 Hostnames are not supported.

Possible values: valid IPv4 or IPv6 address

5.20.7 snmp.X.name

Sets a shortname for SNMP host.

Possible values: text

5.20.8 snmp.X.port

Defines the UDP port to send SNMP requests.

Possible values: 1–65535

Default value: 161

5.20.9 snmp.X.priv_password

Password used to encrypt SNMP communications. Applicable only when SNMP v3 with privacy is used. Refer to [snmp.X.version](#), [snmp.X.secllevel](#).

Possible values: text

5.20.10 snmp.X.priv_protocol

SNMP encryption protocol. Applicable only when SNMP v3 with privacy is used. Refer to [snmp.X.version](#), [snmp.X.secllevel](#).

 DES isn't supported in RedHat Enterprise Linux 9

- 0 - DES
- 1 - AES

Possible values: 0, 1

5.20.11 snmp.X.secllevel

SNMP security services. IRP supports either No security, Authentication only or both Authentication and Privacy. Applicable only when SNMP v3 is used ([snmp.X.version](#)). Depending on security services used further parameters should be configured, for example: [snmp.X.auth_username](#), [snmp.X.auth_password](#), [snmp.X.auth_protocol](#), [snmp.X.priv_password](#), [snmp.X.priv_protocol](#).

- 0 - None - neither Authentication nor Privacy is used
- 1 - Authentication only
- 2 - Authentication and Privacy

Possible values: 0, 1, 2

Default value: 0

5.20.12 snmp.X.timeout

Timeout period before retrying an SNMP request (s).

Possible values: 1 – 5

Default value: 1

5.20.13 snmp.X.version

SNMP version used by host.

- 2 - SNMP v2c
- 3 - SNMP v3

Possible values: 2, 3


Default value: 2

5.21 Routing domain

5.21.1 rd.X.community.local

Specifies a BGP community which will be appended to BGP community attribute of an improvement in the corresponding routing domain.

Parameter represents a valid value for BGP community attribute of the form X:Y. Value in rd.X.community.local is APPENDED to communities attribute.

 The value must be unique across all configured BGP communities.

Refer [global.rd_rtt](#), [peer.X.rd](#), [peer.X.flow_agents](#), [bgpd.rd_local_mark](#).

Possible values: X:Y

5.21.2 rd.X.community_worsening

RD performance worsening community.

The routing domain designated community value is added to the announced global outbound improvement if such improvement resulted in a degraded performance per the allowed worsening thresholds for a specific domain.

Possible values: valid BGP community

5.21.3 rd.X.shortname

A shortname assigned to each routing domain for human readability.

Possible values: text

Chapter 6

Appendixes

6.1 Frontend labels for configuration parameters

AS Path rules for IX ASN `peer.X.aspath_for_ix`
AS-PATH restore priority `bgpd.as_path`
Account ID `pushd.sms.account_sid`
Adaptive packets count `explorer.probing.sendpkts.adaptive_max`
Additional communities `inbound.rule.X.communities`
Allowed IP addresses `global.frontend_acl_ips`
Allowed latency worsening (ms) `core.cost.worst_ms`
Allowed loss worsening (%) `core.commit_control.worst_loss`
Alternative PBR test IPv4 address `peer.X.ipv4_pbr_check`
Analyzed prefixes `collector.ournets`
Announced blackholing localpref value `bgpd.peer.X.blackholing.localpref`
Announced inbound localpref value `bgpd.peer.X.inbound.master_localpref`, `bgpd.peer.X.inbound.slave_localpref`
Autonomous System `bgpd.peer.X.as`
Autonomous system number `policy.X.asn`
Avatar emoji `pushd.webhook.avatar_emoji`
Avatar icon URL `pushd.webhook.avatar_url`
BGP MIB (IPv4) `peer.X.mon.ipv4.internal.mode`
BGP MIB (IPv6) `peer.X.mon.ipv6.internal.mode`
BGP Router ID `bgpd.peer.X.master_router_id`, `bgpd.peer.X.slave_router_id`
BGP community `policy.X.community`
BGP mode `global.nonintrusive_bgp`
BGP redirect communities `irpdetectd.bgp.redirect.communities`
BGP redirect routers `irpdetectd.bgp.redirect.bgp_peers`
BGP session monitoring IPv4 address `peer.X.mon.ipv4.bgp_peer`
BGP session monitoring IPv6 address `peer.X.mon.ipv6.bgp_peer`
BGP session password `bgpd.peer.X.master_password`, `bgpd.peer.X.slave_password`
BGPd monitoring guard time (sec) `bgpd.mon.guardtime`
BGPd monitoring holdtime (sec) `bgpd.mon.holdtime`
BGPd monitoring keepalive (sec) `bgpd.mon.keepalive`
BGPd monitoring long holdtime (sec) `bgpd.mon.longholdtime`
BW reserved for local instance operation `globalgroup.X.outbound.95th.reserve`
Base community for inbound improvements `peer.X.inbound.community_base`
Blackhole threshold kpps `irpdetectd.blackhole.threshold.kpps`

Blackhole threshold mbps `irpdetected.blackhole.threshold.mbps`
 Blackholing IPv4 next hop `bgpd.peer.X.blackholing.ipv4.next_hop`
 Blackholing IPv6 next_hop `bgpd.peer.X.blackholing.ipv6.next_hop`
 Blackholing Routers `peer.X.blackholing.bgp_peer`
 Blackholing community `peer.X.blackholing.community`
 Borrow AS-PATH if can't restore `bgpd.as_path_borrowing`
 Bot name `pushd.webhook.botname`
 CC allowed for `core.commit_control.loss_override`
 CC probing TTL (sec) `core.commit_control.probe_ttl`
 CC probing queue slots `core.commit_control.probing_queue_size`
 CC provider precedence `peer.X.precedence`
 Cascade policy `policy.X.cascade`
 Cascading policy max AS `bgpd.policy.cascade.amount`
 Centile value `peer.X.95th.centile`
 Circuit issues detection `peer.X.circuit.control`
 Commit Control `core.commit_control`
 Commit Control for inbound `core.commit_control.inbound.enabled`
 Commit Control for provider `peer.X.cc_disable`
 Community mark for RD `rd.X.community.local`
 Community marker for local improvements `bgpd.rd_local_mark`
 Confirmation reprobing interval `irpinperfd.probing.confirmation_interval`
 Control `inbound.rule.X.full_control`
 Country `policy.X.country`
 DDoS Mode `irpdetected.mode`
 Default BGP reaction `irpdetected.bgp.reaction`
 Default FlowSpec reaction `irpdetected.flowspec.reaction`
 Delete irrelevant CC improvements `core.commit_control.del_irrelevant`
 Delta loss to restore `core.circuit.recover_loss_diff`
 Delta loss to shutdown `core.circuit.high_loss_diff`
 Delta loss to warn `core.circuit.warn_loss_diff`
 Diag hop uniqueness enforcement `explorer.trace.diag_hop_unique`
 Directory Base DN `directory.X.base_dn`
 Directory CA certificate `directory.X.tls_cacertfile`
 Directory TLS use `directory.X.tls`
 Directory User DN `directory.X.user_dn`
 Directory admin role DNs `directory.X.admin_role_groups`
 Directory bind DN `directory.X.initial_bind_dn`
 Directory bind password `directory.X.initial_bind_password`
 Directory certificate verification `directory.X.verify_cert`
 Directory email attribute `directory.X.email`
 Directory full name attribute `directory.X.fullname`
 Directory hostname `directory.X.hostname`
 Directory name `directory.X.name`
 Directory order `directory.X.order`
 Directory port `directory.X.port`
 Directory state `directory.X.state`
 Directory timeout `directory.X.timeout`

Directory type `directory.X.type`
 Directory user role attribute `directory.X.user_role_attr`
 Directory username attribute `directory.X.username`
 Email server `pushd.email.host`
 Enable policy `policy.X.enabled`
 Explorer algorithms `explorer.algorithms`
 Explorer worker threads `explorer.maxthreads`
 Exploring IPv4 queue slots `core.eventqueueulimit`
 Exploring IPv6 queue slots `core.eventqueueulimit_ipv6`
 Failback timer (s) `global.failover_timer_failback`
 Failed probe lifetime (sec) `core.probes.ttl.failed`
 Failover `global.failover`
 Failover master identity file `global.failover_identity_file`
 Failover timer (s) `global.failover_timer_fail`
 Flapping protection `bgpd.mon.internal.flap_guardtime`
 Flow Collector `collector.flow.enabled`
 Flow agents `peer.X.flow_agents`
 Flow sources `collector.flow.sources`
 FlowSpec redirect IPv4 `irpdetectd.flowspec.ipv4.redirect`
 FlowSpec redirect IPv6 `irpdetectd.flowspec.ipv6.redirect`
 FlowSpec redirect type `bgpd.peer.X.flowspec.redirect_type`
 Flowspec `global.flowspec`
 Flowspec PBR `global.flowspec.pbr`, `peer.X.flowspec.pbr.enabled`
 Flowspec status `bgpd.peer.X.flowspec`
 Flowspec threshold kpps `irpdetectd.flowspec.threshold.kpps`
 Flowspec threshold mbps `irpdetectd.flowspec.threshold.mbps`
 From email `pushd.email.from`
 From phone number `pushd.sms.phone_number`
 Frontend access restriction `global.frontend_acl`
 Global Group member list `globalgroup.X.members`
 Global Group name `globalgroup.X.shortname`
 Global commit improvements `core.global.allow_commit`
 Global improvements max worsening (ms) `core.global.worst_ms`
 Global latency/cost improvements `core.global.allow_latency_cost`
 Group balance check monitors `core.commit_control.group_balance.check_monitor`
 High load limit (%) `globalgroup.X.outbound.rate_high`
 High volume precedence `explorer.high_vol_precedence`
 Holdtime of BGP updates (sec) `bgpd.db.timeout.withdraw`
 ICMP timeout (ms) `explorer.timeout`
 ICMP/UDP ping monitored IPv4 addresses `peer.X.ipv4.mon`
 ICMP/UDP ping monitored IPv6 addresses `peer.X.ipv6.mon`
 IP address for Internal Monitor `provider.X.rule.Y.bgp_peer`
 IP prefix `policy.X.prefix`
 IPv4 External monitor `peer.X.mon.ipv4.external.state`
 IPv4 Internal monitor `peer.X.mon.ipv4.internal.state`
 IPv4 Redirect community `peer.X.flowspec.ipv4.redirect_community`
 IPv4 aggregate size `global.agg_ipv4_max`

IPv4 diagnostic hop `peer.X.ipv4.diag_hop`
 IPv6 External monitor `peer.X.mon.ipv6.external.state`
 IPv6 Internal monitor `peer.X.mon.ipv6.internal.state`
 IPv6 Redirect community `peer.X.flowspec.ipv6.redirect_community`
 IPv6 aggregate size `global.agg_ipv6_max`
 IPv6 diagnostic hop `peer.X.ipv6.diag_hop`
 IPv6 support `global.ipv6_enabled`
 IRP's IPv4 address `bgpd.peer.X.master_our_ip`, `bgpd.peer.X.slave_our_ip`
 IRP's IPv6 address `bgpd.peer.X.master_our_ipv6`, `bgpd.peer.X.slave_our_ipv6`
 IX auto re-configuration `peer.X.auto_config`
 IX auto re-configuration interval (s) `global.exchanges.auto_config_interval`
 Ignore unannounced prefixes `global.ignored.unannounced`
 Ignored ASNs `global.ignored.asn`
 Ignored communities `global.ignored_communities`
 Ignored prefixes `global.ignorednets`
 Improvement MED value `bgpd.peer.X.med`
 Improvement communities `bgpd.peer.X.master_communities`, `bgpd.peer.X.slave_communities`
 Improvement localpref `bgpd.peer.X.master_localpref`, `bgpd.peer.X.slave_localpref`
 Improvement mode `global.improve_mode`
 Inbound overload by (%) `trap.core_cc_overload.inbound_limit_pct`
 Inbound CC moderated mode `core.commit_control.inbound.moderated`
 Inbound Injection method `global.inbound.injection`
 Inbound Performance moderated mode `irpinperfd.moderated`
 Inbound Prefix Next Hop `inbound.rule.X.next_hop`
 Inbound bandwidth estimation algorithm `core.commit_control.inbound.volume_estimation`
 Inbound high load limit (%) `core.commit_control.inbound.rate.high`
 Inbound improvement delay (sec) `core.commit_control.inbound.improvement.delay`
 Inbound low load limit(%) `core.commit_control.inbound.rate.low`
 Inbound overload by (Mbps) `trap.core_cc_overload.inbound_limit_mbps`
 Inbound performance `irpinperfd.enabled`
 Inbound prefix control `bgpd.full_control`
 Inbound traffic distribution history `collector.export.archive_inbound`
 Indirect probing precedence `explorer.probe.indirect_priority`
 Infrastructure IPs `explorer.infra_ips`
 Instance traffic stake within GG (Mbps) `globalgroup.X.outbound.95th`
 Internal monitor SNMP host `peer.X.mon.snmp`
 Irppushd TCP listen port `pushd.listen.port`
 Irpspand interfaces `collector.span.interfaces`
 Issues time horizon (min) `core.circuit.hist_interval`
 Keepalive interval (sec) `bgpd.peer.X.keepalive`
 List of providers `policy.X.providers`
 Load balancing within group `peer.X.group_loadbalance`
 Local IPv4 inbound next_hop `bgpd.peer.X.inbound.ipv4.next_hop`
 Local IPv6 inbound next_hop `bgpd.peer.X.inbound.ipv6.next_hop`
 Local improvement only `policy.X.forcelocal`
 Low load limit (%) `globalgroup.X.outbound.rate_low`
 Management interface `global.master_management_interface`

Master routing domain `global.master_rd`
 Max IPv4 Flowspec rules `core.flowspec.max`
 Max IPv4 Improvements `core.improvements.max`
 Max IPv6 Flowspec rules `core.flowspec.max_ipv6`
 Max IPv6 Improvements `core.improvements.max_ipv6`
 Max collected IPs `explorer.max_collector_ips`
 Max message size `pushd.sms.message_size`
 Maximum load per interface (Mbps) `peer.X.limit_load`
 Maximum probe lifetime (sec) `core.probes.ttl.max`
 Min loss detection packets `explorer.probing.sendpkts.min`
 Mindelay probing queue slots `collector.span.min_delay.probing_queue_size`
 Minimal prefix bandwidth (Mbps) `core.commit_control.agg_bw_min`
 Minimal probe lifetime (sec) `core.probes.ttl.min`
 Minimal traffic volume (%) `collector.export.volume.min_pct`
 Minimal traffic volume (bytes) `collector.export.volume.min`
 Mitigation prefix size (IPv4) `irpdetectd.ipv4.prefix_size`
 Mitigation prefix size (IPv6) `irpdetectd.ipv6.prefix_size`
 NO_EXPORT/NO_ADVERTISE community `bgpd.no_export`
 NetFlow UDP port `collector.flow.listen.nf`
 Outage confirmation timeout (sec) `core.problem.outage_timeout`
 Outage detection `core.outage_detection`
 Outage prefix confirmation rate `core.outage_detection.limit_pct`
 Outage round trip rate (%) `core.problem.rtt.diff_pct`
 Outbound high load limit(%) `core.commit_control.rate.high`
 Outbound low load limit(%) `core.commit_control.rate.low`
 Outbound optimization `global.outbound`
 Outbound traffic matching `collector.flow.all_outbound`
 Overload by (%) `trap.core_cc_overload.limit_pct`
 Overload by (Mbps) `trap.core_cc_overload.limit_mbps`
 Overusage interval (sec) `core.overusage.check_interval`
 Overusage policies `global.bw_overusage`
 Overusage rule retention (sec) `core.overusage.hold_timer`
 Overusage threshold multiplier `core.overusage.out.threshold.trigger`
 Overusage throttle multiplier `core.overusage.out.threshold.throttle`
 PBR check for provider `peer.X.pbr_check`
 Packet size from IP header `collector.span.size_from_ip_header`
 Peer AS number override `provider.X.rule.Y.next_hop_as`
 Peering Partner PBR check `provider.X.rule.Y.pbr_check`
 Peering Partner name `provider.X.rule.Y.shortname`
 Peering Partner next-hop `provider.X.rule.Y.next_hop`
 Peering Partner status `provider.X.rule.Y.enabled`
 Perf/Cost improvements within group `peer.X.improve_in_group`
 Plivo API uri `pushd.sms.uri.plivo`
 Policy notes `policy.X.notes`
 Policy priority `policy.X.priority`
 Policy type `policy.X.policy`
 Prefix `inbound.rule.X.prefix`

Prefix BW average time (hours) `core.overusage.out.average.period`

Prefix aggregation `global.aggregate`

Prefix announcement rate (%) `trap.bgpd_announced_rate_low.limit_pct`

Prefix filter for excessive loss/latency `trap.core_excessive.prefixes`

Prefix latency limit (ms) `trap.core_excessive_latency.limit`

Prefix loss limit (%) `trap.core_excessive_loss.limit`

Prefix relevant BW (Mbps) `core.overusage.out.average.relevant_min`

Prefix status `inbound.rule.X.enabled`

Prepend inbound prefixes `core.circuit.inbound`

Prepend transit prefixes `core.circuit.transit`

Probe shelf life (s) `irpinperfd.probing.shelf_life`

Probing DSCP `provider.X.rule.Y.probing_dscp`

Probing IP address `provider.X.rule.Y.probing_ip`

Probing IPv4 address `peer.X.ipv4.master_probing`, `peer.X.ipv4.slave_probing`

Probing IPv6 address `peer.X.ipv6.master_probing`

Probing algorithm `explorer.probe.algorithm`

Probing failure margin `irpinperfd.probing.failure_margin`

Probing interface(s) `global.master_probing_interface`

Probing timeout `irpinperfd.probing.timeout`

Protected addresses `irpdetctd.protected_addresses`

Provider 95th Calculation Mode `peer.X.95th.mode`

Provider 95th percentile `peer.X.95th`

Provider ASN for IPv4 `peer.X.ipv4.next_hop_as`

Provider ASN for IPv6 `peer.X.ipv6.next_hop_as`

Provider Shutdown `peer.X.shutdown`

Provider cost per Mbps (USD) `peer.X.cost`

Provider description `peer.X.description`

Provider inbound overload by (%) `trap.core_cc_provider_overload.inbound_limit_pct`

Provider overload by (%) `trap.core_cc_provider_overload.limit_pct`

Provider overload by (Mbps) `trap.core_cc_provider_overload.limit_mbps`

Provider overload inbound by (Mbps) `trap.core_cc_provider_overload.inbound_limit_mbps`

Provider probing IPv6 address `peer.X.ipv6.slave_probing`

Provider type `peer.X.type`

Provider's 95th (inbound) `peer.X.95th.in`

Provider's bandwidth max deviation (%) `core.commit_control.rate.group`

Provider's billing day `peer.X.95th.bill_day`

Provider's global group ID `peer.X.global_group`

Provider's routing domain `peer.X.rd`

Provider/link name `peer.X.shortname`

Providers `inbound.rule.X.providers`

Public IPv4 address for PBR tests `explorer.ipv4_test`

Public IPv6 address for PBR tests `explorer.ipv6_test`, `peer.X.ipv6_pbr_check`

RD performance worsening community `rd.X.community_worsening`

RD shortname `rd.X.shortname`

RTT between RDs `global.rd_rtt`

Re-probe on new path change `bgpd.retry_probing.new.bmp_path_change`

Re-probe on old path change `bgpd.retry_probing.old.bmp_path_change`

React on Irpflowd stats `core.commit_control.react_on_collector`
 Relevant RTT difference (%) `core.performance.rtt.diff_pct`
 Relevant RTT difference (ms) `core.performance.rtt.diff_ms`
 Relevant RTT: IX vs transit (%) `core.performance.rtt.ix_diff_pct`
 Relevant RTT: IX vs transit (ms) `core.performance.rtt.ix_diff_ms`
 Relevant loss (%) `core.performance.loss_pct`
 Remove if exact prefix absent in BMP `bgpd.improvements.remove.bmp_exact_aggregate`
 Remove next hop update `bgpd.improvements.remove.next_hop_eq`
 Remove on aggregate withdrawal `bgpd.improvements.remove.withdrawn`
 Reprobing interval `irpinperfd.probing.interval`
 Restore after (sec) `core.circuit.recover_hold_time`
 Restore interval (min) `core.circuit.recover_monitored_intervals`
 Role `global.failover_role`
 Route server's IPv4 addresses `peer.X.ipv4.route_server`
 Route server's IPv6 addresses `peer.X.ipv6.route_server`
 Router `inbound.rule.X.bgp_peer`, `peer.X.bgp_peer`
 Router IPv4 address `bgpd.peer.X.master_peer_ip`, `bgpd.peer.X.slave_peer_ip`
 Router IPv6 address `bgpd.peer.X.master_peer_ipv6`, `bgpd.peer.X.slave_peer_ipv6`
 Router next-hop IPv4 address `peer.X.ipv4.next_hop`
 Router next-hop IPv6 address `peer.X.ipv6.next_hop`
 Routes configuration mode `peer.X.routes_config`
 SMS gateway `pushd.sms.gateway`
 SNMP Traps community `trap.destination.community`
 SNMP Traps destination port `trap.destination.port`
 SNMP Username (traps) `trap.destination.auth_username`
 SNMP authentication (traps) `trap.destination.auth_protocol`
 SNMP authentication password (traps) `trap.destination.auth_password`
 SNMP community `snmp.X.community`
 SNMP encryption `snmp.X.priv_protocol`
 SNMP encryption (traps) `trap.destination.priv_protocol`
 SNMP encryption password `snmp.X.priv_password`
 SNMP encryption password (traps) `trap.destination.priv_password`
 SNMP host address `snmp.X.ip`
 SNMP host short name `snmp.X.name`
 SNMP interfaces `peer.X.snmp.interfaces`
 SNMP port `snmp.X.port`
 SNMP security (traps) `trap.destination.seclvl`
 SNMP v3 Username `snmp.X.auth_username`
 SNMP v3 authentication password `snmp.X.auth_password`
 SNMP v3 authentication protocol `snmp.X.auth_protocol`
 SNMP v3 security level `snmp.X.seclvl`
 SNMP version `snmp.X.version`
 SNMP version (traps) `trap.destination.version`
 SNMPv3 context `snmp.X.context`
 SPAN Collector `collector.span.enabled`
 Safe removal of improvements `core.improvements.safe_removal`
 Scanning attempts `explorer.scanning.sendpkts.factor`

Secret `pushd.sms.auth_token`
 Server port `pushd.email.port`
 Shutdown `bgpd.peer.X.shutdown`
 Slave IPv4 address `global.failover_slave.ip`
 Slave IPv6 address `global.failover_slave.ipv6`
 Slave Management interface `global.slave_management_interface`
 Slave Probing interface(s) `global.slave_probing_interface`
 Slave SSH port `global.failover_slave.port`
 Slave routing domain `global.slave_rd`
 Speaking IP responses for candidates `explorer.scanning.replypkts.min`
 Speaking IPs RTT dispersion (ms) `explorer.scanning.rtt.dispersion_ms`
 Speaking IPs to scan on loss `explorer.scanning.confirm_ips`
 Spike preceeding interval (seconds) `trap.core_cc_improvements_spike.period_sec`
 Spike size (%) `trap.core_cc_improvements_spike.diff_pct`
 Split outbound announcements `bgpd.updates.split`
 Standard reprobing period (sec) `core.improvements.ttl.retry_probe`
 Strip non-IRP communities `bgpd.improvements.strip_non_irp_communities`
 TACACS+ directory secret `directory.X.secret`
 The routing instance index `peer.X.mon.vpn_instance_id`
 Time keep `irpdetectd.time.keep`
 Time monitor `irpdetectd.time.monitor`
 Timeout interval (s) `snmp.X.timeout`
 Top volume prefixes per cycle `collector.export.volume.high.top_n`
 Top volume prefixes per inperf rule `irpinperfd.model.topn_per_rule`
 Top-N relevant volume prefixes `core.outage_detection.volume_top_n`
 Trace all providers `explorer.trace.all`
 Traceroute algorithms order `explorer.trace.algorithms`
 Traceroute max hops `explorer.traceroute.ttl.max`
 Traceroute min hops `explorer.traceroute.ttl.min`
 Traceroute packets per hop `explorer.traceroute.sendpkts`
 Traceroute retry packets `explorer.traceroute.retrypkts`
 Transit ASNs `bgpd.prefixlist.asn`
 Transit Improvement TTL max (sec) `core.improvements.inbound_transit.ttl.max`
 Transit Improvement TTL min (sec) `core.improvements.inbound_transit.ttl.min`
 Transit Improvements Max `core.improvements.inbound_transit.max`
 Transit SNMP hosts `bgpd.peer.X.transit.snmp`
 Transit matching at egress `collector.flow.process_transit_in_outbound`
 Transit peering partner mark `provider.X.rule.Y.transit`
 Transit prefixes `bgpd.prefixlist.prefixes`
 Transit top N prefixes `collector.flow.export.inbound_transit.topn`
 Transit traffic distribution histor `collector.export.archive_transit`
 Transiting traffic `bgpd.peer.X.transit.status, global.inbound_transit`
 Twilio API uri `pushd.sms.uri.twilio`
 Use BGP peer for flowspec PBR `peer.X.flowspec.pbr.use_bgp_peer`
 Use BMP data `peer.X.bmp`
 Use BMP to search for routes `peer.X.bmp.check_routes`
 Users filter `directory.X.user_filter`

VIP probing `policy.X.vip`

VIP reprobing period (sec) `core.vip.interval.probe`

Webhook URL `pushd.webhook.url`

Whitelist `irpdetectd.whitelist`

Withdraw improvements on warn `core.circuit.withdraw_on_warn`

min_delay status `collector.span.min_delay`

sFlow UDP port `collector.flow.listen.sf`

6.2 Configuration parameter index

glo

bgpd.peer.X.flowspec, 9, 56, 89, 96, 151, 265, 338, 416
bgpd.peer.X.flowspec.redirect_type, 9, 338, 416
global.agg_ipv4_max, 7, 160, 317, 336, 416
global.agg_ipv6_max, 7, 160, 317, 336, 417
global.aggregate, 7, 69, 160, 317, 336, 419
global.bw_overusage, 7, 317, 418
global.exchanges, 7, 317
global.exchanges.auto_config_interval, 7, 180, 295, 318, 417
global.failover, 7, 161, 318, 319, 323, 325–327, 340, 395, 396, 416
global.failover.log, 7, 318
global.failover.use_communities, 7, 318
global.failover_identity_file, 7, 98, 318, 416
global.failover_role, 8, 213, 318, 420
global.failover_slave.ip, 8, 161, 213, 319, 421
global.failover_slave.ipv6, 8, 319, 421
global.failover_slave.port, 8, 161, 213, 319, 421
global.failover_timer_fail, 8, 48, 161, 213, 319, 416
global.failover_timer_failback, 8, 48, 161, 213, 319, 416
global.flowspec, 8, 56, 89, 151, 265, 319, 320, 416
global.flowspec.pbr, 8, 21, 76, 89, 320, 416
global.frontend_acl, 8, 320, 416
global.frontend_acl_ips, 8, 320, 414
global.ignored.asn, 8, 94, 160, 276, 320, 321, 417
global.ignored.unannounced, 8, 320, 417
global.ignored_communities, 8, 94, 160, 276, 320, 321, 417
global.ignorednets, 8, 94, 160, 276, 320, 321, 417
global.improve_mode, 8, 69, 160, 204, 276, 321, 361, 417
global.inbound.injection, 8, 322, 417
global.inbound_conf, 8, 322
global.inbound_transit, 8, 164, 202, 286, 322, 421
global.ipv6_enabled, 8, 322, 417
global.master_management_interface, 8, 160, 204, 276, 323, 325, 417
global.master_probing_interface, 8, 150, 160, 204, 215, 276, 323, 326, 419
global.master_rd, 8, 47, 323, 418
global.nonintrusive_bgp, 8, 69, 89, 160, 276, 277, 323, 414
global.offpeak_hour, 8, 323, 363
global.outbound, 8, 324, 418
global.png.datadir, 8, 324
global.policies, 8, 324
global.policy.unpack_max, 8, 324
global.private_key, 8, 324
global.public_key, 8, 324
global.rd_rtt, 8, 38, 41, 325, 334, 401, 413, 419
global.rrd.age_max, 8, 325
global.rrd.datadir, 8, 325
global.slave_management_interface, 8, 323, 325, 421
global.slave_probing_interface, 8, 215, 323, 326, 421
global.slave_rd, 8, 47, 326, 421
global.user_directories_conf, 8, 326

com

db.clickhouse.dbname, 7, 315
db.clickhouse.host, 7, 315
db.clickhouse.http_port, 7, 315
db.clickhouse.password, 7, 315
db.clickhouse.port, 7, 315
db.clickhouse.username, 7, 315
db.dbname, 7, 316
db.host, 7, 316
db.ourhost, 7, 316
db.password, 7, 316
db.port, 7, 316
db.username, 7, 316

bgp

bgpd.as_path, 8, 94, 95, 178, 277, 289, 328–330, 346, 374, 390, 395, 397, 406, 414
bgpd.as_path_borrowing, 8, 329, 415
bgpd.db.timeout.withdraw, 8, 102, 103, 107, 329, 416
bgpd.full_control, 8, 330, 402, 417
bgpd.improvements.remove.bmp_exact_aggregate, 8, 329–331, 420
bgpd.improvements.remove.hold_time, 8, 330
bgpd.improvements.remove.next_hop_eq, 8, 163, 277, 330, 420
bgpd.improvements.remove.withdrawn, 9, 163, 277, 330, 331, 420
bgpd.improvements.strip_non_irp_communities, 9, 163, 331, 421
bgpd.log, 9, 331
bgpd.log.level, 9, 331
bgpd.mon.guardtime, 9, 26, 163, 277, 331, 395, 396, 414
bgpd.mon.holdtime, 9, 26, 163, 277, 331, 332, 395, 396, 414
bgpd.mon.internal.flap_guardtime, 9, 163, 332, 416
bgpd.mon.keepalive, 9, 26, 163, 277, 332, 395, 396, 414
bgpd.mon.longholdtime, 9, 26, 332, 414
bgpd.monitor.type, 9, 332
bgpd.no_export, 9, 90, 333, 418
bgpd.policy.cascade.amount, 9, 333, 415
bgpd.prefix.monitor.interval, 9, 332, 333
bgpd.prefix.monitor.search_interval, 9, 333
bgpd.prefixlist.asn, 9, 202, 286, 334, 421
bgpd.prefixlist.prefixes, 9, 202, 286, 334, 421
bgpd.rd_local_mark, 9, 39, 325, 334, 401, 413, 415
bgpd.retry_probing.new.bmp_path_change, 9, 95, 163, 334, 335, 419
bgpd.retry_probing.old.bmp_path_change, 9, 95, 163, 334, 335, 419
bgpd.scaninterval, 9, 330, 335
bgpd.snmp.concurrent_requests, 9, 335
bgpd.snmp.packets_interval, 9, 335
bgpd.snmp.simultaneous, 9, 336
bgpd.transit.monitor.election_interval, 9, 336
bgpd.transit.monitor.fast_reconfirm_interval, 9, 336
bgpd.transit.monitor.retries, 9, 336
bgpd.transit.monitor.timeout, 9, 336
bgpd.updates.split, 9, 163, 317, 336, 345, 364, 421

bgp1

bgpd.peer.X.as, 9, 164, 205, 278, 337, 414
bgpd.peer.X.blackholing.ipv4.next_hop, 9, 56, 96, 337, 415
bgpd.peer.X.blackholing.ipv6.next_hop, 9, 56, 96, 337, 415
bgpd.peer.X.blackholing.localpref, 9, 56, 337, 414
bgpd.peer.X.cap_4byte_as, 9, 337
bgpd.peer.X.inbound.ipv4.next_hop, 9, 164, 338, 403, 417
bgpd.peer.X.inbound.ipv6.next_hop, 9, 164, 339, 403, 417
bgpd.peer.X.inbound.master_localpref, 9, 164, 339, 414
bgpd.peer.X.inbound.slave_localpref, 9, 164, 339, 414
bgpd.peer.X.keepalive, 9, 164, 278, 339, 417
bgpd.peer.X.listen, 9, 339
bgpd.peer.X.master_communities, 9, 165, 206, 214, 278, 318, 340, 342, 417
bgpd.peer.X.master_localpref, 9, 164, 165, 206, 214, 278, 318, 340, 342, 417
bgpd.peer.X.master_our_ip, 9, 164, 165, 205, 214, 278, 340, 342, 417
bgpd.peer.X.master_our_ipv6, 9, 164, 278, 340, 343, 417
bgpd.peer.X.master_password, 9, 165, 214, 341, 343, 414
bgpd.peer.X.master_peer_ip, 9, 164, 205, 278, 341, 344, 420
bgpd.peer.X.master_peer_ipv6, 9, 164, 278, 341, 344, 420
bgpd.peer.X.master_router_id, 9, 341, 344, 414
bgpd.peer.X.med, 9, 165, 206, 278, 342, 417
bgpd.peer.X.origin, 9, 342
bgpd.peer.X.shutdown, 9, 342, 421
bgpd.peer.X.slave_communities, 9, 165, 214, 318, 340, 342, 417
bgpd.peer.X.slave_localpref, 9, 165, 214, 340, 342, 417
bgpd.peer.X.slave_our_ip, 9, 165, 214, 340–342, 417
bgpd.peer.X.slave_our_ipv6, 9, 341, 343, 417
bgpd.peer.X.slave_password, 9, 165, 214, 341, 343, 414
bgpd.peer.X.slave_peer_ip, 10, 164, 344, 420
bgpd.peer.X.slave_peer_ipv6, 10, 164, 344, 420
bgpd.peer.X.slave_router_id, 10, 165, 278, 341, 344, 414
bgpd.peer.X.transit.mib, 10, 344
bgpd.peer.X.transit.snmp, 10, 164, 345, 421
bgpd.peer.X.transit.status, 10, 345, 421
bgpd.peer.X.updates.limit.max, 10, 317, 337, 345
bgpd.peer.X.updates.limit.ps, 10, 345

col

collector.detect.explorer_ips, 10, 346
collector.export.archive_inbound, 10, 347, 417
collector.export.archive_transit, 10, 347, 421
collector.export.top_volume_ips, 10, 130, 347
collector.export.ttl, 10, 347
collector.export.volume.high.top_n, 10, 166, 280, 347, 370, 421
collector.export.volume.min, 10, 166, 280, 347, 348, 418
collector.export.volume.min_pct, 10, 166, 280, 347, 348, 418
collector.flow.all_outbound, 10, 348, 418
collector.flow.buffer.size, 10, 348
collector.flow.enabled, 10, 71, 203, 348, 351, 416
collector.flow.export.inbound_transit.topn, 10, 202, 286, 348, 421
collector.flow.listen.nf, 10, 71, 166, 203, 280, 349, 418
collector.flow.listen.sf, 10, 71, 166, 203, 280, 349, 422
collector.flow.log, 10, 349
collector.flow.log.level, 10, 349
collector.flow.process_transit_in_outbound, 10, 202, 286, 349, 421
collector.flow.sources, 10, 71, 166, 203, 280, 349, 416
collector.ournets, 10, 71, 75, 167, 203, 280, 348, 350, 403, 414
collector.sessions.max, 10, 350
collector.span.buffer.size, 10, 350
collector.span.enabled, 10, 75, 203, 348, 350, 351, 420
collector.span.interfaces, 10, 74, 75, 166, 203, 280, 351, 417
collector.span.log, 10, 351
collector.span.log.level, 10, 351
collector.span.min_delay, 10, 75, 166, 203, 280, 351, 352, 422
collector.span.min_delay.probing_queue_size, 10, 166, 203, 280, 351, 352, 418
collector.span.size_from_ip_header, 10, 75, 352, 418
collector.span.threshold.blackout, 10, 352
collector.span.threshold.congestion, 10, 352
collector.span.threshold.delay, 10, 352
collector.span.threshold.excessive, 10, 352, 353
collector.speaking_ips, 10, 353

COR

core.circuit.high_loss_diff, 10, 170, 283, 354–356, 415
core.circuit.hist_interval, 10, 170, 283, 354, 417
core.circuit.inbound, 10, 354, 355, 419
core.circuit.recover_hold_time, 10, 171, 283, 354, 420
core.circuit.recover_loss_diff, 10, 170, 283, 355, 356, 415
core.circuit.recover_monitored_intervals, 10, 171, 284, 355, 420
core.circuit.transit, 10, 354, 355, 419
core.circuit.warn_loss_diff, 10, 170, 283, 355, 415
core.circuit.withdraw_on_warn, 11, 170, 283, 356, 422
core.commit_control, 11, 58, 60, 69, 322, 356, 389, 392, 400, 415
core.commit_control.agg_bw_min, 11, 30, 62, 172, 284, 356, 418
core.commit_control.del_irrelevant, 11, 357, 415
core.commit_control.group_balance.check_monitor, 11, 357, 416
core.commit_control.inbound.enabled, 11, 200, 285, 357, 415
core.commit_control.inbound.improvement.delay, 11, 357, 417
core.commit_control.inbound.moderated, 11, 200, 285, 358, 417
core.commit_control.inbound.rate.high, 11, 358, 417
core.commit_control.inbound.rate.low, 11, 358, 417
core.commit_control.inbound.volume_estimation, 11, 358, 417
core.commit_control.probe_ttl, 11, 356, 359, 415
core.commit_control.probing_queue_size, 11, 172, 284, 359, 415
core.commit_control.rate.group, 11, 356, 359, 419
core.commit_control.rate.high, 11, 356, 358–360, 407, 418
core.commit_control.rate.low, 11, 59, 356, 358–360, 408, 418
core.commit_control.react_on_collector, 11, 60, 360, 420
core.cost.worst_ms, 11, 168, 281, 361, 362, 414
core.eventqueue.limit, 11, 168, 281, 361, 416
core.eventqueue.limit.retry_probe_pct, 11, 30, 361
core.eventqueue.limit_ipv6, 11, 361, 416
core.flowspec.max, 11, 151, 265, 361, 418
core.flowspec.max_ipv6, 11, 151, 265, 362, 418
core.global.allow_commit, 11, 362, 416
core.global.allow_latency_cost, 11, 362, 416
core.global.worst_ms, 11, 41, 361, 362, 416
core.improvements.clearslots, 11, 362, 363
core.improvements.clearslots.days_max, 11, 362, 363
core.improvements.inbound_transit.max, 11, 202, 286, 363, 421
core.improvements.inbound_transit.ttl.clean, 11, 363
core.improvements.inbound_transit.ttl.max, 11, 202, 286, 363, 421
core.improvements.inbound_transit.ttl.min, 11, 202, 286, 364, 421
core.improvements.max, 11, 30, 168, 281, 336, 337, 364, 418
core.improvements.max_ipv6, 11, 30, 168, 281, 364, 418
core.improvements.safe_removal, 11, 364, 420
core.improvements.ttl.retry_probe, 11, 30, 168, 281, 364, 421
core.log, 11, 30, 365
core.log.level, 11, 365
core.outage_detection, 11, 69, 365, 369, 418
core.outage_detection.limit_pct, 11, 365, 418
core.outage_detection.volume_top_n, 11, 365, 421
core.overusage.check_interval, 11, 169, 282, 365, 418
core.overusage.hold_timer, 11, 169, 282, 366, 418
core.overusage.out.average.period, 11, 169, 282, 366, 419
core.overusage.out.average.relevant_min, 11, 169, 283, 366, 419
core.overusage.out.threshold.throttle, 11, 170, 283, 366, 418
core.overusage.out.threshold.trigger, 11, 170, 283, 366, 418
core.performance.loss_pct, 11, 168, 281, 358, 360, 367, 420
core.performance.rtt.diff_ms, 11, 168, 281, 367, 420
core.performance.rtt.diff_pct, 11, 168, 281, 367, 420
core.performance.rtt.ix_diff_ms, 11, 367, 368, 420
core.performance.rtt.ix_diff_pct, 11, 368, 420
core.probes.ttl.failed, 11, 368, 416
core.probes.ttl.max, 11, 168, 281, 368, 418
core.probes.ttl.min, 11, 168, 281, 368, 418
core.problem.outage_timeout, 11, 369, 418
core.problem.rtt.diff_pct, 12, 369, 418
core.vip.interval.probe, 12, 29, 168, 281, 369, 422

exp

explorer.aipi, 12, 369
explorer.algorithms, 12, 370, 416
explorer.high_vol_precedence, 12, 173, 288, 370, 416
explorer.infra_ips, 12, 88, 173, 203, 287, 370, 417
explorer.interval.infra, 12, 370
explorer.interval.other, 12, 370
explorer.interval.other.trace, 12, 371
explorer.ipv4_test, 12, 371, 396, 419
explorer.ipv6_test, 12, 371, 397, 419
explorer.log, 12, 371
explorer.log.level, 12, 371
explorer.max_collector_ips, 12, 173, 288, 371, 418
explorer.maxthreads, 12, 173, 287, 370, 371, 416
explorer.probe.algorithm, 12, 173, 287, 372, 374, 419
explorer.probe.indirect_priority, 12, 372, 417
explorer.probing.sendpkts.adaptive_max, 12, 173, 288, 372, 414
explorer.probing.sendpkts.min, 12, 173, 288, 372, 418
explorer.probing.simultaneous, 12, 373
explorer.scanning.confirm_ips, 12, 373, 421
explorer.scanning.replypkts.min, 12, 373, 421
explorer.scanning.rtt.dispersion_ms, 12, 373, 421
explorer.scanning.sendpkts.factor, 12, 373, 420
explorer.timeout, 12, 173, 288, 373, 416
explorer.timeout.infra, 12, 374
explorer.trace.algorithms, 12, 173, 287, 374, 421
explorer.trace.all, 12, 329, 372, 374, 421
explorer.trace.diag_hop_unique, 12, 374, 394, 396, 415
explorer.traceroute.retrypkts, 12, 173, 288, 375, 421
explorer.traceroute.sendpkts, 12, 173, 288, 375, 421
explorer.traceroute.simultaneous, 12, 375
explorer.traceroute.simultaneous.infra, 12, 375
explorer.traceroute.ttl.max, 12, 173, 288, 375, 421
explorer.traceroute.ttl.min, 12, 173, 288, 375, 421

adm

collector.snmp.enhanced.sec, [10](#), [350](#)

dbcron.api.log, [13](#), [387](#)

dbcron.api.socket, [13](#), [388](#)

dbcron.log, [13](#), [388](#)

dbcron.log.level, [13](#), [388](#)

globalcc.log, [13](#), [388](#)

globalcc.log.level, [14](#), [388](#)

irpdetectd.log, [14](#), [388](#)

irpdetectd.log.level, [14](#), [389](#)

ups

peer.X.95th, 14, 175, 208, 211, 213, 290, 356, 384, 385, 389, 390, 397, 400, 407, 419
peer.X.95th.bill_day, 14, 175, 290, 389, 400, 419
peer.X.95th.centile, 14, 175, 290, 390, 400, 415
peer.X.95th.in, 14, 175, 290, 390, 419
peer.X.95th.mode, 14, 200, 285, 358, 390, 400, 419
peer.X.aspath_for_ix, 14, 329, 390, 414
peer.X.auto_config, 14, 180, 295, 318, 390, 417
peer.X.bgp_peer, 14, 174, 182, 207, 289, 295, 391, 420
peer.X.blackholing.bgp_peer, 14, 96, 271, 391, 415
peer.X.blackholing.community, 14, 56, 96, 271, 391, 415
peer.X.bmp, 14, 334, 335, 391, 421
peer.X.bmp.check_routes, 14, 29, 95, 391, 421
peer.X.cc_disable, 14, 60, 174, 175, 208, 211, 213, 288, 290, 322, 391, 400, 415
peer.X.circuit.control, 14, 174, 289, 392, 415
peer.X.cost, 14, 174, 175, 208, 290, 392, 400, 419
peer.X.description, 14, 174, 207, 289, 392, 419
peer.X.diag_hop.interval_max, 14, 369, 392
peer.X.diag_hop.interval_min, 14, 369, 393
peer.X.flow_agents, 14, 70, 174, 195, 289, 325, 334, 348, 393, 401, 413, 416
peer.X.flowspec.ipv4.redirect_community, 14, 151, 265, 393, 416
peer.X.flowspec.ipv6.redirect_community, 14, 151, 265, 393, 417
peer.X.flowspec.pbr.enabled, 14, 393, 416
peer.X.flowspec.pbr.use_bgp_peer, 14, 393, 421
peer.X.global_group, 14, 394, 400, 419
peer.X.group_loadbalance, 14, 60, 61, 394, 400, 417
peer.X.improve_in_group, 14, 175, 208, 213, 290, 394, 400, 418
peer.X.inbound.community_base, 14, 394, 414
peer.X.ipv4.diag_hop, 14, 177, 207, 289, 374, 394, 396, 417
peer.X.ipv4.master_probing, 14, 150, 174, 178, 207, 214, 289, 395, 419
peer.X.ipv4.mon, 14, 27, 28, 177, 210, 291, 395, 416
peer.X.ipv4.next_hop, 14, 26, 178, 207, 289, 394, 395, 398, 420
peer.X.ipv4.next_hop_as, 14, 178, 207, 289, 329, 395, 406, 419
peer.X.ipv4.route_server, 14, 396, 420
peer.X.ipv4.slave_probing, 14, 214, 395, 396, 419
peer.X.ipv4_pbr_check, 14, 396, 397, 414
peer.X.ipv6.diag_hop, 14, 177, 208, 289, 374, 394, 396, 417
peer.X.ipv6.master_probing, 14, 150, 174, 178, 208, 289, 396, 397, 419
peer.X.ipv6.mon, 14, 27, 28, 177, 210, 291, 396, 416
peer.X.ipv6.next_hop, 14, 26, 208, 396–398, 420
peer.X.ipv6.next_hop_as, 14, 208, 329, 397, 406, 419
peer.X.ipv6.route_server, 14, 397, 420
peer.X.ipv6.slave_probing, 14, 397, 419
peer.X.ipv6_pbr_check, 14, 396, 397, 419
peer.X.limit_load, 14, 174, 175, 208, 211, 213, 289, 389, 397, 418
peer.X.mon.ipv4.bgp_peer, 14, 28, 177, 210, 291, 398, 414
peer.X.mon.ipv4.external.state, 14, 27, 176, 210, 291, 398, 416
peer.X.mon.ipv4.internal.mode, 14, 177, 291, 396, 398, 399, 414
peer.X.mon.ipv4.internal.state, 14, 27, 177, 210, 291, 396, 398, 416
peer.X.mon.ipv6.bgp_peer, 14, 28, 177, 291, 398, 414
peer.X.mon.ipv6.external.state, 14, 176, 210, 291, 398, 417
peer.X.mon.ipv6.internal.mode, 14, 177, 291, 322, 397, 399, 414
peer.X.mon.ipv6.internal.state, 14, 177, 291, 397, 399, 417
peer.X.mon.snmp, 15, 28, 177, 182, 291, 295, 398, 399, 417
peer.X.mon.vpn_instance_id, 15, 399, 421
peer.X.pbr_check, 15, 399, 418
peer.X.precedence, 15, 58, 60, 175, 208, 213, 290, 356, 400, 415
peer.X.rd, 15, 38, 174, 289, 325, 334, 401, 413, 419
peer.X.routes_config, 15, 401, 420
peer.X.shortname, 15, 174, 207, 212, 289, 401, 419
peer.X.shutdown, 15, 174, 288, 401, 419
peer.X.snmp.enhanced, 15, 401
peer.X.snmp.interfaces, 15, 182, 209, 295, 356, 397, 402, 420
peer.X.type, 15, 394, 402, 419

rou

policy.X.asn, 15, 404, 414
policy.X.cascade, 15, 404, 415
policy.X.community, 15, 404, 414
policy.X.country, 15, 404, 415
policy.X.enabled, 15, 404, 416
policy.X.forcelocal, 15, 404, 405, 417
policy.X.notes, 15, 404, 418
policy.X.policy, 15, 404, 405, 418
policy.X.prefix, 15, 405, 416
policy.X.priority, 15, 405, 418
policy.X.providers, 15, 405, 417
policy.X.vip, 15, 405, 422

exc

provider.X.rule.Y.bgp_peer, 15, 396, 397, 406, 416
provider.X.rule.Y.enabled, 15, 406, 418
provider.X.rule.Y.next_hop, 15, 406, 418
provider.X.rule.Y.next_hop_as, 15, 395, 397, 406,
418
provider.X.rule.Y.pbr_check, 15, 406, 418
provider.X.rule.Y.probing_dscp, 15, 406, 419
provider.X.rule.Y.probing_ip, 15, 406, 407, 419
provider.X.rule.Y.shortname, 15, 407, 418
provider.X.rule.Y.transit, 15, 407, 421

tra

pushd.email.from, 13, 380, 416
pushd.email.host, 13, 380, 416
pushd.email.port, 13, 380, 421
pushd.listen.port, 13, 381, 417
pushd.log, 13, 381
pushd.log.level, 13, 381
pushd.sms.account_sid, 13, 381, 414
pushd.sms.auth_token, 13, 381, 421
pushd.sms.gateway, 13, 381, 420
pushd.sms.message_size, 13, 381, 418
pushd.sms.phone_number, 13, 382, 416
pushd.sms.uri.plivo, 13, 382, 418
pushd.sms.uri.twilio, 13, 382, 421
pushd.templates.datadir, 13, 382
pushd.webhook.avatar_emoji, 13, 382, 383, 414
pushd.webhook.avatar_url, 13, 382, 383, 414
pushd.webhook.botname, 13, 383, 415
pushd.webhook.url, 13, 383, 422

trap.bgpdpd_announced_rate_low.limit_pct, 13, 383, 419
trap.core_cc_improvements_spike.diff_pct, 13, 383, 384, 421
trap.core_cc_improvements_spike.period_sec, 13, 383, 384, 421
trap.core_cc_overload.inbound_limit_mbps, 13, 384, 417
trap.core_cc_overload.inbound_limit_pct, 13, 384, 417
trap.core_cc_overload.limit_mbps, 13, 384, 418
trap.core_cc_overload.limit_pct, 13, 384, 418
trap.core_cc_provider_overload.inbound_limit_mbps, 13, 385, 419
trap.core_cc_provider_overload.inbound_limit_pct, 13, 385, 419
trap.core_cc_provider_overload.limit_mbps, 13, 385, 419
trap.core_cc_provider_overload.limit_pct, 13, 385, 419
trap.core_excessive.prefixes, 13, 385, 419
trap.core_excessive_latency.limit, 13, 386, 419
trap.core_excessive_loss.limit, 13, 386, 419
trap.destination.auth_password, 13, 386, 387, 420
trap.destination.auth_protocol, 13, 386, 387, 420
trap.destination.auth_username, 13, 386, 387, 420
trap.destination.community, 13, 386, 420
trap.destination.port, 13, 386, 420
trap.destination.priv_password, 13, 387, 420
trap.destination.priv_protocol, 13, 387, 420
trap.destination.seclvl, 13, 386, 387, 420

api

apid.allowed_ips, 8, 326
apid.grpc.port, 8, 326
apid.listen.master_ip, 8, 326, 327
apid.listen.port, 8, 327
apid.listen.slave_ip, 8, 326, 327
apid.log, 8, 327
apid.log.level, 8, 327
apid.maxthreads, 8, 327
apid.path.mtr, 8, 327
apid.path.ping, 8, 328
apid.path.ping6, 8, 328
apid.path.traceroute, 8, 328

globalcc.key_service.port, 13, 388
globalcc.service.port, 14, 388

irpbmpd.grpc.port, 10, 345
irpmng.grpc.port, 14, 389

inb

bgpd.peer.X.inbound.ipv4.next_hop, 9, 164, 338, 403, 417
bgpd.peer.X.inbound.ipv6.next_hop, 9, 164, 339, 403, 417
bgpd.peer.X.inbound.master_localpref, 9, 164, 339, 414
bgpd.peer.X.inbound.slave_localpref, 9, 164, 339, 414

core.commit_control.inbound.enabled, 11, 200, 285, 357, 415
core.commit_control.inbound.improvement.delay, 11, 357, 417
core.commit_control.inbound.moderated, 11, 200, 285, 358, 417
core.commit_control.inbound.rate.high, 11, 358, 417
core.commit_control.inbound.rate.low, 11, 358, 417
core.commit_control.inbound.volume_estimation, 11, 358, 417
core.commit_control.loss_override, 11, 58, 358, 360, 415
core.commit_control.worst_loss, 11, 360, 414

global.inbound.injection, 8, 322, 417
global.inbound_conf, 8, 322

inbound.rule.X.bgp_peer, 15, 196, 402, 420
inbound.rule.X.communities, 15, 402, 414
inbound.rule.X.enabled, 15, 196, 402, 419
inbound.rule.X.full_control, 15, 330, 402, 415
inbound.rule.X.next_hop, 15, 196, 338, 339, 403, 417
inbound.rule.X.prefix, 15, 196, 403, 418
inbound.rule.X.providers, 15, 196, 403, 419

peer.X.95th.in, 14, 175, 290, 390, 419
peer.X.95th.mode, 14, 200, 285, 358, 390, 400, 419
peer.X.aspath_for_ix, 14, 329, 390, 414
peer.X.circuit.control, 14, 174, 289, 392, 415
peer.X.inbound.community_base, 14, 394, 414

ud

directory.X.admin_role_groups, 15, 408, 410, 415
directory.X.base_dn, 15, 408, 410, 415
directory.X.email, 15, 408, 415
directory.X.fullname, 15, 408, 415
directory.X.hostname, 15, 408, 415
directory.X.initial_bind_dn, 15, 409, 415
directory.X.initial_bind_password, 15, 409, 415
directory.X.name, 15, 409, 415
directory.X.order, 16, 409, 415
directory.X.port, 16, 409, 415
directory.X.secret, 16, 409, 421
directory.X.state, 16, 409, 415
directory.X.timeout, 16, 409, 415
directory.X.tls, 16, 410, 415
directory.X.tls_cacertfile, 16, 410, 415
directory.X.type, 16, 410, 416
directory.X.user_dn, 16, 408, 410, 415
directory.X.user_filter, 16, 410, 421
directory.X.user_role_attr, 16, 408, 410, 416
directory.X.username, 16, 408, 410, 416
directory.X.verify_cert, 16, 410, 415

snmp

snmp.X.auth_password, 16, 176, 183, 291, 411, 412, 420
snmp.X.auth_protocol, 16, 176, 183, 291, 411, 412, 420
snmp.X.auth_username, 16, 176, 183, 291, 411, 412, 420
snmp.X.community, 16, 183, 411, 420
snmp.X.context, 16, 411, 420
snmp.X.ip, 16, 183, 411, 420
snmp.X.name, 16, 183, 399, 411, 420
snmp.X.port, 16, 412, 420
snmp.X.priv_password, 16, 176, 183, 291, 412, 420
snmp.X.priv_protocol, 16, 176, 183, 291, 412, 420
snmp.X.seclvl, 16, 176, 183, 291, 411, 412, 420
snmp.X.timeout, 16, 412, 421

rou1

rd.X.community.local, [16](#), [334](#), [413](#), [415](#)
rd.X.community_worsening, [16](#), [41](#), [413](#), [419](#)
rd.X.shortname, [16](#), [413](#), [419](#)

bmp

[irpbmpd.log](#), 10, 346

[irpbmpd.log.level](#), 10, 346

[irpbmpd.port](#), 10, 346

[irpbmpd.sources](#), 10, 346

glo1

globalgroup.X.members, [15](#), [407](#), [416](#)
globalgroup.X.outbound.95th, [15](#), [407](#), [417](#)
globalgroup.X.outbound.95th.reserve, [15](#), [407](#), [414](#)
globalgroup.X.outbound.rate_high, [15](#), [407](#), [408](#),
[416](#)
globalgroup.X.outbound.rate_low, [15](#), [407](#), [408](#), [417](#)
globalgroup.X.shortname, [15](#), [408](#), [416](#)

inb1

irpinperfd.enabled, [12](#), [287](#), [376](#), [417](#)
irpinperfd.log, [12](#), [376](#)
irpinperfd.log.level, [12](#), [376](#)
irpinperfd.model.debug, [12](#), [376](#)
irpinperfd.model.topn_per_rule, [12](#), [287](#), [376](#), [421](#)
irpinperfd.moderated, [12](#), [287](#), [376](#), [417](#)
irpinperfd.probing.confirmation_interval, [12](#), [376](#),
[415](#)
irpinperfd.probing.failure_margin, [12](#), [287](#), [377](#), [419](#)
irpinperfd.probing.interval, [12](#), [51](#), [287](#), [377](#), [420](#)
irpinperfd.probing.shelf_life, [12](#), [287](#), [377](#), [419](#)
irpinperfd.probing.timeout, [12](#), [287](#), [377](#), [419](#)

tms

irpdetectd.bgp.reaction, 12, 96, 377, 415
irpdetectd.bgp.redirect.bgp_peers, 12, 96, 377, 414
irpdetectd.bgp.redirect.communities, 12, 96, 377, 414
irpdetectd.blackhole.threshold.kpps, 12, 96, 378, 414
irpdetectd.blackhole.threshold.mbps, 12, 96, 378, 415
irpdetectd.flowspec.ipv4.redirect, 12, 96, 378, 416
irpdetectd.flowspec.ipv6.redirect, 12, 96, 378, 416
irpdetectd.flowspec.reaction, 12, 96, 378, 415
irpdetectd.flowspec.threshold.kpps, 12, 96, 379, 416
irpdetectd.flowspec.threshold.mbps, 13, 96, 379, 416
irpdetectd.ip_batch_size, 13, 379
irpdetectd.ipv4.prefix_size, 13, 96, 379, 418
irpdetectd.ipv6.prefix_size, 13, 96, 379, 418
irpdetectd.mode, 13, 96, 379, 415
irpdetectd.protected_addresses, 13, 380, 419
irpdetectd.time.keep, 13, 380, 421
irpdetectd.time.monitor, 13, 380, 421
irpdetectd.whitelist, 13, 380, 422

List of Figures

1.2.1	IRP Components	18
1.2.2	System Dashboard	27
1.2.3	Error: Internal BGP Monitor not configured	28
1.2.4	Support for Centralized Route Reflectors	34
1.2.5	IRP configuration in a multi-homed network connected to transit providers as well as and Internet Exchange	35
1.2.6	City wide network	36
1.2.7	Regional network	36
1.2.8	Intercontinental network	36
1.2.9	Multiple routing domains	37
1.2.10	Failover high level overview	46
1.2.11	Unbalanced inbound/outbound peering	50
1.2.12	Fluctuating inbound traffic	50
1.3.1	Performance optimization algorithm	57
1.3.2	Cost optimization algorithm	57
1.3.3	More aggressive when actual exceeds schedule	59
1.3.4	Provider load balancing	60
1.3.5	Commit control of aggregated groups	61
1.3.6	Sorted Commit Control improvements	62
2.7.1	Flow export configuration	71
2.7.2	Span port configuration (on the router)	75
2.7.3	Span port configuration (on the switch)	75
2.8.1	PBR configuration: single router and separate probing Vlan	77
2.8.2	PBR configuration: two routers and separate probing Vlan	80
2.8.3	PBR configuration via GRE tunnels	81
2.8.4	PBR configuration: single router and separate probing Vlan	88
3.1.1	Frontend login form	108
3.1.2	Default Dashboard	109
3.2.1	Menu bar	110
3.2.2	Status bar	110
3.2.3	Common sidebar buttons	110
3.2.4	Email schedule configuration	111
3.2.5	Report filters	111
3.2.6	Time period selection	112
3.2.7	Adjusting the number of results per page	112
3.2.8	Dashboard buttons	113
3.2.9	Failover status of master node	113
3.2.10	Failover status of slave node	114
3.2.11	Search box	114
3.2.12	Dynamical global search results	114
3.2.13	Global search results	115
3.2.14	Precedence and improve in a group warning	115
3.2.15	Internal Monitor warning	115
3.2.16	Explorer performance warning	115

3.2.17	License expiration warning	115
3.3.1	Platform status	116
3.3.2	Interfaces status	116
3.3.3	List of improved prefixes	116
3.3.4	AS-Path problems	117
3.4.1	Reports menu	118
3.4.2	Platform overview	118
3.4.3	Platform performance	119
3.4.4	Loss improvements	120
3.4.5	Loss relative rates	120
3.4.6	Loss-improved destinations	121
3.4.7	Latency values (ms)	122
3.4.8	Latency values (%)	122
3.4.9	Latency destinations	122
3.4.10	Probe conversion rates	123
3.4.11	Current improvements	124
3.4.12	Current improvements map view	125
3.4.13	Improvements to Exchanges	125
3.4.14	Historical records	126
3.4.15	Providers overall	126
3.4.16	Provider performance	127
3.4.17	Provider daily performance	127
3.4.18	Cost overview	128
3.4.19	Congestion and outages	129
3.4.20	Top volume prefixes	129
3.4.21	Top volume AS	130
3.4.22	Top problem AS	130
3.4.23	Prefix statistics	131
3.4.24	Prefix statistics map view	131
3.4.25	AS statistics	132
3.4.26	Country statistics	132
3.4.27	Country statistics map view	133
3.4.28	Exchanges statistics	133
3.4.29	Current inbound improvements	134
3.4.30	Inbound Improvements history	134
3.4.31	Probes today	135
3.4.32	Monitors status	135
3.4.33	Monitor history	136
3.5.1	Graphs menu	137
3.5.2	Improvements by type	137
3.5.3	Performance improvements	138
3.5.4	Probed prefixes and Improvements	138
3.5.5	New and retry improvements	139
3.5.6	Improvements by probing source	139
3.5.7	Prefixes rerouted from provider	140
3.5.8	Prefixes rerouted to provider	140
3.5.9	Improvements by provider	141
3.5.10	Performance improvements by provider	142
3.5.11	Commit improvements by provider	142
3.5.12	Total and improved traffic	143
3.5.13	Bandwidth usage by provider	143
3.5.14	Providers bandwidth usage	144
3.5.15	Bandwidth usage and improvements	144
3.5.16	Total bandwidth usage	145
3.5.17	Inbound traffic distribution	145
3.5.18	Provider performance history	146
3.6.1	Routing Policies	147
3.6.2	Routing policy window	148

3.6.3	Static Routes	148
3.6.4	VIP policies	149
3.7.1	Flowspec policies	149
3.7.2	Flowspec policy popup	150
3.8.1	Throttling excessive bandwidth use with Flowspec policies	151
3.9.1	Maintenance windows	152
3.9.2	Maintenance window popup	153
3.9.3	Maintenance window sidebar	153
3.10.1	System events	154
3.11.1	Traceroute	155
3.11.2	Traceroute results	155
3.11.3	Traceroute exploring details results	156
3.11.4	Looking glass	157
3.11.5	Whois	157
3.11.6	Prefix Probing	158
3.11.7	Prefix probing results	158
3.12.1	Configuration editor: Global settings	160
3.12.2	Configuration editor: Ignored prefixes	161
3.12.3	Configuration editor: Global failover	162
3.12.4	Configuration editor: BGP settings	163
3.12.5	Configuration editor: New Router, general settings	164
3.12.6	Configuration editor: New Router, inbound settings	164
3.12.7	Configuration editor: New Router, advanced settings	165
3.12.8	Configuration editor: New Router, advanced settings	166
3.12.9	Configuration editor: Collector settings	167
3.12.10	Configuration editor: Core configuration	168
3.12.11	Configuration editor: Outage detection	169
3.12.12	Configuration editor: Overusage	169
3.12.13	Configuration editor: Circuit issues detection	170
3.12.14	Configuration editor: Commit Control	172
3.12.15	Providers Overall	172
3.12.16	Configuration editor: Explorer settings	173
3.12.17	Configuration editor: Providers configuration	174
3.12.18	Configuration editor: Adding a new provider, General settings	175
3.12.19	Configuration editor: Adding a new provider, Commit Control	176
3.12.20	Configuration editor: Adding a new provider, SNMP settings	176
3.12.21	Configuration editor: Adding a new provider, External Monitor	177
3.12.22	Configuration editor: Adding a new provider, Monitoring	177
3.12.23	Configuration editor: Adding a new provider, Commit Control	178
3.12.24	An Internet Exchange is similar to a provider and it requires configuration of the Router, the Probing IPs and the other usual attributes	179
3.12.25	The list of peering partners for the Exchange shows details about each of them and offers access to IRP Autoconfiguration feature and PBR ruleset generator for the router	180
3.12.26	PBR Generator	181
3.12.27	Warning about new peering partners that have been added to the Exchange	181
3.12.28	Configuration editor: SNMP hosts configuration	182
3.12.29	Configuration editor: Events configuration	183
3.12.30	Configuration editor: Email configuration	184
3.12.31	Configuration editor: SMS configuration	184
3.12.32	Plivo account ID and secret	185
3.12.33	Configuration editor: SNMP Traps configuration	185
3.12.34	Configuration editor: Web Hooks configuration	185
3.12.35	Notifications	186
3.12.36	Subscribe notifications	187
3.12.37	Subscribe notifications	188
3.12.38	Configuration editor: Security	189
3.12.39	Configuration editor: Users and Directories	189
3.12.40	Configuration editor: LDAP User Directory configuration	190

3.12.41	Configuration editor: Users and Directories	190
3.12.42	Configuration editor: Users and Directories	191
3.12.43	Configuration editor: TACACS+ configuration	192
3.12.44	Configuration: Adding a new user account	193
3.12.45	Configuration editor: Frontend	194
3.12.46	Configuration Editor: Manage currencies	194
3.12.47	Main menu: Inbound	195
3.12.48	Configuration editor: Inbound prefixes	197
3.12.49	Configuration editor: Control of individual inbound prefix	198
3.12.50	Configuration editor: Inbound prefixes control	198
3.12.51	Configuration editor: Base communities	199
3.12.52	Configuration editor: Inbound commit control	200
3.12.53	Configuration editor: Inbound BW estimation algorithms	200
3.12.54	Configuration editor: Inbound BW estimation algorithms	201
3.12.55	Configuration editor: Inbound optimization of transiting traffic parameters	201
3.12.56	Configuration editor: Initial Setup	203
3.12.57	Configuration editor: Setup Infrastructure IP addresses and Analyzed networks	203
3.12.58	Configuration editor: Configure Collector	204
3.12.59	Configuration editor: Improvement mode	204
3.12.60	Configuration editor: Management Interface	204
3.12.61	Configuration editor: Probing Interface	205
3.12.62	Configuration editor: Router name and AS	205
3.12.63	Configuration editor: Router IPv4 addresses	206
3.12.64	Configuration editor: Router announcing	206
3.12.65	Configuration editor: Choose a Router	207
3.12.66	Configuration editor: Provider name	207
3.12.67	Configuration editor: Provider IPv4 addresses	208
3.12.68	Configuration editor: Provider IPv6 addresses	208
3.12.69	Configuration editor: Provider Commit Control	209
3.12.70	Configuration editor: Provider Monitoring setup IPv4	209
3.12.71	Configuration editor: Provider Monitoring setup IPv6	209
3.12.72	Configuration editor: External Monitor setup IPv4	210
3.12.73	Configuration editor: External Monitor setup IPv6	210
3.12.74	Configuration editor: Internal Monitor setup	211
3.12.75	Configuration editor: Provider pre-checks	211
3.12.76	Configuration editor: Commit Control basic setup	212
3.12.77	Configuration editor: Commit Control provider precedence setup	212
3.12.78	Configuration editor: Commit Control groups setup	212
3.12.79	Configuration editor: Commit Control Provider overall	213
3.12.80	Configuration editor: Failover setup	213
3.12.81	Configuration editor: Failover probing IPs for providers	214
3.12.82	Configuration editor: Failover BGP session settings	214
3.12.83	Configuration editor: BGP announcement attributes	214
3.12.84	Configuration editor: Probing and management and interfaces	215
4.1.1	Creating initial user in GMI	216
4.1.2	Registering IRP instance in GMI	217
4.1.3	GMI Frontend login form	217
4.2.1	Frontend Sections	217
4.2.2	Time period selection	218
4.2.3	Report filters	219
4.2.4	IRP Instance Performance Stats Quick Reference	220
4.2.5	Global Search	220
4.2.6	Global Search results	221
4.2.7	Recent Searches	221
4.2.8	IRP instance license expiration warning example	221
4.3.1	Dashboard Selection	222
4.3.2	List of Dashboards	222

4.3.3	Creating a new dashboard	223
4.3.4	Rearranging Widgets on Dashboards	224
4.3.5	Widgets Library	224
4.3.6	Editing Widgets	225
4.3.7	Deleting Dashboards	225
4.4.1	Reports menu	226
4.4.2	Report options	226
4.4.3	Report export and print options	226
4.4.4	Report email delivery configuration	227
4.4.5	Save for Later Compare Button	227
4.4.6	Save for Later Compare Results	227
4.4.7	ASN Path Problems	228
4.4.8	ASN Statistics	228
4.4.9	ASN Traffic Percentage	229
4.4.10	Country Statistics	229
4.4.11	Country Statistics map view	230
4.4.12	Current Improvements	230
4.4.13	Exchanges Statistics	231
4.4.14	Historical Records	231
4.4.15	Improved Prefixes Within Probed	232
4.4.16	Improved Traffic Volume	232
4.4.17	Improvements to Exchanges	233
4.4.18	Instances Status	233
4.4.19	Monitor History	234
4.4.20	Monitor Status	234
4.4.21	Platform Overview	235
4.4.22	Prefix Improvement Rate by Cause	235
4.4.23	Prefix Statistics	235
4.4.24	Prefix Statistics map view	236
4.4.25	Probes Today	236
4.4.26	Providers Efficiency	237
4.4.27	Providers Performance and Probes (last 24 hours)	237
4.4.28	Top Problem ASN	238
4.4.29	Top Volume ASN	238
4.4.30	Top Volume Prefixes	239
4.5.1	Graphs menu	239
4.5.2	ASN Traffic Heatmap	240
4.5.3	Bandwidth Usage and Improvements	240
4.5.4	Bandwidth Usage per Provider	241
4.5.5	Commit Control Improvements by Provider	241
4.5.6	Improvements by Cause	242
4.5.7	Improvements by IP Version	242
4.5.8	Improvements by Probing Source	243
4.5.9	Latency Destination	243
4.5.10	Latency Improvements by Provider	244
4.5.11	Latency Values (ms)	244
4.5.12	Loss Improvements by Provider	245
4.5.13	Loss Rates	245
4.5.14	New and Retry Improvements	246
4.5.15	Overall Improvements Type by Provider	246
4.5.16	Overall Rerouted Prefixes by Provider	247
4.5.17	Performance Improvements	247
4.5.18	Prefixes Rerouted from Provider	248
4.5.19	Prefixes Rerouted to Provider	248
4.5.20	Probed and Improved Volumes	249
4.5.21	Probed Prefixes and Improvements	249
4.5.22	Probes per Provider	250
4.5.23	Probes Today	250

4.5.24	Problem Destinations	251
4.5.25	Provider Performance	251
4.5.26	Providers Bandwidth Usage	252
4.5.27	Providers Efficiency	252
4.5.28	Total and Improved Traffic	253
4.5.29	Total Bandwidth Usage	253
4.5.30	Unique Probed/Improved Prefixes	254
4.6.1	Inbound section	254
4.6.2	Current Inbound Improvements	255
4.6.3	Historical Inbound	255
4.6.4	Inbound Improvements History	256
4.6.5	Inbound Traffic Distribution	256
4.6.6	Moderated Inbound Improvements	257
4.6.7	Inbound Performance Feed	258
4.6.8	Inbound Performance History	258
4.6.9	Inbound Performance Rules	258
4.6.10	Creating an Inbound Performance Rule	260
4.6.11	Inbound Performance Report Bandwidth Graph	260
4.6.12	Inbound Performance Report Loss/Latency Graph	260
4.7.1	Routing Policies	261
4.7.2	Detailed routing policy view	262
4.7.3	Routing policy window	262
4.7.4	Routing Policy Type	263
4.8.1	Flowspec policies	263
4.8.2	Flowspec policy popup	264
4.8.3	Policies by Country	265
4.8.4	Policies by Country - Affected Prefixes view	266
4.9.1	Looking glass	267
4.9.2	Prefix Probing	267
4.9.3	Traceroute	268
4.9.4	Traceroute results	268
4.9.5	Traceroute exploring details results	268
4.9.6	Whois	269
4.10.1	Threat Monitor	270
4.10.2	Mitigation Rules	270
4.10.3	Add new rule	271
4.10.4	Mitigation feed	272
4.10.5	History	273
4.11.1	Global Commit groups configuration	273
4.11.2	IRP instance individual settings within the Global Group	274
4.11.3	Global Commit Group Settings	274
4.11.4	Global Commit summary	275
4.11.5	Global Commit Graphs	275
4.12.1	System events	275
4.12.2	Concurrent editing of the Configuration Settings	276
4.12.3	Configuration editor: Global settings	276
4.12.4	Configuration editor: Ignored prefixes	277
4.12.5	Configuration editor: Routing Domains	277
4.12.6	Configuration editor: BGP settings	278
4.12.7	Configuration editor: New Router, general settings	278
4.12.8	Configuration editor: New Router, advanced settings	279
4.12.9	Configuration editor: New Router, Inbound settings	279
4.12.10	Blackholing configuration settings	279
4.12.11	Configuration editor: Collector settings	280
4.12.12	Configuration editor: Collector Inbound Prefixes settings	281
4.12.13	Configuration editor: Core configuration	282
4.12.14	Configuration editor: Outage detection	282
4.12.15	Configuration editor: Overusage	282

4.12.16	Configuration editor: Circuit issues detection	283
4.12.17	Configuration editor: Commit Control	284
4.12.18	Providers Overall	285
4.12.19	Configuration editor: Inbound commit control	285
4.12.20	Configuration editor: Inbound BW estimation algorithms	285
4.12.21	Configuration editor: Inbound optimization of transiting traffic parameters	286
4.12.22	Configuration editor: Inbound configuration settings	287
4.12.23	Configuration editor: Explorer settings	288
4.12.24	Configuration editor: Providers configuration	288
4.12.25	Configuration editor: Adding a new provider, General settings	289
4.12.26	Configuration editor: Adding a new provider, IPv4 / IPv6 settings	290
4.12.27	Configuration editor: Adding a new provider, Commit Control	290
4.12.28	Configuration editor: Adding a new provider, SNMP settings	291
4.12.29	Configuration editor: Adding a new provider, External Monitor	291
4.12.30	Configuration editor: Adding a new provider, Internal Monitor	292
4.12.31	Configuration editor: Adding a new provider, Inbound	292
4.12.32	Configuration editor: Adding a new provider, Flowspec	292
4.12.33	Provider Blackholing configuration details	293
4.12.34	An Internet Exchange is similar to a provider and it requires configuration of the Router, the Probing IPs and the other usual attributes	294
4.12.35	The list of peering partners for the Exchange shows details about each of them and offers access to Autoconfiguration feature and PBR ruleset generator for the router	294
4.12.36	Configuration editor: SNMP hosts configuration	296
4.12.37	Configuration editor: Events configuration	296
4.12.38	Configuration editor: Email configuration	297
4.12.39	Configuration editor: SMS configuration	297
4.12.40	Plivo account ID and secret	298
4.12.41	Configuration editor: SNMP Traps configuration	298
4.12.42	Configuration editor: Web Hooks configuration	298
4.13.1	IRP Instances	299
4.13.2	Adding an instance	300
4.13.3	User Management	300
4.13.4	Adding users	301
4.13.5	User Directory configuration	302
4.13.6	User Directory Advanced Options	302
4.13.7	User Directory Bindings	303
4.13.8	User Directory Access Options	304
4.13.9	Access Restriction	304
4.13.10	Frontend	305
4.13.11	Senders	305
4.14.1	Reports and Graphs emailing subscription	305
4.14.2	Emailing subscription details	306
4.15.1	System events	306
4.16.1	Notifications	310
4.16.2	Subscribe notifications	310
4.17.1	User Profile	311
4.18.1	Maintenance windows menu button	312
4.18.2	Maintenance windows	313

4.18.3 Maintenance window popup	314
---	-----

List of configuration examples

1.1 Stop IRP Core and purge existing improvements	23
1.2 Switch to Intrusive Mode and adjust IRP Core and Bgpd parameters	23
1.3 Inserting test improvements	23
1.4 Remove next-hop-self route-map (RM-NHS) example	24
1.5 Restart IRP Bgpd	24
1.6 Show BGP information for specified IP address or prefix	24
1.7 Traceroute destination networks	25
1.8 Delete test improvements	25
1.9 Configure the maximum improvements limit and revert Bgpd configuration	25
1.10 Restart IRP Core and Bgpd	25
1.11 Delete test improvements	25
1.12 Switch the system to non-intrusive mode	25
1.13 Restart IRP Core and Bgpd	25
2.1 Global MLS settings configuration	71
2.2 Global NetFlow settings and export configuration	72
2.3 Per-interface NetFlow settings configuration	72
2.4 Flexible NetFlow monitor configuration	72
2.5 Flexible NetFlow exporter configuration	72
2.6 Flexible NetFlow sampler configuration	72
2.7 Per-interface Flexible NetFlow settings configuration	72
2.8 NetFlow configuration on Cisco 7200/3600 series routers	73
2.9 Ingress/egress flow export configuration on peering interfaces	73
2.10 Ingress flow export configuration	73
2.11 Configuring NetFlow export on Vyatta	73
2.12 Configuration of an interface for the flow accounting	74
2.13 Juniper flow export configuration	74
2.14 Per-interface NetFlow sampling configuration	74
2.15 Cisco IPv4 PBR configuration: 1 router, 2 providers	77
2.16 Cisco IPv4 PBR configuration for IOS XR: 1 router, 2 providers	77
2.17 Juniper IPv4 PBR configuration: 1 router, 2 providers	78
2.18 Vyatta IPv4 PBR configuration example	79
2.19 Vyatta (VC6.5 and up) IPv4 PBR configuration example	79
2.20 IRP server source-based routing using the 'ip' command	80
2.21 IRP server source-based routing using standard CentOS configuration files	80
2.22 Cisco IPv4 PBR configuration: 2 routers, 2 providers, separate Vlan	81
2.23 IRP-side IPv4 GRE tunnel CLI configuration	81
2.24 IRP-side IPv4 GRE tunnel configuration using standard CentOS configs	82
2.25 Router-side IPv4 GRE tunnel configuration (Vyatta equipment)	82
2.26 Router-side IPv4 GRE tunnel configuration (Cisco equipment)	82
2.27 Router-side IPv4 GRE tunnel configuration (Juniper equipment)	82
2.28 IRP server IPv4 source-based routing via GRE using the 'ip' command	82
2.29 IRP server IPv4 source-based routing via GRE using standard CentOS configuration files	83
2.30 Cisco IPv4 PBR configuration: 2 routers, 2 providers, separate Vlan	83
2.31 Cisco IPv4 PBR configuration template	83
2.32 Cisco IPv4 PBR configuration template	84
2.33 Juniper IPv4 PBR configuration template	85

2.34	PBR validation using ‘traceroute’	87
2.35	PBR validation using Explorer self check	87
2.36	PBR validation using ‘iptables’ and ‘traceroute’	87
2.37	Sample provider configuration	88
2.38	SNMP parameters validation	89
2.39	iBGP session configuration example	91
2.40	Vyatta IPv4 iBGP session configuration example	91
2.41	Vyatta IPv6 iBGP session configuration example	91
2.42	Vyatta IPv4 route-map for setting local-pref on the router	91
2.43	Vyatta IPv6 route-map for setting local-pref on the router	91
2.44	Cisco IPv4 iBGP session configuration example	91
2.45	Cisco IPv6 iBGP session configuration example	92
2.46	Cisco IPv4 route-map for setting local-pref on the router	92
2.47	Cisco IPv6 route-map for setting local-pref on the router	92
2.48	Limiting the number of received prefixes for an IPv4 neighbor on Cisco	92
2.49	Limiting the number of received prefixes for an IPv6 neighbor on Cisco	92
2.50	Juniper IPv4 iBGP session configuration example	92
2.51	Juniper IPv6 iBGP session configuration example	93
2.52	Juniper IPv4 route-map for setting local-pref on the router	93
2.53	Limiting the number of received prefixes for an IPv4 neighbor on Juniper	93
2.54	Reload IRP Bgpd configuration	94
2.55	Cisco: Configure BGP Additional Paths	95
2.56	Juniper: Configure BGP Additional Paths	95
2.57	Generate keys on \$IRPMaster	98
2.58	Install public key on \$IRPSlave	98
2.59	Check SSH certificate-based authentication works	98
2.60	Generate CA and certificates	98
2.61	Verify certificates	99
2.62	Install certificates in designated directories	99
2.63	Cross copy client key and certificates	99
2.64	Set file permissions for keys and certificates	100
2.65	Example \$IRPSlave configuration from template	100
2.66	Check MySQL on \$IRPSlave works correctly	100
2.67	Set \$IRPMaster as a first node for Multi-Master replication	101
2.68	Check MySQL on \$IRPMaster works correctly	101
2.69	Replication user and grants	101
2.70	Copy database root user configuration file	101
2.71	Copy database configuration files	102
2.72	Copy database data files	102
2.73	Copy differences of database files (OS with Systemd)	102
2.74	Copy differences of database files (OS with RC-files)	102
2.75	Set \$IRPMaster as replication slave	103
2.76	Starting replication slave on \$IRPMaster	104
2.77	Set \$IRPSlave as replication slave	104
2.78	Starting replication slave	105
2.79	Starting IRP services and Frontend (OS with Systemd)	105
2.80	Starting IRP services and Frontend (OS with RC-files)	105
2.81	Synchronize RRD	106
5.1	Error log	338