FRAFOS ABC SBC Handbook

Release 5.6

FRAFOS GmbH

May 06, 2026

Table of Contents

ABC SBC 5.6 Documentation index 6
1 5.6 Release Notes 6
1.1 Common container changes 6
1.2 Cluster Config Manager / GUI changes 6
1.3 ABC SBC changes 6
2 ABC SBC changelog 8
2.1 v5.6.0 8
2.1.1 Added 8
2.1.2 Fixed 9
2.1.3 Changed 11
2.1.4 Removed 12
3 Installation guide 13
3.1 Hardware Requirements 13
3.2 Deployment Modes 13
3.2.1 Single Node Mode 13
3.2.2 High Available (HA) Pair Mode 14
3.2.3 Cluster based solution 14
3.3 Container Installation 14
3.3.1 Configuring the SBC container 14
3.3.2 Podman 16
3.4 Post-installation steps 23
4 User guide 25
4.1 About the ABC Session Border Controller 25
4.1.1 How to Start? 26
4.1.2 Credits 27
4.2 Introduction 27
4.2.1 A Brief Introduction to History and Architecture of SIP 27
4.2.2 What is a Session Border Controller (SBC)? 31
4.2.3 Do You Need an SBC? 34
4.2.4 ABC SBC Networking Concepts 35
4.3 Practical Guide to the ABC SBC 40
4.3.1 Network Planning Guidelines 40
4.3.2 Planning Checklists 47
4.3.3 A Typical SBC Configuration Example 48
4.4 Initial Configuration 65
4.4.1 SBC Interfaces Overview 65
4.4.2 Web GUI Configuration (Cluster Config Master) 66
4.5 Setting Up Web Interface Access and User Accounts 67
4.5.1 Default User Accounts 68
4.6 ABC SBC License 68
4.7 Interface Configuration 69
4.7.1 Physical and System Interfaces 69
4.7.2 SBC Interfaces 71
4.7.3 Retro Compatibility 73

4.8 TLS profiles Configuration ... 75
4.8.1 TLS profile options ... 75
4.8.2 Certificate requirements ... 76
4.8.3 Let's encrypt gocertbot ... 77
4.9 Hardware Specific Configurations ... 80
4.9.1 Network adapters ... 80
4.9.2 Configuration of SBC Number of Threads ... 81
4.9.3 Configuration of sysctl settings ... 81
4.10 General ABC Configuration Guide ... 82
4.10.1 Physical, System and SBC Interfaces ... 82
4.10.2 Defining Rules ... 83
4.10.3 Using Replacements in Rules ... 90
4.10.4 Using Regular Expression Backreferences in Rules ... 96
4.10.5 Binding Rules together with Call Variables ... 97
4.10.6 SIP Routing ... 98
4.10.7 View A-B-C rules ... 113
4.10.8 SIP Mediation ... 114
4.10.9 SDP Mediation ... 124
4.10.10 Media Handling ... 128
4.10.11 NAT Traversal ... 142
4.10.12 Registration Caching and Handling ... 146
4.10.13 Call Data Records (CDRs) ... 157
4.10.14 Advanced Use Cases with Provisioned Data ... 159
4.10.15 SIP-WebRTC Gateway ... 173
4.10.16 Amazon Elastic Cloud Configuration Cookbook ... 186
4.10.17 Template parameters ... 195
4.11 ABC SBC System administration ... 197
4.11.1 User Management ... 197
4.11.2 Server Administration ... 200
4.11.3 Backup and Restore Operations ... 201
4.11.4 How to setup a Semi-redundant CCM on ABC SBC ... 206
4.11.5 Upgrade Procedure ... 209
4.11.6 Migration from 4.5/4.6 to 5.0 ... 210
4.11.7 SBC Dimensioning and Performance Tuning ... 213
4.11.8 Removing SBC Node ... 215
4.11.9 High Availability administration ... 216
4.11.10 API tokens ... 217
4.12 Monitoring and Troubleshooting ... 218
4.12.1 Overview of Monitoring and Troubleshooting Techniques ... 218
4.12.2 Live ABC SBC Information ... 219
4.12.3 Measurements and Monitoring ... 226
4.12.4 Additional Sources of Diagnostics Information ... 231
4.12.5 Viewing ABC SBC Logs ... 232
4.12.6 Coredumps ... 232
4.13 Securing SIP Networks using ABC SBC and ABC Monitor (optional) ... 233
4.13.1 SIP Security Principles: Collect, Analyze and Police ... 233
4.13.2 Police: Devising Security Rules in the ABC SBC ... 235
4.13.3 Collect Events: Gathering Usage Data in the ABC Monitor ... 255
4.13.4 Practices for Devising Secure Rule-basis ... 257
4.14 Preview of Experimental Features ... 263
4.14.1 Using Two-Factor Authentication for Users ... 263

4.14.2 AWS: Reputation Lists ... 271
4.14.3 Server Transaction limits ... 272
4.14.4 New restify CDR process ... 274
5 Reference book ... 276
5.1 Reference of Actions ... 276
5.1.1 SIP Mediation ... 278
5.1.2 SDP Mediation ... 294
5.1.3 Monitoring and Logging ... 299
5.1.4 Traffic Shaping ... 301
5.1.5 Media Processing ... 302
5.1.6 SIP Dropping ... 311
5.1.7 Scripting ... 312
5.1.8 Register Processing ... 312
5.1.9 External Interaction ... 312
5.1.10 NAT Handling ... 313
5.1.11 Other ... 313
5.2 Reference of Global Configuration Parameters ... 313
5.2.1 AWS Parameters ... 314
5.2.2 Backup Parameters ... 315
5.2.3 CDR Parameters ... 316
5.2.4 Conference Parameters ... 317
5.2.5 Eventbeat Parameters ... 318
5.2.6 Event Parameters ... 319
5.2.7 Firewall Parameters ... 322
5.2.8 Lawful Interception Parameters ... 325
5.2.9 Low-level Parameters ... 326
5.2.10 Miscellaneous Parameters ... 330
5.2.11 PCAP Parameters ... 331
5.2.12 RTP handling Parameters ... 331
5.2.13 SEMS Parameters ... 334
5.2.14 SIPREC Parameters ... 338
5.2.15 SIP Parameters ... 339
5.2.16 SRTP Parameters ... 343
5.2.17 Signaling SSL ... 346
5.2.18 Syslog Parameters ... 347
5.2.19 System Monitoring Parameters ... 349
5.2.20 Probe Parameters ... 351
5.3 Reference of Log Level Parameters ... 353
5.4 Reference of Call Agent Configuration Parameters ... 355
5.4.1 Destination Monitor Parameters ... 355
5.4.2 Failover Parameters ... 355
5.4.3 Registration Agent Parameters ... 356
5.4.4 Topology Hiding Parameters ... 356
5.4.5 Firewall Blacklisting Parameters ... 357
5.4.6 Security Parameters ... 357
5.4.7 SIP Timer Parameters ... 357
5.4.8 Resolver Parameters ... 357
5.5 Default Audio Files ... 358
5.5.1 Join meet-me conference ... 358
5.5.2 Meet-me set PIN audio prompts ... 360
5.5.3 Two-Factor authentication ... 361

5.6 Reference of Default Port Numbers ... 361
5.7 Reference Interface Parameters ... 363
5.8 Reference Application Interface Options ... 363
5.8.1 Supported by the current release ... 364
5.9 Command Line Reference ... 367
5.9.1 Configuration Management ... 368
5.9.2 User Management ... 368
5.9.3 Low-Level CLI ... 369
5.9.4 HA CLI ... 369
5.10 Reference of Used Open-Source Software ... 369
5.11 Reference Userdata Parameters for AWS Instances ... 371
5.12 Reference XML-RPC functions ... 371
5.12.1 Provisioned Tables ... 372
5.12.2 Call agents ... 373
5.12.3 TLS profiles ... 373
5.12.4 Nodes ... 373
5.12.5 Logical interfaces ... 374
5.12.6 System interfaces ... 374
5.12.7 Maintenance mode ... 374
5.13 Reference of CCM Configuration Parameters ... 374
5.13.1 Login ... 375
5.13.2 LDAP Parameters ... 376
5.13.3 Backup Parameters ... 382
5.13.4 Management access Parameters ... 383
5.13.5 SBC security Parameters ... 384
5.13.6 Email Parameters ... 385
5.13.7 Certbot Parameters ... 385
5.13.8 Miscellaneous Parameters ... 386
5.14 Reference of Token Capabilities ... 387
5.14.1 CCM API token capabilities ... 388
5.14.2 SBC API token capabilities ... 388
5.15 CCM configuration API ... 388
5.16 Reference of Supported Codecs ... 389
5.17 SIP Timers ... 389
5.17.1 Timer Definitions ... 389
5.17.2 Customizing Timer Values ... 390
5.17.3 SIP Timer Precision ... 390
5.17.4 Destination Monitoring ... 390
5.18 Legacy application ... 391
5.18.1 SSH ... 391
5.18.2 SNMP ... 391
5.18.3 TURN server for websocket ... 391
5.18.4 Local monitoring query service ... 392
5.18.5 PCAP query service ... 393
5.18.6 Local webconf API ... 393
5.18.7 Management for host ... 394
5.18.8 Log files provider ... 394
5.18.9 Local packet classifier ... 395
5.18.10 HTTP proxy ... 396
5.18.11 HTTP redirect ... 396
5.18.12 Prometheus Pull Service ... 397

6 Glossary 398

Chapter 1

5.6 Release Notes

For a detailed list of all changes and fixes, please check the changelog.

1.1 Common container changes

1.2 Cluster Config Manager / GUI changes

> **Warning**
>
> Upgrading the CCM to 5.6 is only supported from 5.5 and requires a fresh backup including provisioned tables to prevent data loss. Any data not included in the latest backup will not be migrated.

- Changed: the CCM has been rebuilt on a new architecture based on FrankenPHP and SQLite, replacing Nginx, PHP-FPM, and MariaDB. This brings significantly faster startup, shutdown, backup, and restore operations. Configuration publish, backup, and restore now run within a single SQL transaction. Upgrade is supported from version 5.5 only; see the warning above.
- Added: API tokens - allow creating bearer tokens for REST API access with configurable capabilities and optional expiration, see API tokens.
- Added: The CCM printable configuration report has been enhanced to include node information and the values of the cluster configuration parameters for each node.
- Added: support for multiple user branch LDAP authentication - previously, the `extra_group` setting could only be set to a single value, therefore enforcing all users that need to access the CCM to belong to the same organizational unit (LDAP wise). Now, `extra_group` may be set to one or more values, thus allowing authentication from multiple organizational units.
- Extended: the CCM REST API has been significantly expanded compared to 5.5, adding endpoints for configuration management, config activation and verification, push configuration, global config options, cluster config parameters, call agent attributes, template parameters, node hosts, logical interfaces, admission control rules, firewall, HA, licenses, file management, and CCM settings.
- Added: possibility to customize CDR format using global config variable.
- Fixed: LDAP and Active Directory login now works for users whose Distinguished Name contains spaces or other special characters (for example `CN=Thomas Danzer,CN=Users,DC=frafos,DC=net`), which previously failed silently.

1.3 ABC SBC changes

- Added: Role Based Access Control (RBAC) for the ABC SBC API. The ABC SBC API now requires bearer token authentication, with a new `sbc` token type and fine-grained capabilities (`sbc-config`, `sbc-provtables`, `sbc-status`, `sbc-system`). The Cluster Config Manager automatically provisions its own token; additional tokens can be created from the Cluster Config Manager API tokens screen for restricted, capability-based access (e.g. read-only status, provisioning tables push).
- Added: new Rust-based PCAP logger enabling continuous monitoring. Traces are now written in PCAPng format and automatically chunked into short time slices (10s default), making captures available faster and allowing long-running calls to be traced without waiting for them to end. The new

format also supports Name Resolution Blocks and opens the way for additional metadata in future releases.

- Changed: upload of PCAPs and recordings to the Monitor now uses a dedicated upload service instead of rsync, lifting a range of limitations related to rsync. This requires Monitor 10.2 or later (upload API support).
- Added: ABC SBC prometheus metrics now include redis' HA statistics. Please refere to documentation chapter Prometheus Statistics from HA redis for more details about the new entries.

Chapter 2

ABC SBC changelog

2.1 v5.6.0

2.1.1 Added

- Add support for AMR-WB transcoding. (G#9741)
  - Only single-channel is supported.
  - Supporting codec modes 0-8.
  - Supporting `mode-set` attribute.
  - Supporting `octet-aligned: 0|1`.
  - Supporting `mode-change-period` and signaling `mode-change-capability: 2`.
  - No support for `robust-sorting`.
  - No support for `interleaving`.
  - No support for `crc`.
- Prometheus: Added call-agent MOS QoS average s/m/l values as a new gauge. (G#9810)
- CCM GUI instantly display config changes made by concurent user (G#9639)
- Exposed number of active sessions as `abcsbc_sessions` in prometheus. (G#9859)
- Add four new better behaving operators to "SRTP Types" filter: all contain, partial contain, all do not contain, partial do not contain. (G#9420)
  Marked the old broken operators with "(deprecated)" suffix. The old ones gave wrong results with multiple m-lines such as failing to match `sdes` unless all m-lines contained `sdes`. I.e. "none" also matched if there're one sdes and one dtls m-line. The new operators are "Any media contain" and "All media do not contain" that behave correctly.
- In addition to weekly GeoIP database updates, SBC will also try to immediately download/update the GeoIP database when credentials are changed. (G#10080)
- New global configuration: DTLS CA list file (G#9855)
  When given, the remote DTLS certificates will be verified against the given CA list file.
- New filter condition: IP Address Operation. (G#10165)
  It supports 4 operators:
  - `==`: equal
  - `!=`: not equal
  - `is in`: left hand side IP is in the right hand side CIDR
  - `is not in`: left hand side IP is not in the right hand side CIDR
- More granular CDR log rotate. (G#10146)
- Add PCMU/16000 to transcoding. (G#10317)
- new signaling interface option `no_rport` to disable adding `rport` to Via header (G#10300)
- Inject redis replication statistics within prometheus metrics (G#10197)
- New options in "Log Received Traffic" action to log only RTP/DTMF or RTP/non-DTMF. (G#10338)
- New replacement expression: $_g (G#10325) Computes a reproducible unique string (UUIDv5) based on the provided value. (Ex: $_g($ci))
- Added new action "Add Contact header field parameter" for adding custom parameters to the Contact header field of outgoing SIP messages. (G#10325)
- Possibility to reset global config and CCM config options to default value (G#10387)

2.1.2 Fixed

- Codec whitelist/blacklist did not work well with dynamic payload numbers and transcoding. (G#9811)
  Consider the case where caller sends codec A on payload number 97 and codec whitelist/blacklist removes this codec on callee side. This causes the SBC to forget about payload number 97 and put 96 in the answer (to the caller) instead. If "Strict SDP answer" option of transcoding is used, the payload is not added to the response at all.
- RTP does not use configured interface after HA restore/failover after call-pickup with INV w/ Replaces due to wrong media interface being assigned. (Z#1999) (G#9825)
  The configured interface for the caller is forgotten and default (first) media interface is used.
- HA: Backup SBC created empty traffic log pcaps. (G#9869)
- Traffic logs are not put in intermediary directories after SBC restoration/fail-over. (G#9869)
- HA: For active calls, traffic was not recorded any more after failover. (G#9869)
- RTP traffic is not logged after SBC restoration/fail-over. (G#9869)
- "Change SSRC" flag was forgotten on HA/restoration. (Z#1999) (G#9903)
  Also affected "Sticky Stream SSRC" action as it needs the flag.
- "Map Replaces Header" action was unable to parse multiple headers in Refer-to uri-headers. (G#8971)
- "Sticky Stream SSRC" printed a warning and did not work with multiple streams. (G#9925)
- HA: Counters used by actions/filters such as "Limit Parallel Calls" were decremented twice on backup node which caused integer underflow, hitting the limits upon failover. (Z#2058) (G#9942)
- SBC `Content-Type` body filtering is buggy when the filtering ends up with a top body part that has header fields. It caused the SBC to either reject the request with 500 or put garbage data in the `Content-Type` header. (G#9957)
- SIPREC recording did not work after SBC HA failover/restoration. (G#9970)
- Call agent's blacklist queries do not use the "force transport" configuration when sending SRV requests. (G#9177)
- SIPREC signaling/media interface configuration in global config could fail silently. (Z#2061) (G#9799)
  The configuration would be applied even if the configured SIPREC signaling/media interface was not assigned to the node, falling back to the respective default interface. Now it will generate a configuration error on the SBC node.
- HA: After restoration, SBC forgets about "Media far end NAT traversal" and "Lock on addresses learned from RTP" setting in RTP anchoring. The former leads to RTP being sent to the original address in the SDP. (G#9948)
- SBC does not forward `rtpmap` attributes in m-lines with unknown transport types (i.e. `TCP/RTP/AVP`). (G#9537)
- Provisioned tables `range` queries did not work due to configuration generation problem. (G#10012)
- "Date and time" filter does not work properly. (G#10011)
  - When HH:MM:SS or YYYY-MM-DD format is used, it sometimes generated wrong results due to residual data from previous parsing attempts.
  - HH:MM:SS format did not take daylight saving into account.
- 5.4.18 and 5.5.5 introduced a resolver bug where TLS SRV is queried as `_sip._tcp` instead of `_sips._tcp`. (G#10051)
- SIPREC: `a=inactive` mark was put on the wrong stream. (G#9755)
- SIPREC: Did not consider the answer side (i.e. offer `sendrecv` & answer `sendonly`) when deciding whether the stream was on hold. (G#9755)
- SIPREC: In re-invites, SDP `o=` version was not increased even when SDP changes. (Z#1963) (G#9755)
- SIPREC: Did not consider `c=0.0.0.0` method when deciding whether the stream was on hold. (G#9755)
- HA: On SBC restoration, calls with SRTP lost audio due to a problem about restoring SRTP keys. (G#10018)

- HA - removed global config options for using also signaling interfaces for vrrp adverts and using macvlan interface for VIP and using unicast, replaced by one new option for HA mode, and fixed the macvlan sub interfaces scenario (G#9843)
- Use latest Redis 7.2.11. (G#10090)
- GeoIP filter lookup did not properly return the country iso code (also seen in DEBUG logs) due to wrong usage of a non-null-terminated string. This resulted in filter condition matching failure when it would normally match or vice-versa. (G#10080)
- End to end encryption doesn't work with HA / SBC restoration. (G#10138)
- End-to-end Encryption action prevented establishing calls with DTLS-SRTP. (G#10125)
- Call termination/cancellation caused by SIPREC "stop on error" option are now reflected in call events. (Z#1987) (G#9804)
  - `originator` field in `call-end` and `call-attempt` will have `recording-error`,
  - `hangup_cause` field in `call-end` will have `Recording Server  Error`.
  - `description` field in `call-end` and `call-attempt` will have `SIPREC call recording failed`.
  - `hangup_initiator` field in `call-end` events generated for calls terminated by a recording failure is changed to `internal`.
- Use new Redis 7.2.12. (G#10162)
- `dtls_legacy_server_connect` defaulted to `False` instead of `0` (G#10170)
- "DTLS handshake timeout" option not applied properly (G#10167)
- `Content-Type` Whitelist/Blacklist is not applied to initial INVITE when it is on C-rules. (G#10095)
- Destination Monitor could crash the SBC when trying to publish CA state without any information available. (Z#2147) (G#10192)
- Allow websocket port below 1024. (G#10222)
- Allow IP+prefix for SBC JSON configuration's HA destination (G#10245)
- SBC complained about being unable to load snmp libraries at start-up. (G#10179)
- SBC failed to handle SDP crypto lifetime parameters that are over 32-bit. (G#10224)
- SBC used Register Agent credentials in calls even when register agent was disabled. (G#10219)
- Events disk persistence. (G#10268)
- Potential deadlock with replacing filters such as Replace Header or Replace SDP Attribute. (G#10229)
- Make evenbeat use client certificate even if no CA certs is configured. (G#9812)
- Use global CRL with signaling interfaces. (G#10135)
- Wait until Redis is finished loading. (G#10063)
- HA: SUBSCRIBE/NOTIFY messages cause error logs on standby SBC (G#10199)
- Backup SBC terminates calls. (G#10273)
  It caused abrupt rekeying on backup SBC and that in turn caused SDES handshake failure. This caused randomly attempts to send BYE message on backup SBC to terminate the call (and even successfully sending the message depending on the container network configuration).
- Backup SBC tries to send BYE messages on start-up in case call data is missing (i.e. calls that were already being terminated). (G#10273)
- HA: Backup SBC tries to merge recent call updates are received in a small window but the debouncing logic did not work, possibly consuming higher CPU and memory than needed on the backup SBC. (G#10208)
- Attempt at resolving possible hanging calls with call replacement. (G#10193)
  With this change, when the call being replaced is missing its peer during processing a `Replaces:`, it sends a terminate request to the referrer and stops itself. Previously, this caused hanging of the call leg that's being replaced, if the referrer hanged up the call.
- SBC configuration failed when there are a lot of interfaces due to config reader line buffer limitation. (G#10336)
- GUI: Allow wraping of text without spaces in A/B/C rule list (G#10195)
- Allow template parameters in 'Nameserver IP addresses' call agent attribute (G#9997)

- Fixing tooltip of "Enable new version of CDRs (CDR-NG)" global config option (G#10237)
- CCM GUI: fixing problem with removing actions from rule (G#10392)
- CCM GUI: Re-Register all for RegAgents causes internal server error (G#10397)
- "Insert or Replace SIP Message Body (on leg)" action did not work for out of dialog messages and subscribe/notify messages. (G#8086)
- CCM: For High Availability, improving detection of split-brain. Do not report split-brain for offline nodes. (G#9871)
- CCM GUI: multiple restart warnings while adding system interface (G#9940)
- Fix name of "List of table to watch for expired rooms" global config attribute (G#9977)
- When changing `host` of a node, the node is highlighted in yellow color (G#9758)
- CCM: Discrad changes loads config backup also for CCM config (G#9919)
- Fixing popup about unsaved changes when leaving HA screen (G#10393)
- Fix wrong parameter name offered for call agent destination (G#9990)
- thread safety while reloading TLS certificates (G#10424) When reloading TLS certificates, a race condition could have enabled an old and defunct pointer to be used for a new connection (use-after-free).
- memory safety issue in SEMS configuration reader (G#10432)
- potential minor memory leak in SRTP/DTLS (G#10433)
- GUI: Value for monitoring interval is replaced by 0 (G#9994)
- CCM GUI: allow hostport in Set RURI host, Set To host, Set From host actions (G#10439)
- Printable config report: show defaults from gconfig for call agent attributes (G#10391)
- Failure to trigger LE profile renewal from UI toggle (G#10368)
- HA routes were deduplicated inside sbc JSON config (G#10296)
- SBC does not bind outbound created TCP connections to call-agents. (G#10395) On incoming requests SBC creates a correlation between the used source call agent and the newly created TCP socket and keeps it. This allows SBC to find the call agent by looking at the socket so that it's resillient against DNS changes as long as the TCP connection lives. The same mechanism is now used on the connections that are created through outbound connections that the SBC creates.
- GUI: display for A/B/C rules improved. Better highlighting of inactive rules. (G#10312)
- CCM GUI: improve placement of tooltip for `release` field to not cover buttons (G#10093)
- CCM GUI: fix input values for "Read call variables from table" rule action (G#10463)
- Call with INVITE w/ Replaces does not work after failover or restoration (broken audio or call is terminated). (Z#1999) (G#10428)
- REFER w/ Replaces does not work after failover or restoration (broken audio). (G#10438)
- Possible crash while using "Enable SIP Session Timers (SST) - caller leg" action without the "let remote refresh" option. (G#10455)
- Use after free during redis reconnection. (G#10491)
- enforcing "Required: 100rel" would lead to replying following re-INVITEs with "491 Request Pending" (G#10489)

2.1.3 Changed

- GUI: Info/error messages are displayed as toast (G#9621)
- Status screen improvements (G#9590)
  - The node table is narrower:
    - Buttons are replaced by icons
    - Moving 'master addr' value to status details
    - Long header titles are split to two rows
    - 'Status' and 'Configuration' columns are joined
  - Nodes are grouped by config group
- In "Add header" action, if the header value ends up empty as a result of replacement (i.e. empty call variable), header will not be added instead of resulting in an error. (G#9944)

- SBC now reloads the new geoip database as soon as it is updated. Previously a restart was required. (G#10101)
- AMR default mode is now bandwidth efficient (octet-aligned=0). (G#10231)
- The SBC will not log CDRs to container output any more when "Enable CDR" is unchecked. (G#10234)
- Support users from different branches (LDAP) (G#9174)
  Previously, the extra_group setting had to be set to a single value (e.g. ou=People) to allow users such as uid=john,ou=People,dc=example,dc=org to authenticate. With this change, extra_group can now be configured with one or more groups. This allows authentication from multiple organizational units, for example ou=People and ou=external,ou=contractors, enabling users like: uid=john,ou=People,dc=example,dc=org
  uid=june,ou=external,ou=contractors,dc=example,dc=org
- Restore legacy CDR format customization capability (G#10168)
- GUI: change some RTP global config options to checkboxes (G#10286)
  Following options are changed from text input to checkbox:
  * Drop unnegotiated RTP payloads
  * Do not drop unnegotiated comfort noise RTP payload
  * Force symmetric RTP for mediaserver apps
- Allow concurrent login to CCM GUI by default (G#10361)
- GUI - system interfaces - disable dropdown options if HA VIPs or AWS nodes are not available (G#10097)
- Enable "live calls" & "registration cache" redis related endpoints on HA Backup nodes (G#10194)
- GUI status screen: status icon is now clickable and display status details when clicked (G#9906)
- More friendly messages on LE certbot errors (G#10238)
- Update SBC go projects' toolchain to latest (G#10374)
- Prepare CI for new continious release process (G#10475)
- do not enforce 100rel required on re-INVITEs (G#10488)

2.1.4 Removed

- Date & Time filter condition: Removed interpretation of numeric-only values on the left operand as "days" since unix epoch. (G#10021)
- Removing no more used "Time for terminal session automatic logout if idle" global config option (G#6799)
- Mariadb timeout for "Read call variables" queries (s) global config option (G#10431)

Chapter 3

Installation guide

The ABC SBC is distributed in form of a container of OCI type. It can be run on operating system of customer choice, if the OS supports running of that container type.

FRAFOS also offers an Amazon cloud based solution where the ABC SBC is running as an instance in AWS (EC2). This is by far the fastest installation, the software can be started by several clicks. See Amazon Elastic Cloud Configuration Cookbook.

FRAFOS can also provide a hardware based solution with preinstalled ABC SBC software on a reference hardware, see Hardware Requirements for more details.

3.1 Hardware Requirements

FRAFOS ABC SBC is provided as container, internally based on Debian 13 64bit operating system (x86_64 architecture).

Capacity and performance of the system depends mainly on the number and type of processors (CPU), available operating memory (RAM) and the number and performance of network cards (NIC).

There are no specific constraints for vendors of hardware and components, but we do have some suggestions and recommendations for the used hardware and its settings. Generally amount of memory and CPU power increases system resilience against load peaks, Ethernet cards with high packet rate facilitate high media anchoring throughput and fast solid-state drives facilitate WAV and PCAP recording.

Minimum hardware:
- CPU: 1x processor - 64bit architecture
- RAM: 4 GB
- NIC: 1x 1Gb network card
- HDD: 10 GB

Recommended / reference hardware: Fujitsu Primergy RX1330 (M3 or M4)
- CPU: Intel® Xeon® processor E3-1200 family, 4GHz
- RAM: 64 GB (DDR4 2666 MHz, dual channel)
- NIC: 2 x Intel 1Gb adapter
- SSD: 2 x 256 GB

For more details about system capacity and dimensioning, see Sec. SBC Dimensioning and Performance Tuning.

3.2 Deployment Modes

According to the system dimensioning and high-availability requirements, ABC SBC can be deployed in different modes:

3.2.1 Single Node Mode

In single node mode a single ABC SBC container is used. It is necessary to connect this node to a CCM module which provides GUI for administration of system and serves as a configuration master to all connected ABC SBC nodes. The CCM module is distributed as a single container and can be either deployed on the same host together with ABC SBC node or on a different server.

3.2.2 High Available (HA) Pair Mode

IMPORTANT note: up to ABC SBC release 4.1, it was using HA solution based on Pacemaker. The ABC SBC 4.2 was a transitional release that removed Pacemaker based HA solution. The new HA solution based on keepalived was introduced in ABC SBC 4.3 release.

HA pair ABC SBC installation is formed by two containers running on two physically identical host servers, running in an active/hot-standby configuration. Only the active (HA master) server processes signaling and media traffic. In case of any device or network failure, the internal management system performs a failover where the originally standby (HA backup) machine becomes active. Switching the active and standby operation modes is also useful during the upgrades of the system.

Both the HA master and backup servers share virtual IP addresses (VIPs) and communicate with each other over the "Internal Management Interface" - IMI. The HA backup server can check the network availability of the HA master server. Once the HA backup server determines that the HA master server is no longer reachable over IMI interface then the backup server will assume the role of HA master and take over the VIPs used for receiving and sending the media and signaling messages.

Further, the HA master server replicates state information about running sessions to the HA backup machine. Thereby, after a failover the backup server will be able to continue processing already established calls and the failure of the server will not result in dropping of already established calls. Note however that calls are replicated after they are established and calls unanswered yet may be dropped. Also only signaling over connection-less transport protocols is certain to reach the Call Agents as transport protocol context gets lost during failover.

Note: status about non running SEMS process reported by SNMP from nodes in 'BACKUP' state should be ignored

3.2.3 Cluster based solution

For very high traffic and performance requirements, ABC SBC instances (in a single or HA pair mode) can create a cluster. A SIP load balancer is put in front of these cluster nodes so as to distribute the SIP traffic to a particular ABC SBC instances.

3.3 Container Installation

Before actual installation admin should consider enabling coredumps on the host, see Coredumps. It may be beneficial to have them enabled in advance, to ease later troubleshooting but it needs to be considered that whole host and all containers running there will be influenced.

3.3.1 Configuring the SBC container

The recommended way of configuring an ABC SBC container is via environment variables passed when creating the container (e.g. with `-e VAR=value` for podman or docker). They tell the ABC SBC how to reach the Cluster Config Manager and how to receive its initial configuration, and allow the container to come up unattended without any interactive setup step.

`MASTER` (required for new deployments)

Address of the Cluster Config Manager (FQDN or IP) used by the ABC SBC node to reach the configuration master. Can be omitted when a previously persisted configuration already exists under `/data`; otherwise the container starts with an empty configuration.

`CONFIG_USER` / `CONFIG_PASS` (required)

Credentials used for HTTP basic authentication on the Cluster Config Manager pullconf API, both for pulling configurations and for pushing node status to the Cluster Config Manager. They must match the credentials configured in the Cluster Config Manager under SBC security Parameters.

`CONFIG_MODE` (optional)

Configuration synchronization mode, either `pull` (default — the ABC SBC periodically pulls activated configurations from the Cluster Config Manager) or `push` (the Cluster Config Manager pushes configurations to the ABC SBC on demand). See Web GUI Configuration (Cluster Config Master) for the difference between the two modes.

`NO_BOOTSTRAP_PULL` (optional)

When set to `1` *and* `CONFIG_MODE=push`, the ABC SBC skips the initial configuration pull at startup and instead waits for the Cluster Config Manager to push its first configuration.

Because pushing a configuration to the ABC SBC API requires a valid ABC SBC API token, this option additionally disables token authorization on the ABC SBC API until a configuration containing tokens is activated on the node, avoiding a chicken-and-egg situation on a freshly started node.

In this mode, `CLIENT_CERT` / `CLIENT_KEY` are **required**: with no initial configuration there is no `IMI` interface yet, so the ABC SBC API has no TLS profile to fall back on and reuses `CLIENT_CERT` / `CLIENT_KEY` as its server certificate. This prevents the ABC SBC API from listening over plain HTTP while the node is still un-configured.

`GROUP` (optional)

Name of the configuration group the node will be assigned to. Defaults to `default`.

`NODE_UUID` (optional)

The node's UUID. If not set, a random value is generated on first start. Setting this variable allows the same node identity to be reused across container re-creations even when `/data` is not persisted. Otherwise, mount a named volume (or bind mount) under `/data` to persist the UUID and other state across container restarts.

`CLIENT_CERT` / `CLIENT_KEY` (optional, conditionally required)

Path (inside the container) to the TLS certificate and private key used by the ABC SBC to authenticate to the Cluster Config Manager, both when pulling configurations and when receiving config pushes. `CLIENT_KEY` defaults to `CLIENT_CERT`, in which case `CLIENT_CERT` must contain both the certificate and the private key.

Required when the Cluster Config Manager has mTLS enabled (see SBC security Parameters), and also required whenever the ABC SBC starts with no initial configuration yet (i.e. `CONFIG_MODE=push` together with `NO_BOOTSTRAP_PULL=1`): in that case the ABC SBC API reuses these files as its own server certificate, so the API can listen over HTTPS even before any `IMI` interface has been configured. In every other case, the certificates carried by the initial configuration are used as soon as the `IMI` interface is applied, and `CLIENT_CERT` / `CLIENT_KEY` may be omitted.

`CA_CERT` (optional)

Path (inside the container) to a CA certificate bundle used by the ABC SBC to verify the Cluster Config Manager's TLS certificate.

`SERVER_IP` (optional)

IP address that the ABC SBC API listens on before the first configuration is received. Defaults to `0.0.0.0` so that the Cluster Config Manager can reach the ABC SBC API to push the initial configuration. See Unified SBC management service.

> **Info**
>
> In earlier releases, ABC SBC containers had to be configured by running the interactive `sbc-init-config` script inside the container. That approach is no longer the preferred method; new deployments should use the environment variables described above. `sbc-init-config` is still available as an alternative and is documented in Initial Configuration.

> **Warning**
>
> If `sbc-init-config` was run at any point, the configuration file it creates (`/data/sbc/etc/sbc-pullconf.conf`) takes precedence over environment variables. When migrating from the old configuration method to environment variables, remove that file from the data volume before (re)creating the container.

3.3.2 Podman

In this section it is described the OCI container installation process under podman for the ABC SBC and the Cluster Config Manager. The procedure is based on podman 5.x used on Debian 13.

Please refer to the official documentation for more detailed information.

Installing podman

Since Debian 13 the podman package is available in official repositories and can be installed on host with Debian OS via:

```
% apt install podman
```

OCI images download

Start by downloading the OCI images directly from Frafos's docker registry (registry.frafos.net) using the following:

```
% podman pull registry.frafos.net/abc/sbc:5.6
% podman pull registry.frafos.net/abc/ccm:5.6
```

Frafos's OCI tagging strategy is the following:

- `5.0.[0,100]`: the tag matches the image exact version (`5.0.1`, `5.0.2`)
- `5.0`: alias to the latest `5.0.X` image for the major release
- `5.0-XX`: alias to an exact `5.0.X` image (`5.0-rc1`, `5.0-dev`)

One may also run the following to pull the desired exact images:

```
% # exact minor release
% podman pull registry.frafos.net/abc/sbc:5.0.42
% # test the release candidate
% podman pull registry.frafos.net/abc/ccm:5.0-rc1
% # test the "latest" 5.0
% podman pull registry.frafos.net/abc/ccm:5.0
```

The access to container registry requires authentication using username and password. The account can be self-created at https://registry.frafos.net/, but the account has to be then assigned to proper project by Frafos. Please contact Frafos support, if your account is not approved yet for access to the ABC SBC repository project.

To log in to the Frafos repository, use the following command. You will be prompted to enter the username and password associated with your customer account:

```
% podman login registry.frafos.net
```

Persistent data

It is highly recommended to use persistent `/data` directory for the container, where all configuration, that needs to be preserved across upgrades, is stored.

This can be reached either by mounting an external directory from the host into the container's `/data` directory:

```
% mkdir -p /var/data/ccm
```

or by using a *podman volume* for Cluster Config Manager and ABC SBC separately:

For Cluster Config Manager installation:

```
% podman volume create ccm-data
```

For ABC SBC installation:

```
% podman volume create sbc-data
```

In the first case, the storage is fully under administrator's control and can be accessed not only from containers, but directly by other processes as well. This can be beneficial for downloading CDRs from ABC SBC, for backups generated from host or for mounting a specific partition.

In the second case, the storage is fully managed by podman and can be accessed only from containers or via podman commands. This may be sometimes limiting and thus in general rather not the recommended way.

Podman Container Installation

There are many ways how to configure networking with podman which may vary use case by use case. For simplification we will focus on the most common scenario only. Cluster Config Manager and Monitor are run on a bridged network, which is the recommended setup. For the ABC SBC, separate MACVLAN networks are used for management and VoIP traffic. This configuration allows multiple containers (Cluster Config Manager and ABC SBC) running on the same host or on different hosts, depending on desired performance and other requirements.

On AWS, Azure or GCP, only bridge and host network types are supported, so MACVLAN is not used there for the ABC SBC either.

Podman installation for AWS - Azure - GCP

After pulling Cluster Config Manager image to the host, the network can be created by the following command:

```
% podman create network mgmt
```

The container can be created by the following command:

```
% podman create --name ccm --hostname ccm --tz=local --tty \
 --mount type=volume,src=ccm-data,target=/data,rw=true \
 --cap-add=AUDIT_CONTROL --cap-add=AUDIT_WRITE --cap-add=NET_RAW \
 --network mgmt:interface_name=eth0 \
 -p 443-444:443-444 \
 registry.frafos.net/abc/ccm:5.6
```

For ABC SBC container creation:

```
% podman create --name sbc --hostname sbc --tz=local --tty \
 --mount type=volume,src=sbc-data,target=/data,rw=true \
 --cap-add=NET_RAW \
 --cap-add=NET_ADMIN \
 --pids-limit 65536 \
 --network host \
 -e MASTER=ccm.example.com \
 -e CONFIG_USER=sbcnode \
 -e CONFIG_PASS=<password> \
 -e CONFIG_MODE=pull \
 registry.frafos.net/abc/sbc:5.6
```

Notes about parameters:

- The `--pids-limit` can be updated according to the threads number that is used by sems. By default it is 2048, however with higher number of threads (used by sems) sbc requires higher limit.
- The `--tz=local` enables container to use the same timezone with host.
- When HA is used, adding `--cap-add=SYS_NICE` to the SBC container could prevent false failovers that could happen when the system is under heavy load.
- The `-e` options tell the ABC SBC how to reach the Cluster Config Manager at startup. `MASTER`, `CONFIG_USER` and `CONFIG_PASS` are required; `CONFIG_MODE` defaults to `pull` and may be set to `push` instead. See Configuring the SBC container for the full list of supported variables (TLS certificates, `NO_BOOTSTRAP_PULL`, `NODE_UUID`, `GROUP`, …).

When using the host network type, ensure SSH access on port 22 is allowed in the initial configuration by adding a rule through the “Manual low-level rules” section in the Cluster Config Manager GUI (System > Firewall). This prevents accidental loss of SSH connectivity for ABC SBC host.

After container creation, a service should be created by the following commands to manage container, and then this will create a service which needs to be started and enabled:

```
% cd /etc/systemd/system/
```

For Cluster Config Manager:

```
% podman generate systemd --new -f -n ccm
% systemctl start container-ccm
% systemctl enable container-ccm
```

For ABC SBC:

```
% podman generate systemd --new -f -n sbc
% systemctl start container-sbc
% systemctl enable container-sbc
```

Please note that each change (container name, adding a network, adding capabilities) must be added in service file (/etc/systemd/system/container-ccm.service).

To check logs, this command can be run:

```
% podman logs -f ccm
% podman logs -f sbc
```

Regular podman installation

Interface configuration like IP assignment, DNS, etc. can be passed to the container by podman.

Deciding interface numbers should be according to the topology. In basic topology Cluster Config Manager should have only a management (mgmt) interface. On the other hand, ABC SBC should have these interfaces by default:

mgmt: will be used for communication between Cluster Config Manager, Monitor and other ABC SBC in case HA is used. For mgmt interface, one IP should be assigned on the host, one IP should be assigned for container.

internal: will be used for SIP/MEDIA communication with internal sip servers. For internal interface, there is no need to assign an IP on the host, only for container.

external: will be used for SIP/MEDIA communication with public sip servers (like providers). For external interface, there is no need to assign an IP on the host, only for container.

If more interfaces are needed for cluster, then external2, external3 etc. can be added as well. Obviously, the names of the interfaces can be updated as customers wish.

For Cluster Config Manager a specific routing is not needed as it supposed to have only one interface. If additional routing rule is necessary, please check the note related to routing part under ABC SBC network configuration. For Cluster Config Manager container creation, only one network is required for the Cluster Config Manager, even if the Manager and Monitor can be installed together on the same host. A bridged network is used here with the driver netavark. The definition is made using a simple command:

```
% podman network create mgmt
```

Display a list of existing Podman networks:

```
% podman network ls
```

An individual network can also be displayed for checking purposes. The exact definition within the container definition is displayed here:

```
% podman network inspect mgmt
```

Then please set proper parameters while creating the Cluster Config Manager container:

```
% podman create --name ccm --hostname ccm --tz=local --tty \
  --mount type=volume,src=ccm-data,target=/data,rw=true \
  --cap-add=NET_RAW --cap-add=AUDIT_CONTROL --cap-add=AUDIT_WRITE \
  --network mgmt:interface_name=eth0 \
  -p 443-444:443-444 \
  --dns=172.22.31.2 \
  registry.frafos.net/abc/ccm:5.6
```

Notes about parameters:

- The `--pids-limit` can be updated according to the threads number that is used by sems. By default it is 2048, however with higher number of threads (used by sems) sbc requires higher limit.
- The `--tz=local` enables container to use the same timezone with host.

- The `NET_RAW` is needed for raw sockets usage and additionally for troubleshooting purposes (for example to allow `ping` utility).
- The `--dns=none` flag tells podman not to create `/etc/resolv.conf` in the container.

Please note, that to make per-interface DNS configuration working in an ABC SBC container, it is necessary to allow resolvconf when performing the SBC initial config. To make it working in Cluster Config Manager, a CCM configuration option "Enable to use resolvconf for dns" needs to be set.

For ABC SBC container network creation:

Several networks are required for the SBC, here no bridged network is used, but macvlan network. As driver netavark is used.

VMWare hints: By default, a VNIC only receives unicast traffic targeted at its own MAC address. As MACVLAN is used for ABC SBC installation, promiscuous mode must be enabled on the networks so that the VNIC can receive traffic for other MAC addresses.

With VMWare ESXi, you should set "Accept" for "Promiscuous mode" option, like in following screenshot:

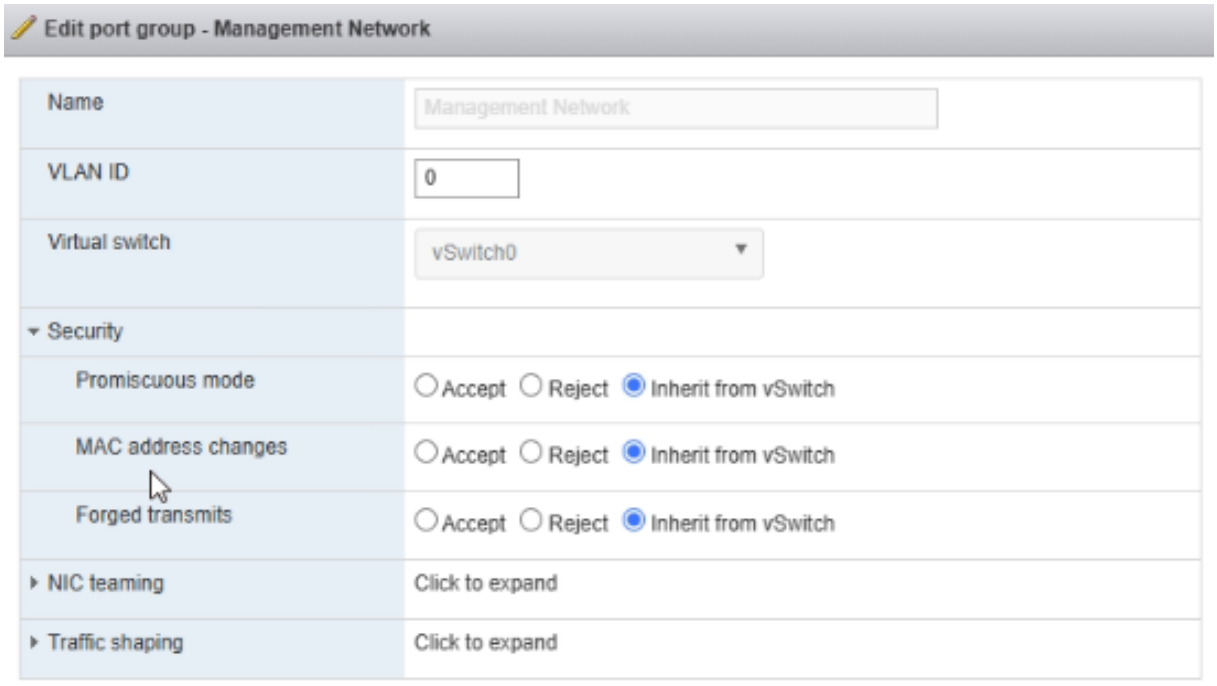

The Podman networks are defined via the command line as examples, please note that customer should decide for each subnet according to the customer network.

The `--internal` must be added to avoid adding default gateway for the network.

Management network mgmt 10.10.10.0/24 via enp0s8

```
% podman network create --driver macvlan --opt parent=enp0s8 --gateway 10.10.10.1 --internal --
subnet 10.10.10.0/24 mgmt
```

Voice network internal 10.10.20.0/24 via enp0s9:

```
% podman network create --driver macvlan --opt parent=enp0s9 --gateway 10.10.20.253 --internal --
subnet 10.10.20.0/24 internal
```

Voice network (with default gateway) External 10.10.30.0/24 via enp0s10:

```
% podman network create --driver macvlan --opt parent=enp0s10 --gateway 10.10.30.253 --subnet
10.10.30.0/24 external
```

A note related to routing:

The `--route` option can be used during network creation: A static route in the format <destination in CIDR notation>,<gateway>,<route metric (optional)>. This route will be added to every container in this network. It can be specified multiple times if more than one static route is desired.

A sample:

```
% podman network create --driver macvlan --opt parent=enp0s9 --gateway 10.10.30.253 --subnet
10.10.30.0/24 --route 10.40.20.0/24,10.10.30.1 external
```

Then please set proper parameters while creating the container:

```
% podman create --name sbc --hostname sbc --tz=local --tty \
  --mount type=volume,src=sbc-data,target=/data,rw=/true \
  --cap-add=NET_ADMIN --cap-add=NET_RAW \
  --pids-limit 65536 \
  --network mgmt:interface_name=eth0,ip=10.10.10.100 \
  --network internal:interface_name=eth1,ip=10.10.20.101 \
  --network external:interface_name=eth2,ip=10.10.30.102 \
  --dns=172.22.31.2 \
  -e MASTER=ccm.example.com \
  -e CONFIG_USER=sbcnode \
  -e CONFIG_PASS=<password> \
  -e CONFIG_MODE=pull \
  registry.frafos.net/abc/sbc:5.6
```

Notes about parameters:

- The `--pids-limit` can be updated according to the threads number that is used by sems. By default it is 2048, however with higher number of threads (used by sems) sbc requires higher limit.
- The `--tz=local` enables container to use the same timezone with host.
- The `NET_RAW` is needed for raw sockets usage and additionally for troubleshooting purposes (for example to allow `ping` utility).
- The `--dns=none` flag tells podman not to create `/etc/resolv.conf` in the container.
- The `-e` options tell the ABC SBC how to reach the Cluster Config Manager at startup. `MASTER`, `CONFIG_USER` and `CONFIG_PASS` are required; `CONFIG_MODE` defaults to `pull` and may be set to `push` instead. See Configuring the SBC container for the full list of supported variables (TLS certificates, `NO_BOOTSTRAP_PULL`, `NODE_UUID`, `GROUP`, …).

After container creation, a service should be created by the following commands to manage container, and then this will create a service which needs to be started and enabled:

```
% cd /etc/systemd/system/
```

For Cluster Config Manager:

```
% podman generate systemd --new -f -n ccm
% systemctl start container-ccm
% systemctl enable container-ccm
```

For ABC SBC:

```
% podman generate systemd --new -f -n sbc
% systemctl start container-sbc
% systemctl enable container-sbc
```

Please note that each change (container name, container image, adding capabilities) must be updated in service file (/etc/systemd/system/container-ccm.service).

To check logs, this command can be run:

```
% podman logs -f ccm
% podman logs -f sbc
```

Container management

To list running containers use:

```
% podman ps
```

To list even stopped containers:

```
% podman ps -a
```

To list volumes:

```
% podman volume ls
```

To list all images:

```
% podman image ls
```

To start, stop or restart container:

```
% systemctl start container-ccm
% systemctl stop container-ccm
% systemctl restart container-ccm
```

Executing commands within the container

To open a shell inside the container, use the following command on the host server:

```
% podman exec -it sbc bash
```

Another commands can be executed directly similar way:

```
% podman exec -it sbc ip route
```

Checking journal

As ABC SBC does not contain journal anymore, container's systemd journal or podman logs can be checked:

```
% podman logs -f ccm
```

or:

```
% journalctl -u container-ccm
```

Upgrade Procedure

For a container upgrade, please start by pulling the newest images:

```
% podman pull registry.frafos.net/abc/sbc:5.6
% podman pull registry.frafos.net/abc/ccm:5.6
```

You may check the latest images metadata using:

```
% podman image ls
REPOSITORY                      TAG     IMAGE ID      CREATED       SIZE
registry.frafos.net/abc/ccm     5.6     bc3c1b61316a  2 hours ago   1.03 GB
registry.frafos.net/abc/sbc     5.6     574b44e60f30  5 hours ago   1.25 GB
<none>                          <none>  655cf4d30cbf  24 hours ago  1.25 GB
<none>                          <none>  8e0eb0df41c0  24 hours ago  1.03 GB
```

To re-create and re-start the container (ABC SBC in this case), using the newest image, update the version part (end of the "ExecStart" line) in container service:

```
% vi /etc/systemd/system/container-sbc.service
```

and then reload the daemon and restart the container service:

```
% systemctl daemon-reload
% systemctl restart container-sbc
```

One may then remove the unused images (old, untagged variant), using:

```
% podman image prune
WARNING! This will remove all dangling images.
Are you sure you want to continue? [y/N] y
8e0eb0df41c0c9c8cb6c8154b5fc7f4979869ede92febc7f220b4b6a08ebf133
655cf4d30cbf018198f096bf059283eda27c9f9b0073f92ae748af74316e6be4
```

3.4 Post-installation steps

> **Info**
>
> IMPORTANT: AFTER THE INSTALLATION PROCESS IS COMPLETE AND BEFORE CONFIGURATION AND TESTING BEGINS WE URGE YOU TO WHITELIST THE IP ADDRESS FROM WHICH THE ABC SBC WILL BE ADMINISTERED.

Failure to whitelist the administrator's IP address may – especially during the initial configuration and testing – easily block the administrative access to the machine. Various automated blacklisting techniques can block the whole IP address if they spot unexpected traffic from the IP address. See more details in Section Automatic IP Address Blocking.

To whitelist the IP address, visit the administrative GUI under **"Config ▸ Firewall ▸ Exceptions to automatic Blacklists ▸ Add"** as shown in the Figure bellow:

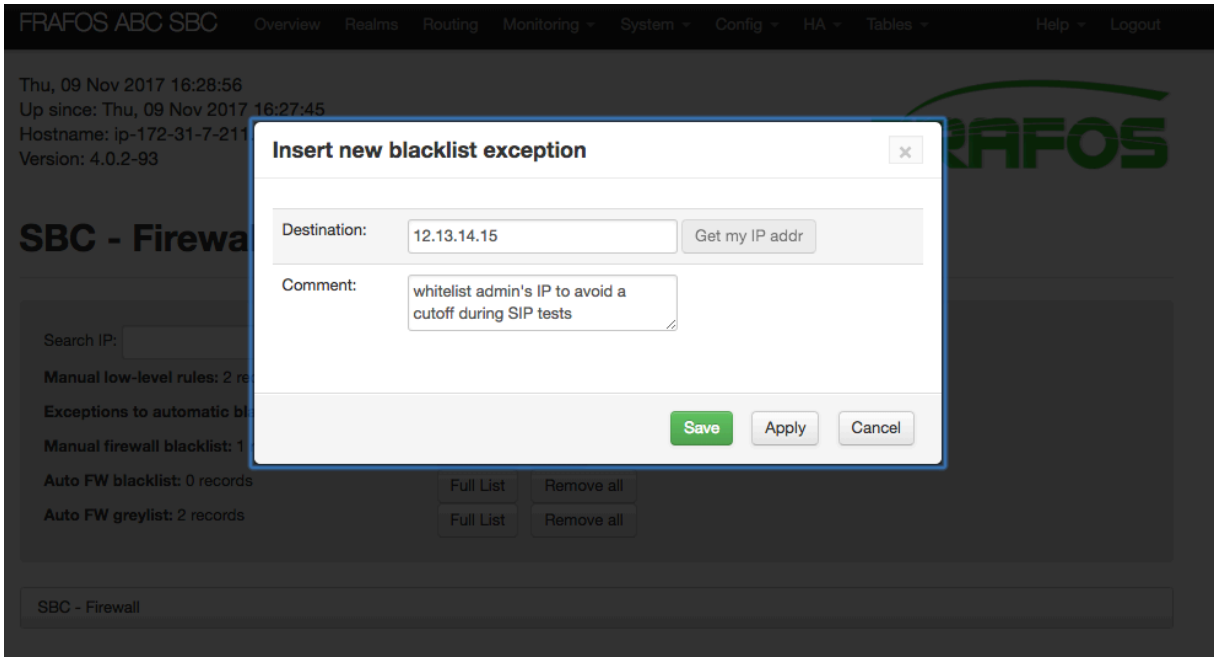

Figure 1: Warning: Whitelist Administrator's IP Address

Chapter 4

User guide

4.1 About the ABC Session Border Controller

This manual is a complete handbook for the ABC Session Border Controller (ABC SBC). It documents network planning, SBC installation, policy configuration and the best current practices for operating the SBC.

The ABC Session Border Controller (ABC SBC) is a SIP Back-2-Back User Agent (B2BUA) that provides operators and enterprises with a scalable session border control solution for secure connections with Voice over IP (VoIP) operators and users. With the ABC SBC VoIP service providers and enterprises deploy a session border controller that is designed to run on top of high end hardware as well as appliances and virtual machines. Thereby, the ABC SBC enables VoIP providers to gradually scale up their infrastructure and covers the needs of enterprises of all sizes.

The ABC SBC provides the following features:

- Infrastructure Security: The ABC SBC serves as the first line of defence, fending off attacks coming over the Internet, hiding internal topology, applying rate limits and performing Call Admission Control, limiting number of parallel calls and call length, and off-loading registrar and registration throttling.
- Confidentiality. The ABC SBC implements cryptographic protocols TLS and SRTP that make it incredibly hard for unauthorized third parties to intercept VoIP calls. Secured telephony is possible even without exotic telephones using off-the-shelf webRTC browsers (See the next point). The ABC SBC can also combine cryptographically secured RTC telephony with traditional policy-based IT practices like VPNs, so that confidentiality can be achieved in a practicable end-to-end way between all kinds of equipment.
- Browser Telephony. The ABC SBC includes a built-in SIP/WebRTC gateway. The gateway allows users to interconnect WebRTC browser telephony with SIP telephony and even PSTN telephony users behind SIP/telephony gateways. The browser telephony allows for easy integration with web applications and provides a level of privacy previously unprecedented before in fixed and mobile networks.
- Network Functions Virtualization (NFV). The ABC SBC also comes in a virtualized form that allows administrators to run the SBC without managing the physical infrastructure. More than that, a whole auto-scaling load-balanced RTC gateway cluster can be started in Amazon Elastic Cloud by a single button using the Cloud Formation launching facility. Such a cluster adapts to network conditions, growing and shrinking with network traffic. It can be geographically dispersed for best QoS worldwide and it can be launched in less than five minutes – compare that to the effort of placing your own equipment in multiple geographically distributed air-conditioned data-centers!
- Mediation: The ABC SBC connects disconnected unroutable networks and VLANs, different transport protocols, secured and plain RTP, facilitates NAT traversal, steers codec negotiation, translates identities and adapts SIP headers and bodies for best interoperability between incompatible devices and networks policies.
- Rapid IT integration. The ABC SBC dramatically reduces the time-to-deploy. Studies show that in the vast majority of new network deployments inadequate time and cost is spent in designing data integration concepts. ABC SBC reduces the time-to-deploy with its built-in integration capabilities. Administrators can place external logic to web-servers and govern how the SBC behaves through a RESTful interface. Large amounts of pre-provisioned data can be used to govern the SBC logic, such

as routing tables, peering characteristics, SIP bulk registration, blacklists or whitelists, or subscriber information.
- SIP Routing: The ABC SBC's competitive design allows administrators to route SIP traffic based on any message element. Routing methods like source-based, destination-based, least-cost-route based, proprietary-header-field-based and others can be easily configured and cascaded behind each other to find the most-proper destination for SIP traffic.
- Real-time monitoring. The ABC SBC allows its administrator to permanently know what is going on in their SIP networks. Due to the centralized nature of SBCs, the ABC SBC enables you to gain deep insight into the traffic it steers and constantly reports on it using "events" and "Call Detail Records" (CDRs). This data can be further used to perform troubleshooting, backwards analysis and future predictions of the system as whole as well as that of its individual users. The ABC SBC also reports on its status using Prometheus and SNMP (legacy).
- Media processing: The ABC SBC includes built-in audio recording, transcoding, announcements and conferencing.
- Web management. Remote management allows rapid and convenient adaptation to ever changing network conditions. ABC SBC's policies can be easily changed through the web interface.
- Non-stop service. The ABC SBC is designed to provide high-availability by running in redundant hot-standby pairs. Alternate route definitions and built-in monitoring conceal scheduled and unplanned outages of network elements behind the ABC SBC.

4.1.1 How to Start?

This book is intended for everyone interested in installing and using the ABC SBC. Knowledge of SIP, RTP and IP networking is of an advantage and would ease the reading and use of the book. Of essential value is, however, a good understanding of the VoIP environment in which the ABC SBC is to be deployed. Depending on your goal, there are these options for how to get the most out of this book in the shortest time:
- Cloud RTC Trial: Trialing the RTC gateway using amazon Elastic Cloud allows you to start the WebRTC/SIP gateway service within minutes and establish connectivity between web browsers and an existing SIP service. See the Section Amazon Elastic Cloud Configuration Cookbook and visit our trial site at https://go.frafos.com/.
- Installing the ABC SBC: Before installing the ABC SBC it is advisable to go through the practical guide to have a better understanding of the needed infrastructure. After installing the ABC SBC, the practical guide can be used for a quick configuration of the solution. In case certain issues need to be solved that are not covered in the guide, then a look into the reference chapter (Section Reference of Actions) will be helpful. The administrator should also go through the administration, monitoring and security chapters to develop a better understanding and control of the installed system.

The book is structured in the following parts:
- Introduction: This section provides an overview of the basic technologies addressed here, namely SIP and SBC. Furthermore, the basic concepts and terminology of the ABC SBC are described. If you are knowledgeable with SIP and VoIP deployments you can skip the introductions to SIP and SBCs.
- Practical Guide to the ABC SBC: This section provides first an overview of what a future user of the ABC SBC - or actually any other SBC as well - must consider before purchasing and installing an SBC. Moreover, this guide can be seen as a short cut for configuring and using the main features of the ABC SBC without having to go through the entire manual.
- Installation guide: This section covers the steps needed to deploy the ABC SBC. Firstly, one needs to determine whether to do a complete installation from the FRAFOS repository or to use ABC SBC container version.
- General ABC Configuration Guide: This section provides the details about the different features of the ABC SBC, what they are good for and how they can be used. For the more complex parts, additional

sections with examples are provided. These example sections are intended as short cuts for solving common issues.

- ABC SBC System administration: This chapter explains a set of features available for the administrator of the ABC SBC. These features include the capability to create and manage users of the ABC SBC and define their rights, the list of commands that can be used to run certain tasks that might not be available through the GUI as well as conduct upgrades and updates of the ABC SBC software.
- Monitoring and Troubleshooting: The ABC SBC collects various measurement values and call traces and generates alarms. These features provide the administrator of the ABC SBC with the information needed to detect errors and problems in the processed VoIP traffic as well as in the operation of the ABC SBC.
- Securing SIP Networks using ABC SBC and ABC Monitor (optional): This chapter provides an overview of the security capabilities of the ABC SBC as well as a guide for configuring blacklists, traffic shaping and limiting the call duration.

4.1.2 Credits

The initial version of this book was written by the FRAFOS team with support from Sipwise in a three day Book Sprint facilitated by Barbara Rühling. The illustrations were provided by Juan Camilo Cruz. This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License .

4.2 Introduction

4.2.1 A Brief Introduction to History and Architecture of SIP

The Session Initiation Protocol (SIP, **RFC 3261**) is a backbone of every VoIP network nowadays. Its "language" is used by telephony devices to find each other, signal who is calling whom, negotiate which audio/video codecs to use and even more. The telephony devices are typically SIP desktop phones, but it may be also smartphones, softphones, or massive PSTN gateways with PSTN infrastructure and users behind them. In between, there are intermediary SIP network devices that help to locate the end-user devices, perform Call Admission Control, help often with various imperfections of the end-devices and perform other useful functions. The ABC Session Border Controller is one of such, however other kinds of SIP proxy servers, Back-to-Back User Agents, specialized application servers, and more are common.

VoIP began to reach market back in mid nineties. By then Internet had established itself as a consumer product. The number of users buying PCs and subscribing to an Internet Service Provider (ISP) for a dial-up access was increasing exponentially. While mostly used for the exchange of Email, text chatting and distribution of information VoIP services based on proprietary solutions as well as H.323 started to gain some popularity. The standardization organization Internet Engineering Task Force, IETF, began to devise its own protocol suite. Some protocols existed already by then. The Real-Time Transport Protocol (RTP) **RFC 1889** enabled the exchange of audio and video data. The Session Description Protocol (SDP) **RFC 2327** enabled the negotiation and description of multimedia data to be used in a communication session. The first applications, often open-source, for sending and receiving real-time audio and video data emerged. A signaling protocol was missing, however.

In those days, the procedure for establishing a VoIP call between two users based on the IETF standards would look as follows: The caller starts his audio and video applications at a certain IP address and port number. The caller then either calls the callee over the phone or sends him an Email to inform him about the IP and port address as well as the audio and video compression types. The callee then starts his own audio and video applications and informs the caller about his IP and port number. While this approach was acceptable for a couple of researches wanting to talk over a long distance or for demonstrating some research on QoS this was clearly not acceptable for the average Internet user.

The Session Initiation Protocol (SIP) **RFC 3261** was the attempt of the IETF community to provide a signaling protocol that will not only enable phone calls but can also be used for initiating any kind of communication sessions. SIP has been contemplated for use by audio and video calls, as well as for setting up a gaming session or controlling a coffee machine.

The SIP specifications describe three types of components: user agents (UA), proxies and registrar servers. The UA can be the VoIP application used by the user, e.g., the VoIP phone or software application. A VoIP gateway, which enables VoIP users to communicate with users in the public switched network (PSTN) or an application server, e.g., multi-party conferencing server or a voicemail server are also implemented as user agents. The registrar server maintains a location database that binds the users' VoIP addresses to their current IP addresses. The proxy provides the routing logic of the VoIP service. When a proxy receives a SIP request from a user agent or another proxy it also conducts service specific logic, such as checking the user's profile and whether the user is allowed to use the requested services. The proxy then either forwards the request to another proxy or to another user agent or rejects the request by sending a negative response.

While the server roles prescribed by the SIP specification are functional, actual implementations found in networks tend to integrate multiple roles in a server product. A registrar is often co-located with a proxy server so that they can share user-location databases. A server can also present itself as User-Agent to both sides of a signaling session to be able to manipulate SIP messages more extensively than the proxy specification would permit.

Every signaling SIP transaction consists of a request and one or more replies. The three most commonly used request types are REGISTER, INVITE and BYE. The REGISTER request makes a SIP phone's address known to a SIP server so that it knows where to forward incoming SIP requests. The INVITE request initiates a dialog between two users. A BYE request terminates this dialog. Responses can either be final or provisional. Final responses can indicate that a request was successfully received and processed by the destination. Alternatively, a final response can indicate that the request could not be processed by the destination or by some proxy in between or that the session could not be established for some reason. Provisional responses indicate that the session establishment is in progress, e.g. the destination phone is ringing.

In general one can distinguish between three types of SIP message exchanges, namely registrations, dialogs and out of dialog transactions.

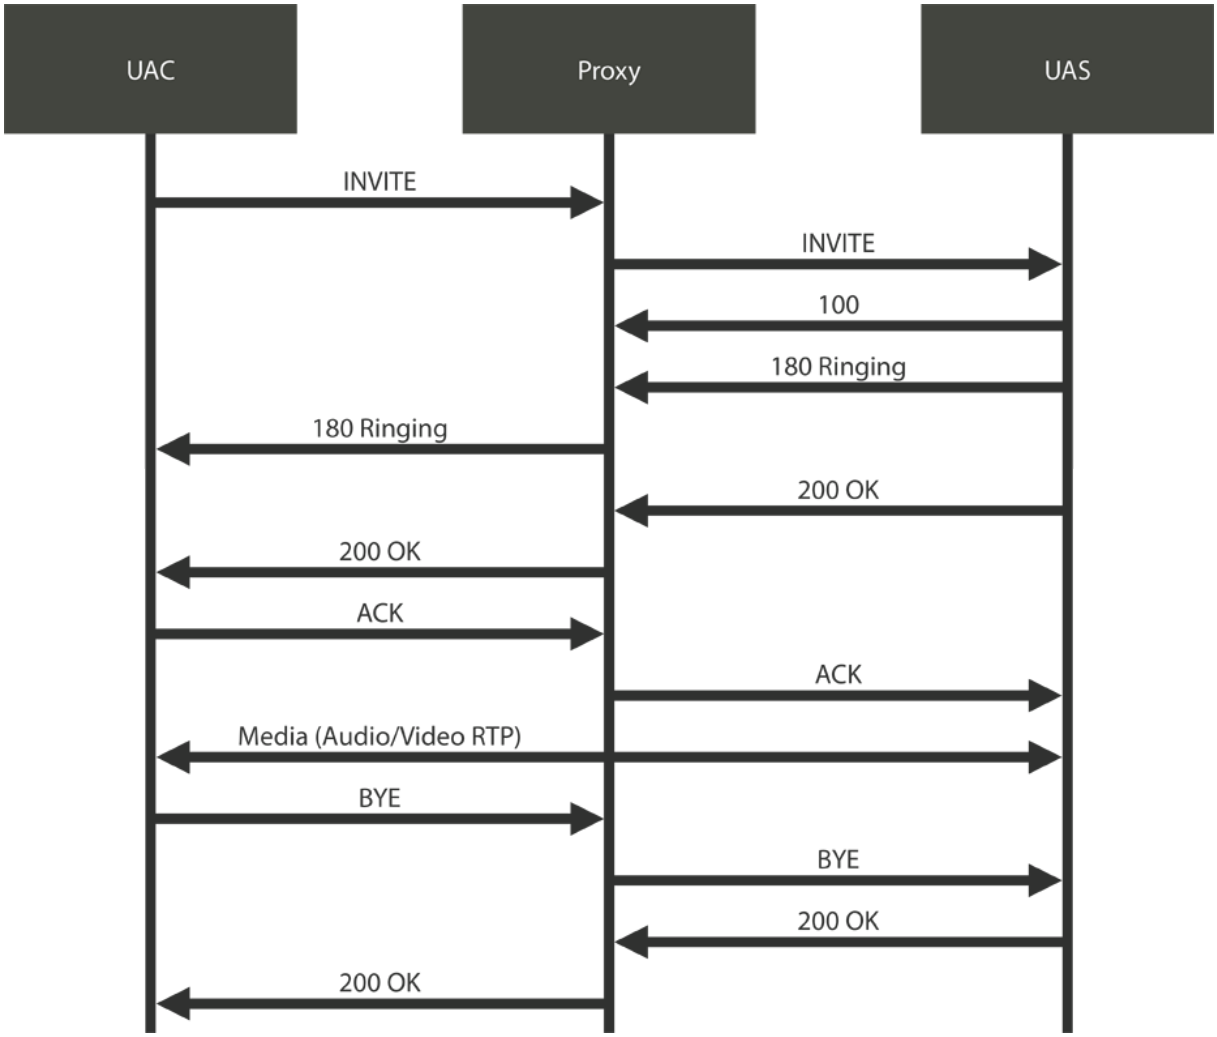

Figure 2: SIP Call flow

A SIP registration enables a user agent to register its current address, IP address for example, at the registrar. This enables the registrar to establish a correlation between the user agent's permanent address, e.g. sip:user@frafos.com, and the user agent's current address, e.g., the IP address used by the user's user agent. In order to keep this correlation up to date the user agent will have to repeatedly refresh the registration. The registrar will delete a registration that is not refreshed for a while.

A SIP dialog, a call for example, usually consists of a session initiation phase in which the caller generates an INVITE that is responded to with provisional and final responses. The session initiation phase is terminated with an ACK, see SIP Call flow. A dialog is terminated with a BYE transaction. Depending on the call scenario the caller and callee might exchange a number of in-dialog requests such as reINVITEs or REFER.

The last type of SIP interactions is SIP transactions that are not generated as part of a dialog. Examples of out of dialog SIP requests include OPTIONS and INFO that are often used for exchanging information between SIP nodes or as an application level heartbeat.

Every SIP message consists of three parts: First line, message header and message body, see Content of SIP messages. The first line states the purpose of the message. For requests it identifies its type and the destination address. For replies the first line states the result as a numerical 3-digit status code together with a textual human-readable form. The second part of the message, the header part, includes a variety of useful information such as identification of the User Agent Client and the SIP path taken by the

request. The third part includes a message body that contains application specific information. This can be for example session description information (SDP) indicating the supported codecs.

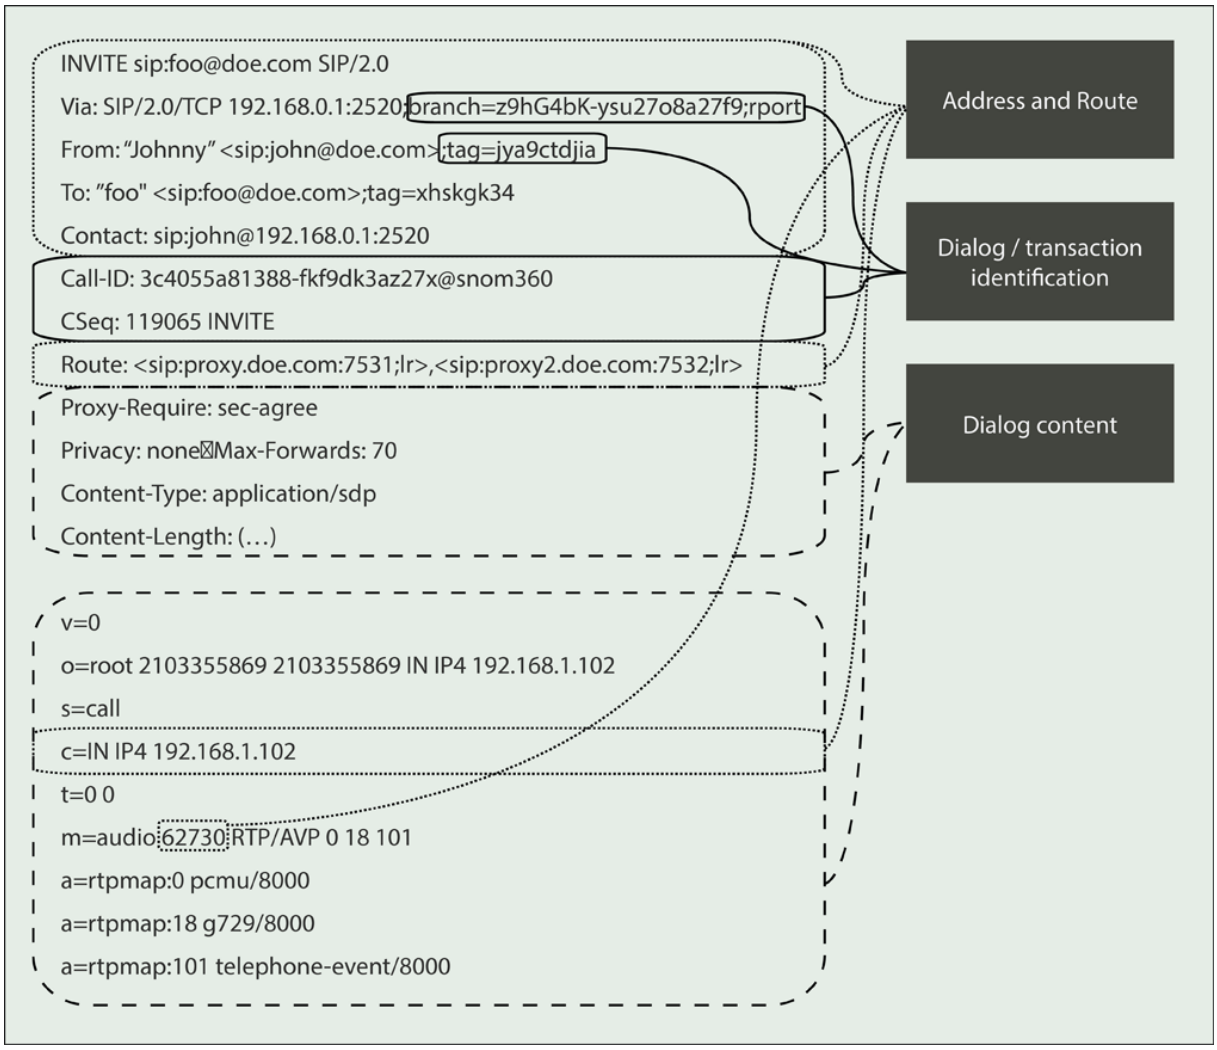

Figure 3: Content of SIP messages

The information contained in these three parts can be roughly divided into three categories, see Content of SIP messages:

- Addressing and routing information: This includes information about who has sent the message and where it is destined to, the next hop to be sent to as well as the hops it has traversed. This information is included in the first line as well as in different headers such as From, To, Contact, P-Asserted-Identity, Via, Route, Path and other headers. The message body can contain information about where the media traffic should be sent to or is expected to come from.
- Dialog and transaction identification: This part of a SIP message is used to uniquely identify a SIP dialog or transaction. This information is included in SIP headers such as Cseq, Call-Id and tags included in From, To and Via headers.
- Dialog content: With dialog content we categorize data that is included in a SIP message that is either used to describe certain features of a dialog or indicates how a node receiving the message should process the message. This can include parts of the SIP message body carrying SDP, which includes description about which audio or video codes to use. Certain headers such as Privacy for example indicate the user's wishes with regard to the way private information such as user address should be dealt with.

4.2.2 What is a Session Border Controller (SBC)?

Historically Session Border Controllers emerged after publication of the SIP standard as a panacea to early protocol design mistakes: ignorance of Network Address Translators (NATs), unclear data model, liberal syntax, reluctance to standardize legal interception and more.

Probably the single biggest mistake in the design of SIP was ignoring the existence of network address translators (NAT). This error came from a belief in the IETF leadership that IP address space would be exhausted more rapidly and would necessitate global upgrade to IPv6 which would eliminate the need for NATs. The SIP standard has assumed that NATs do not exist, an assumption, which turned out to be a failure. SIP simply didn't work for the majority of Internet users who are behind NATs. At the same time it became apparent that the standardization life-cycle is slower than how the market ticks: SBCs were born, and began to fix what the standards failed to do: NAT traversal.

Yet another source of mistakes has been the lack of a clear data model behind the protocol design. Numerous abstract notions, such as dialog or session, transaction or contact simply didn't have unique unambiguous identifiers associated with them. They were calculated or almost guessed out of various combinations of header-fields, decreasing the interoperability. Some message elements, such as *Call-ID*, have been overloaded with multiple meanings. While some of these were fixed in the later SIP revision and its extensions (*rport* **RFC 3581**, *branch*, gruu **RFC 5628**, session-id) the market forces jumped in quickly. SBCs began to implement "protocol repair".

The other class of mistakes emerged from implementations. Many SIP components were built under a simplifying assumption that security comes for free. Numerous implementations were found to be vulnerable to malformed SIP messages or excessive load. The SBCs began to play a security role.

The reality in today's real time communication networks is that, contrary to the end-to-end design of the Internet and its protocols, service operators can achieve the best user experience by exerting tight control - over the endpoints and over the interface to peering networks.

Over several years, Session Border Controllers became a de facto standard for which ironically no normative reference existed. A non-normative information reference on the subject, **RFC 5853** was published as late as in 2013. Session Border Controllers nowadays handle NATs, fix oddities in SIP interoperability and filter out illegitimate traffic. They began to incorporate elements of the standardized SIP components. For example, routing functionality contemplated by the standards for proxy servers is nowadays part of SBC products. Similarly the SBCs often incorporate media recording and processing functions, whether that's for quality assurance, archiving or legal-compliance purposes.

General Behavior of SBCs

Purist SIP call flow depicts the message flow of an INVITE request between a caller and a callee. This is the simplest message sequence that one would encounter with only one proxy between the user agents. The proxy's task is to identify the callee's location and forward the request to it. It also adds a *Via* header with its own address to indicate the path that the response should traverse. The proxy does not change any dialog identification information present in the message such as the tag in the *From* header, the *Call-Id* or the *CSeq*. Proxies also do not alter any information in the SIP message bodies. Note that during the session initiation phase the user agents exchange SIP messages with the SDP bodies that include addresses at which the agents expect the media traffic. After successfully finishing the session initiation phase the user agents can exchange the media traffic directly between each other without the involvement of the proxy.

SBCs come in all kinds of shapes and forms and are used by operators and enterprises to achieve different goals. Actually even the same SBC implementation might act differently depending on its configuration and the use case. Hence, it is not easily possible to describe an exact SBC behavior that would apply to all SBC implementations. However, in general one we can still identify certain features that are common for most of SBCs. For example, most SBCs are implemented as "Back-to-Back User Agent" (B2BUA).

A B2BUA is a proxy-like server that splits a SIP transaction in two pieces: on the side facing the User Agent Client, it acts as server; on the side facing the User Agent Server it acts as a client. While a proxy usually keeps only state information related to active transactions, B2BUAs keep state information about active dialogs, e.g., calls. That is, once a proxy receives a SIP request it will save some state information. Once the transaction is over, e.g., after receiving a response, the state information will soon after be deleted. A B2BUA will maintain state information for active calls and only delete this information once the call is terminated.

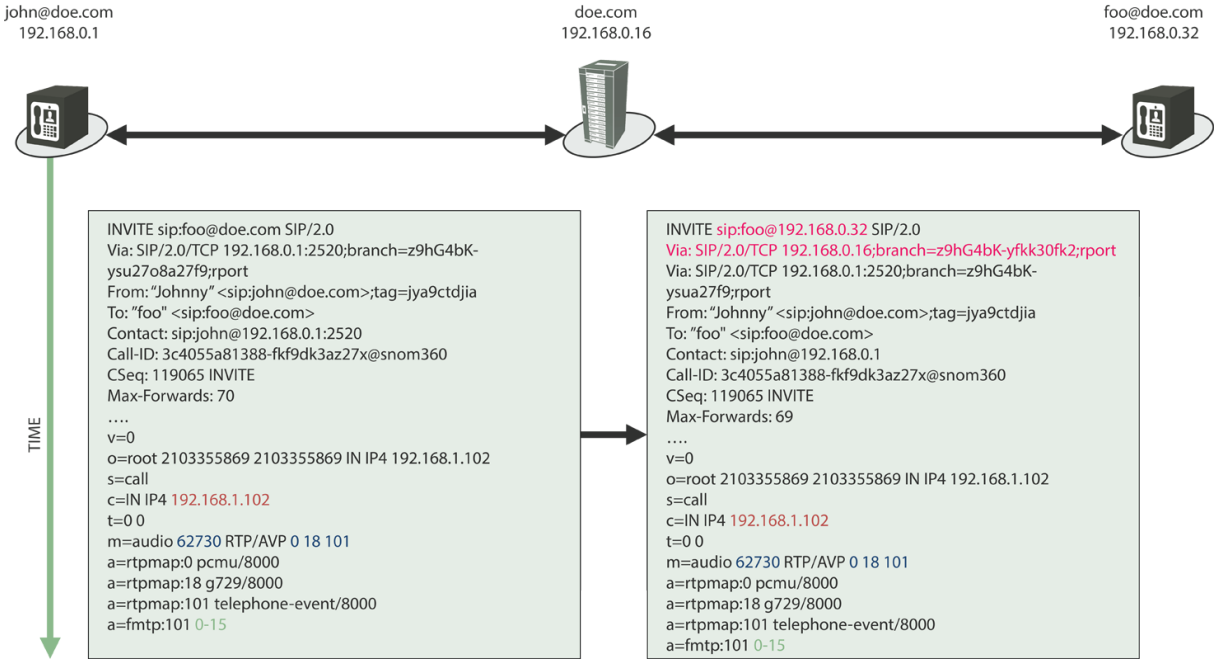

Figure 4: Purist SIP call flow

SIP call flow with SBC depicts the same call flow as in Purist SIP call flow but with an SBC in between the caller and the proxy. The SBC acts as a B2BUA that behaves as a user agent server towards the caller and as user agent client towards the callee. In this sense, the SBC actually terminates that call that was generated by the caller and starts a new call towards the callee. The INVITE message sent by the SBC contains no longer a clear reference to the caller. The INVITE sent by the SBC to the proxy includes *Via* and *Contact* headers that point to the SBC itself and not the caller. SBCs often also manipulate the dialog identification information listed in the *Call-Id* and *From* tag. Further, in case the SBC is configured to also control the media traffic then the SBC also changes the media addressing information included in the *c* and *m* lines of the SDP body. Thereby, not only will all SIP messages traverse the SBC but also all audio and video packets. As the INVITE sent by the SBC establishes a new dialog, the SBC also manipulates the message sequence number (*CSeq*) as well the *Max-Forwards* value.

Note that the list of header manipulations listed in SIP call flow with SBC is only a subset of the possible changes that an SBC might introduce to a SIP message. Furthermore, some SBCs might not do all of the listed manipulations. If the SBC is not expected to control the media traffic then there might be no need to change anything in the SDP lines. Some SBCs do not change the dialog identification information and others might even not change the addressing information.

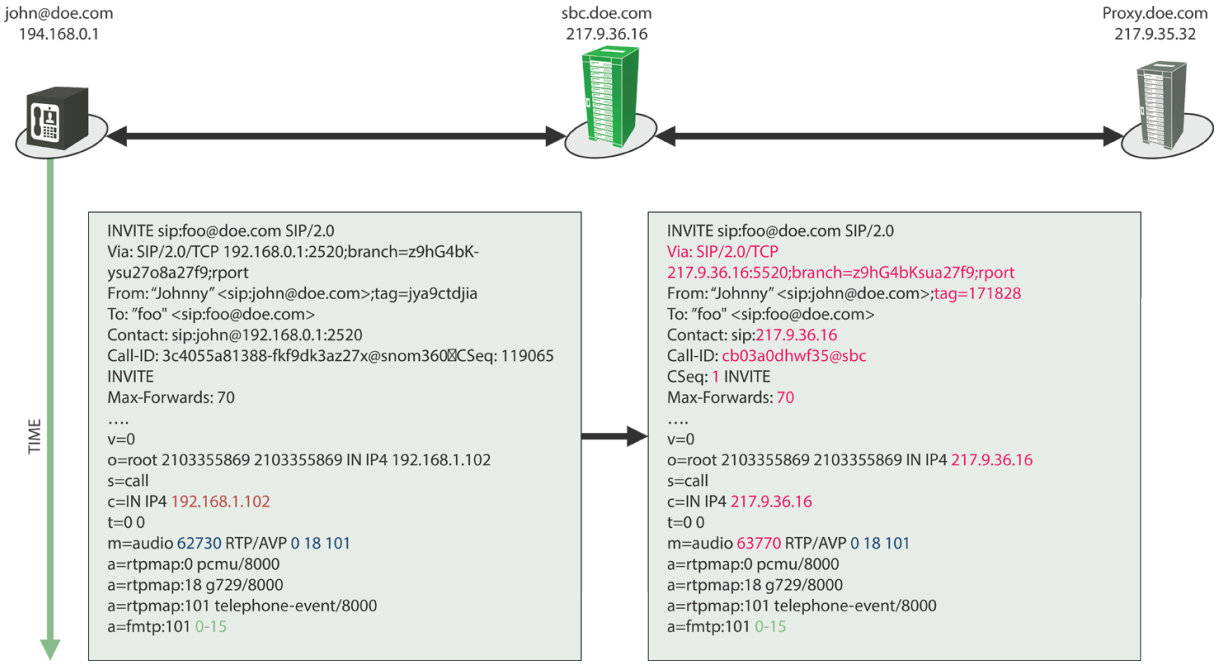

Figure 5: SIP call flow with SBC

General Deployment Scenarios of SBCs

Session border controllers are usually deployed in a similar manner to firewalls, namely with the goal of establishing a clear separation between two VoIP networks.

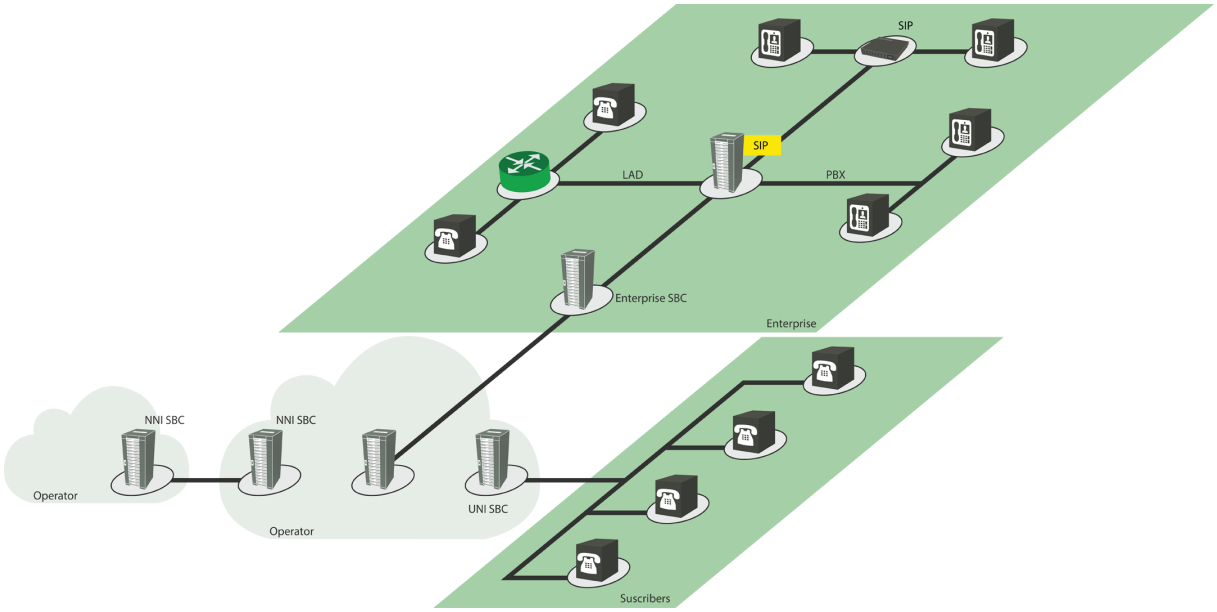

Figure 6: SBC deployment scenarios

In general one can distinguish three deployment scenarios, see SBC deployment scenarios:

- User-Network-Interface (UNI): Operators use SBCs to establish a secure border between their core VoIP components and subscribers. The core components consists of PSTN gateways, media servers, SIP proxy and application servers. Subscribers use SIP hardphones and softphones, Internet Access Devices that connect analog and digital phones to the IP network, and newly web browsers deploying the WebRTC standard. Most important administrative tasks in this scenario include facilitation of NAT traversal (see NAT Traversal), achieving interoperability among multiple types of clients (see SIP Mediation), security against attacks coming from the public Internet (see Securing SIP Networks

using ABC SBC and ABC Monitor (optional)) and off-loading registrar (see Section Registration Caching and Handling).
- Network-Network-Interface (NNI): In the NNI interface, two operators connect to each other directly over SIP. Most important administrative concerns in this scenario include mediation of different network policies (see SIP Mediation), enforcement of service-level-agreements between providers by traffic shaping (see Traffic Limiting and Shaping) and multi-provider SIP routing (see SIP Routing).
- Enterprise SBC (E-SBC): Enterprises are increasingly replacing their PBXs with VoIP PBX or are extending their PBX with a VoIP module to benefit from attractive VoIP minute prices. Enterprise SBCs are used to secure the access to the PBX. The enterprise SBC is also expected to secure the communication to the VoIP operator, which is offering the VoIP service to the enterprise. Typical administrative concerns include harmonization of dialing plans between an enterprise and its trunking partners using the mediation feature see SIP Mediation), and setting up secured VoIP connectivity for web-users (see Securing SIP Networks using ABC SBC and ABC Monitor (optional)).

4.2.3 Do You Need an SBC?

Before installing an SBC it might be worth thinking whether an SBC is needed in the first place. To answer this, here are a couple of questions:

- **will you deal with SIP devices behind a NAT?**

  If the answer is yes then deploying an SBC is most likely the right choice. While there are already a number of NAT traversal solutions such as STUN **RFC 5389**, TURN **RFC 5766** or ICE **RFC 5245** these solutions either do not solve all issues or require certain extensions at the end devices which are not always available.
- **do you deploy SIP devices that you would not want other users or operators to be able to send SIP or media traffic directly to?**

  This is usually the case when a PBX or a PSTN gateway is deployed. If the answer is yes then an SBC would be the right choice as it would hide the IP addresses of these devices and prevent direct communication to them.
- **do you deploy a heterogeneous set of VoIP devices?**

  If yes then an SBC can be the proper point in the network to fix interoperability issues by normalizing the traffic and solving issues created by protocol implementation peculiarities.
- **do you want to protect your VoIP devices from Denial of Service attacks?**

  If there is the danger that an attacker might overload your network and VoIP devices by generating a large amount of SIP requests and RTP packets then an SBC would act as a first line of defense and filter the malicious traffic before it reaches the core VoIP components.
- **do you want to reduce the possibilities of fraud?**

  If there is a danger that a fraudulent user might try to make more calls than allowed then an SBC would be the best approach. With an SBC it is possible to reliably limit the number of calls made by a user.
- **do you want to protect your users from a bill shock?**

  When a user calls an expensive number and fails to terminate the call in a proper manner then he will most likely get a shock when receiving the bill for a call lasting for hours. An SBC on the border of the network can be configured so as to cut calls after a certain period of time and hence limit the damage.
- **will you need to transcode the media?**

  If different users are using different codecs - which is especially the case when connecting mobile to fixed networks - then media transcoding will be needed. Media transcoding is often an integral part of SBCs.
- will your users be using browser telephony using WebRTC? Then you need to connect them to the rest of SIP world and SIP-PSTN gateways using the built-in WebRTC gateway.

4.2.4 ABC SBC Networking Concepts

This section provides an overview of the main concepts and terms of the ABC SBC. It shows the overall model of SBC-managed networks, how the SBC connects to the individual networks using "Interfaces", models SIP devices as "Call Agents", and groups these in "Realms". Eventually A-B-C rules are described that define how the ABC SBC manages SIP traffic as it passes through it.

Network Topology

As depicted in the Figure ABC SBC Concepts Overview, the ABC SBC communicates with VoIP phones, media servers and other entities that act as SIP user agent. We call these entities Call Agents (CA) and group them into so called Realms. The ABC SBC associates rules with Call Agents and Realms. These rules fully describe how every single session traversing the ABC SBC from one CA/Realm to another is processed.

The rule processing occurs in three steps. When receiving a SIP message from a Call Agent, the ABC SBC will first execute inbound rules ("A-Rules") associated with the Call Agent and the Realm it belongs to. These rules typically implement all kinds of admission control. Once the message is accepted the ABC SBC applies routing rules ("B-Rules") to determine the Call Agent where to send the message to. Before actually forwarding the message to the destination, the ABC SBC executes its outbound "C-rules". The C-Rules are typically used to transform the SIP messages to conform to practices used by the destination, such as local specific dialing conventions.

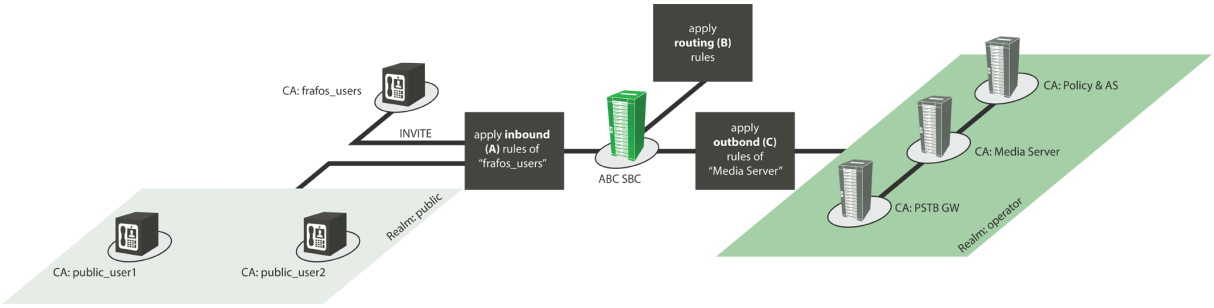

Figure 7: ABC SBC Concepts Overview

SBC Interfaces

SBC Interfaces define how the ABC SBC connects to the adjacent IP networks. They are are an abstraction layer on top of the network interfaces. Specifically, the SBC Interfaces define through which IP addresses, port numbers and network interfaces the ABC SBC offers its services.

There are the following types of SBC interfaces:

- Internal management (IMI): used in high availability (HA) and cluster setups as the communication channel between the SBC node servers and between CCM and SBC nodes.
- Media (MI): used for receiving and sending media payload.
- Signaling (SI): used for receiving and sending SIP signaling messages.
- Websocket Signaling (WS): used for receiving and sending SIP over websocket from and to WebRTC clients.
- Custom Interface (CI): used for different applications depend on admin (e.g. SSH, SNMP, HTTP proxy/redirect)

Each of the SBC interfaces is mapped to a physical or system network interface that is used for the actual sending and receiving of the data. Multiple SBC interfaces can be mapped to the same network interface. If VLANs are used, they are administered under management of physical interfaces and remain otherwise transparent to the rest of the system.

Administration of the SBC interfaces is described in the Section Interface Configuration.

Call Agents

Call Agents (CA) are the smallest type of peering entities the ABC SBC can differentiate. They represent logical end-points. They can be defined based on several addressing mechanisms:

- IP address and port
- Domain or host name and port
- IP network and mask

Additionally, a Call Agent is assigned to a signaling and a media interface. These interfaces are used whenever SIP signaling or media packets are sent to or received from a Call Agent.

For security reasons, the SBC communicates by default only with well-known and defined Call Agents. When an incoming SIP request cannot be attributed to a Call Agent, it is rejected.

To determine the source Call Agent, the SBC uses the source IP address and port of the request to search among the configured Call Agents. If the definitions of Call Agents are overlapping (for example when some Call Agents are defined with an IP address which belongs to a subnet used to define another Call Agent), the following descending order is used to determine the Call Agent:

- Call Agents with matching IP address and port.
- Call Agents with matching IP address but a port equal to 0.
- Call Agents with matching IP network (including mask) in descending order of mask length

Routing rules determine the target Call Agent. In this case, the interface used to send the SIP signaling is the one assigned to the target Call Agent. In case media relay is used, the media interface assigned to this target Call Agent is also used accordingly. The target Call Agent is used to determine the set of applicable rules on the outbound side as well. Note that Call Agents specified by subnet address cannot be used for routing.

Realms

Every Call Agent belongs to one Realm. Realms are the logical groupings of one or more Call Agents. They allow multiple Call Agents to share the same SIP processing logic without defining it individually multiple times.

In a classical context where the SBC is placed on the border of an internal network, it is common to define one Realm for the outside world, and one for the internal network. This way, all the restrictive rules to protect the internal network are defined for the outside Realm, while the internal one can be safely trusted.

In a peering use case, usually one Realm per peering partner is defined.

A-B-C rules

The ABC SBC is fundamentally rules driven. This means that almost all features can be activated based on certain conditions evaluated at run-time, based on parts of the signaling messages or media payload.

All rules are constructed using the same pattern. They consist of a set of one or more conditions. If all conditions apply (logical conjunction), a set of one or more actions is executed.

It is important to understand that rules are generally applied only on dialog-initiating requests or out-of-dialog requests. However, some actions have a scope that goes beyond these dialog-initiating requests or out-of-dialog requests. For example, header filters apply to all requests exchanged, including in-dialog requests. Action descriptions include their scopes where applicable.

There are three types of rules that are always executed in the same order: A, B, and C. A-rules describe how incoming traffic for a Call Agent/Realm is handled, B-rules determine destination for the SIP request, and C-rules describe SIP processing behavior specific to that chosen destination.

A and C rules are associated with Call Agents and/or Realms. The realm rules allow to have shared logic for all Call Agents that are to be handled the same way, while CA rules are suitable for individual logic. Often, rules are associated with both Realms and Call Agents. The Realm rules are executed first, and their results can be overridden by more specific Call Agent Rules.

B rules are different in that they are global. They are not associated with a specific realm or call agent. When processing of A-rules completes, the B-rules determine the next hop. That's is the only action the B-rules can perform. Then Realm and Call Agent specific C-rules are processed.

The FRAFOS ABC SBC handles calls according to the schema shown in the figure Call handling algorithm.

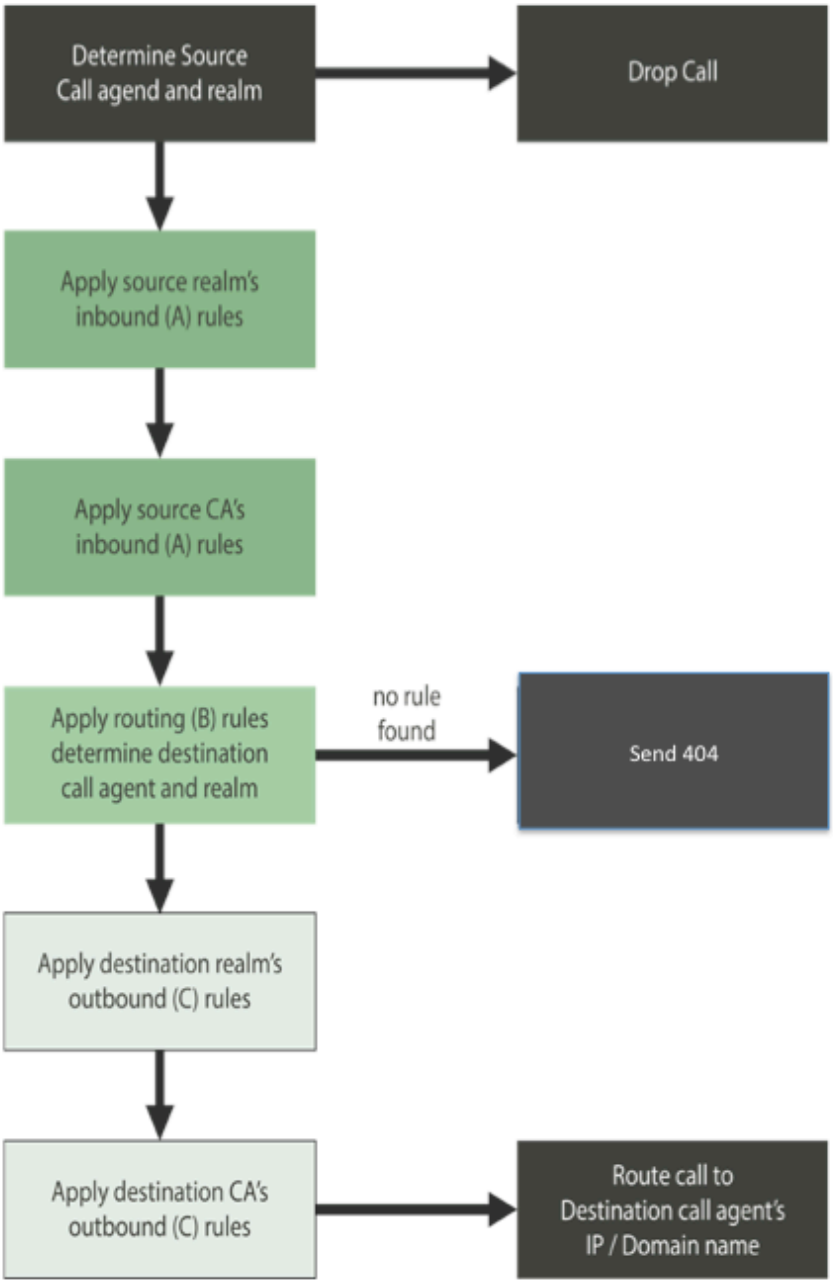

Figure 8: Call handling algorithm

It is a good practice to place rules in realms rather than in Call Agents, unless they are clearly specific to Call Agents. Rules in realms don't have to be repeated Call Agent by Call Agent and don't expose administrator to Copy-and-Paste administrative errors.

Conditions and Actions

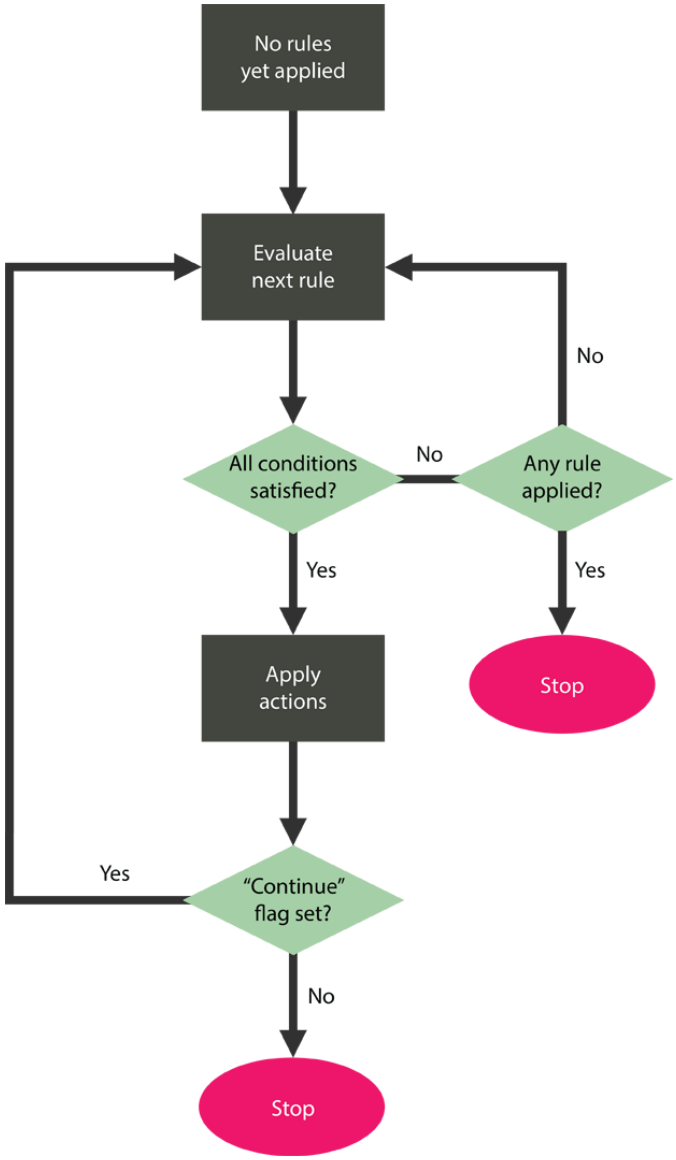

Figure 9: Rule evaluation sequence

Every A/B/C rule may have one or multiple conditions. The conditions can check incoming message content (method, R-URI, headers, codecs), source (IP, port, realm), values stored in call variables etc.

Rules have zero or more actions. If all its conditions are satisfied ('AND' combination), the actions of a rule are executed in the order in which they are defined. In case a rule does not contain any conditions, the rule's actions are always applied.

Actions can have parameters depending on the action type. For example, the action "Add Header" that appends a SIP header to an outgoing message takes two parameters - a header name and a header value.

Within a rule set (Realm A or C rules; Call Agent A or C rules), the SBC evaluates each rule by evaluating the condition set first. If all conditions match, the set of actions is executed. As part of the rule definition a "continue flag" is defined. If the "continue flag" is checked, the next rule is evaluated. Otherwise, the rule evaluation within this block stops. Irrespectively of the state of the "continue flag",

the rule evaluation continues with the next block. This means that if the "continue flag" is not checked in a Realm A rule where the conditions match, the Call Agent A rule will still be executed, see Rule evaluation sequence.

Routing rules

Routing rules have the same set of conditions as A & C rules, but only one possible action: route the request to a target Call Agent.

Figure 10: B-Rule evaluation

The essential role of B rules is to determine to which target Call Agent the processed request will be sent to. This also influences the outbound signaling and media interface used to send out the forwarded request.

To determine a target Call Agent, the SBC will evaluate the conditions for each routing (B) rule. Once a match is found, the Call Agent associated with the routing rule is determined as the target Call Agent. Then, the processing continues with the C rules assigned to the target Call Agent.

As depicted in Figure B-Rule evaluation, all active B rules are traversed and evaluated sequentially. In case all conditions of a rule are satisfied, the destination call agent and routing method are successfully determined. In case that the conditions of a rule are not satisfied, processing continues with the next rule. If no matching routing rule can be found, then the call is refused with a *404 Not Found* error code and an *error* event type is produced.

4.3 Practical Guide to the ABC SBC

4.3.1 Network Planning Guidelines

This section provides you with a list of steps every network administrator shall walk through carefully before deploying an SBC-powered network. Early planning helps to create robust network that well serves the needs of its users and make administrator's life free of surprises.

Each planning step starts with a question about a particular network planning aspect. Administrators need to to ask themselves this question to determine their configuration needs. It is then followed by a short debate of the most important configuration options and trade-offs. Not all available options are included, yet those present are of major importance and shall be answered in the early planning phase.

The subsequent section includes a checklist that reiterates question raised in the section. We recommend going through it thoroughly before a deployment is commenced, and also completing the answers in the "Customer Site Survey" document available from our customer care.

The steps in this section are grouped as follows:

- The Section Topology Model describes how an ABC SBC connects to the IP networks, how it models SIP devices and groups them administratively.
- The Section SBC Logic describes the anticipated behavior of the ABC SBC and what needs to be considered when configuring it: routing, media processing, NAT handling, and more.
- The Section Security Policies summarizes configuration steps needed to protect both the connected SIP networks and the SBC itself.
- The Section Capacity planning provides guidelines for estimating the SBCs cluster dimensions.
- Eventually the Section IT Integration discusses the configuration steps that need to be considered if the SBC connects to other network management elements.

Topology Model

The key function of the SBC is to securely connect various SIP elements and peering networks together. This is not trivial because the networks and devices may use different security policies, SIP protocol extensions, dialing plans, codecs, etc. To deal with this variety the SBC uses a network model in which the SIP devices are represented as abstract "Call Agents" that are grouped in "Realms". With Call Agents (CAs) and Realms there are rules associated that characterize how their traffic from and to them is treated.

The topology planning process includes the following steps:

IP layer topology

To which IP networks does the SBC connect?

The first step in defining your topology is located at the IP layer: you have to specify which IP networks connect with each other and how the SBC connects to them. This is captured in the specification of "interfaces".

Interfaces describe in detail how an SBC connects to IP networks. In the simplest case all traffic can be routed through a single Ethernet card using a single IP address and dedicated UDP port range. Many deployments use multiple Ethernet cards or VLANs to connect two or more physical networks with different levels of security.

A widely used practice we recommend is use of three network cards for three networks: unprotected public, protected private, and highly-protected administrative. Then for example residential SIP phones and peering providers connect to the SBC over the public network, operator's PSTN gateways are located in a private IP subnet, and administrators access the SBC over the administrative networks.

Note that using fewer cards makes a clean separation and security of traffic more difficult and also reduces total throughput.

More detailed discussion of interface configuration is described in Section Physical, System and SBC Interfaces.

IP layer security

Which firewall rules shall be used in firewalls and the SBC?

Once the IP connectivity is specified, you need to specify L3/L4 restrictions for your deployment. This consists of two parts: if you deploy additional L3/L4 firewalls in front of the SBCs you must make sure that they do not restrict legitimate traffic. You must allow SIP traffic (typically UDP/TCP port 5060), and if media anchoring is used also unprivileged UDP ports. On the SBC you may take the opposite approach and restrict critical ports, especially if no L3/L4 firewall is used. At least you shall make sure that traffic to privileged ports used for administration is permitted only from trusted IP addresses.

More can be found in the Section Manual IP-layer Blocking .

Call Agents (CAs)

What SIP Devices does the SBC talk to?

Identify CAs by IP address, IP address range or DNS name. A CA may be physically a PSTN gateway, a whole "cloud" of identical IP phones located in a subnet, a peering party, just anything with unique identification and characteristics. Also specify if there is some specific treatment a CA shall obtain and that needs to be specified in SBC rules. These frequently include:

- traffic limitations, i.e. will you impose one of these constraints on the traffic: RTP bandwidth, signaling rate, number of parallel calls?
- mediation rules, i.e., do you to reconcile dialing plans, identity (URI) usage, header usage? Note that some of these rules cannot be easily planned for ahead of time as they are used to fix protocol imperfections discovered after fact in operation.
- NAT handling, i.e. does presence of NATs necessitate use of media relays, and shall the SBC handle traffic symmetrically? (normally the answer is yes to both these questions if there are NATs present).

More details about traffic limitations are described in the Section Traffic Limiting and Shaping, mediation is described further in Section SIP Mediation and media anchoring is described in more detail in the Sections NAT Traversal and Media Anchoring (RTP Relay).

Realms

What SIP Networks does the SBC talk to?

Most CAs belong to an administrative zone, whose traffic the SBC handles the same way. It would be impractical to define traffic rules for every single CA in such a zone. Therefore the SBC uses the concept of Realms that group all CAs sharing the same characteristics. For example a total bandwidth maximum restriction may be applied to a whole cluster of peering partner's PSTN gateways modeled as a Realm. Also identical header-field manipulation and routing may be applied to all the machines in a Realm. Therefore the administrator needs to assign CAs to the Realms and associate the common rules with the Realm. The functionality of Realms' rules is the same as for CAs.

SBC Logic

It is important to plan what the SBC will actually do for your network in precise terms because particular features have further impact on capacity planning, integration with other components, interoperability and administration.

Routing

What will be the routing criteria used in your network?

Routing is a mandatory part of every SBC configuration. Once the topology is established, you must define how traffic flows between the Realms and Call Agents. That is described in routing tables. The key decision to be made is what is the criteria used to determine the next hop for a new session. The most common examples of criteria include:

- prefix-based routing. This is frequently used when you have a number of PSTN gateways serving different regions. Technically you match area codes against beginning of the user-part of the request URI.
- source-based routing. This is frequently used when you connect multiple IP networks and want to make sure that all traffic from one network is forwarded to the other and vice versa. The criteria is then the source IP address, source Call Agent or Realm.
- method-based routing. Sometimes specialized servers are used for processing specific traffic, like message stores for keeping messages for off-line recipients.
- The SBC configuration options include even more criteria and these can be also combined with each other.

Some functionality is only present in some deployments and whether to use it or not depends on used equipment, network characteristics and network policies. More information about administration of SIP routing may be found in the Section SIP Routing.

Media Anchoring

Do you need the SBC to anchor media so that all RTP traffic visits your site?

The ideal answer is no due to latency and bandwidth concerns, the most common answer is yes due to NAT traversal and controlling media. Relaying media costs considerable bandwidth and implies more SBC boxes. Yet if any of the following conditions applies, you will have to enable media relay:

- There are SIP clients behind the NATs. That's the common case in residential VoIP.
- You wish to record calls. Obviously you can only record media that visits the SBC.
- You want to implement topology hiding consequently and make sure that no party sees media coming from any other IP than that of the SBC.
- The SBC connects two networks that are mutually unroutable.

More administrative details can be found in the Section Media Anchoring (RTP Relay).

Media Restrictions

Do media-restricting rules need to be placed?

The need for media restrictions arises mostly when bandwidth is scarce. This may be the case if media anchoring is used on a link from/to the SBC or on the SBC itself. It may be also the case on the link to the client, particularly if it is a mobile one.

The simplest solution is to restrict media negotiated by Call Agents by putting desirable codecs on a whitelist. All other codecs will be removed from codec negotiation. It may happen though that the resulting codec subset is empty and the Call Agents would not be able to communicate with each other.

If there are no codecs left, you may extend the codec set by transcoding. The SBC then adds additional codecs to the negotiation process and if the Call Agents choose it, the SBC will convert media to the chosen codec. The penalty that needs to be considered is degraded throughput of the SBC.

In addition to the codec-based proactive bandwidth saving approach, the SBC can also limit bandwidth retroactively and put bandwidth limits on CAs or Realms (or some portions of its traffic). This helps to stay on the bandwidth budget even if SIP devices exceed traffic signaled in SIP. However, it remains a

reactive measure. That is, it does not prevent excessive traffic, it just drops it and impairs the affected media streams.

Codec handling is described in the Sections Media Type Filtering, CODEC Filtering, CODEC Preference and Transcoding, administration of media limits is described in Section Traffic Limiting and Shaping.

Registrar Cache

Does the SIP traffic include REGISTER messages?

If so, we recommend that you do enable registrar cache. The cache is optimized to reduce the REGISTER traffic that is passed down to the registrar. This is particularly important if the clients are behind NATs. Then the cache must be configured to force SIP clients to re-register every minute to stay connected from behind NATs. Also the ability to track registration status of users allows the SBC logic to differentiate call processing for online and offline users. This can be used for example for voicemail routing.

Further administrative details are described in the Section Registration Caching and Handling.

NAT Handling

Are there some CAs behind NATs?

If so, you not only have to anchor media as described above, but also make sure that the signaling protocol traverses NATs successfully. Also registrar-cache must be used to force clients to refresh their connectivity using frequent re-registrations. Some deployments with STUN-capable SIP phones also set up a STUN server to assist these phones.

NAT configuration is described in further details in Sections NAT Traversal and Media Anchoring (RTP Relay).

SBC High Availability

Shall the SBC be operated in high-availability mode?

While this is normally the case, small enterprise deployments may prefer buying and administering fewer boxes. Introducing high-availability requires a standby spare host machine for every active SBC and effectively doubles the number of machines.

It is recommended that in a high-available configuration setup an administrative network is used for internal inter-node communications and the availability protocol used between the machines in the active/standby pair.

More administrative details about HA mode are available in the Section High Available (HA) Pair Mode.

Downstream Failover and Load-Balancing

Shall the SBC seek alternative destinations when primary destinations become unavailable?

Handling downstream failover may or may not be needed. For example if the downstream telephones are single-user SIP telephones, there are usually no backup devices. Some high-density devices like PSTN gateways implement automated failover in a way which is invisible to the SBC and the SBC doesn't need to handle it either. However if the primary destinations have spare backup machines without automated failover, the SBC can still detect a failure and try the alternate destinations.

If there are multiple alternate destinations, it may be also practical to spread the load among them.

There are several ways how to define a set of alternate destinations and their priorities: it can be defined in DNS maps or in the Call Agent specification. If the definition is managed in DNS, the SBC resolves DNS names automatically in compliance with **RFC 3263**. If using DNS is not practical, the same effect can be achieved by associating multiple IP addresses with a Call Agent. Additionally, a backup Call

Agent may be also associated with a Call Agent: in that case traffic to the backup destination will be processed by additional C-rules specific to the destination.

Procedures for determining the next hop are described in Section Determination of the IP destination and Next-hop Load-Balancing.

Dialing Plan Mediation

Do different CAs and Realms connect to the SBC use different dialing plans?

Often SBCs connecting different sites that use different numbering conventions: short-dials, regional dialing plans, special services numbers. To enable interconnection of such sites and avoid number overlaps, the SBC must bring all the numbers to a common denominator, mostly the E.164 numbering format.

More about mediation can be found in the Section SIP Mediation.

Security Policies

Generally, the SBC has two ways for protecting networks: putting various restrictions on traffic and concealing network internals. The latter is sometimes a double-edged sword as obfuscation of SIP traffic makes it hard to troubleshoot.

Restricting Traffic from Unwanted Sources

How do you identify and discard illegitimate traffic?

There are several ways the SBC recognizes and drops undesirable traffic.

At the SIP-level you may set a variety of criteria which if it is met results in declining a session request. The conditions may include:

- unusual message patterns such as User-Agents of a type known to offend other SIP devices, URIs to premium numbers or simply anything else that can be matched
- unusual traffic patterns, such as call excessive call rate, number of parallel calls, or RTP bandwidth consumptions
- The traffic patterns apply statically to a whole Realm or Call Agent. However they may be also tied dynamically to "traffic from any single IP coming from the Realm" or "traffic to any single phone number". This way you could for example impose a Realm condition "maximum one parallel call from a single IP address to a 900 phone number".

Additionally, if traffic from some specific IP address begins to take really excessive dimensions, you can drop it straight at the IP layer before it reaches the SBC logic.

More information about filtering unwanted traffic can be found in Section Police: Devising Security Rules in the ABC SBC.

Topology Hiding

Do you prefer SIP transparency across networks or concealing network information?

This is indeed an operational dilemma. If you process SIP traffic "by standards", the traffic will be passing the SBC with minimum changes. This approach will reveal lot of information about one network to the other: which IP addresses are being used, which port ranges, what type of equipment and potentially even more. This makes life easier for attackers seeking security holes in networks and therefore some operators chose to obfuscate this information.

The penalty for traffic obfuscation is significant however: operators' administrators will find it similarly hard to find out what's is going on in their own networks. That doesn't make troubleshooting easier. Some complicated applications in which SIP messages tend to refer to each other (such as in call transfer) may also fail.

The choice to obfuscate or not is eventually to be taken by the operator. The ABC SBC has the following means of doing that:

- Topology hiding rewrites known SIP header fields in which use of IP addresses is mandatory. The downside is that troubleshooting becomes more difficult.
- Use of non-transparent mode will rewrite dialog-identifying information: from-tag, to-tag and Call-ID which in some older implementations also includes IP addresses. The downside is some applications which refer in protocol to a call may fail.
- Header whitelisting drops all header-fields that may potentially carry additional sensitive information, standardized (*Warning*, *User-Agent* for example) or proprietary (*Remote-Party-ID* for example). The downside is that sometimes "a baby can be thrown out with the bath water", when the header-fields include potentially useful information.
- Media anchoring can be used to obfuscate where media flows from and to. The downside is high bandwidth consumption and increased latency if media anchoring wouldn't be used otherwise.

Additional information can be found in the Sections Topology Hiding, SIP Header Processing and Media Anchoring (RTP Relay).

Capacity planning

Capacity planning is a key part of the planning exercise. Failure to provision resources sufficiently can lead to network congestion and low quality of services. Overprovisioning way too far increases cost. The goal is to find the right measure of network size that serves the anticipated traffic. This section provides rules of thumb to estimate needed capacity and makes simplifying assumptions about state-of-the-art hardware, "normal" traffic patterns and no dependencies on external servers. A more detailed discussion of dimensioning can be found in the Section SBC Dimensioning and Performance Tuning.

Cluster Size

How many SBCs are required for a deployment?

There are two major factors that determine how many hosts you need to serve your traffic: anticipated performance bottleneck and organization of clusters.

Which bottleneck is the most critical strongly varies with actual traffic patterns and services configured on the SBC. A rule of thumb for a rough estimate of the performance of the ABC SBC on PC with three-1GB-Ethernet and 12 GB of memory is this:

- If transcoding is used, the bottleneck of a single box in terms of the maximum number of parallel calls which is about 1000. Otherwise…
- … if media-anchoring is used, the bottleneck in terms of the the maximum number of parallel calls which is about 5000. (Media overhead prevails even over heavy registration load.)
- Otherwise the limit is a call rate of 480 calls per second.

We advice to add at least additional 35% of buffer capacity to deal with variances in hardware performance, increasing traffic patterns, too conservative traffic forecasts and DoS attacks.

Once you determined the per-box capacity, you need to take cluster organization in account. There are the following three cases:

- a single SBC deployment: no scaling, no high-availability.
- high-available active/standby pair: the pair has still the total capacity of a single box, however it can survive scheduled and unscheduled outages without service impairment
- high-available cluster: the number of boxes is determined by number of boxes needed to serve the target capacity, doubled to achieve high availability plus two more boxes for a highly-available load-balancer: cluster_capacity / box_capacity * 2 + 2.

Bandwidth

How much bandwidth needs to be allocated to serve the deployment?

To determine needed bandwidth you need to discriminate between two cases: using SBC with and without media anchoring. The more bandwidth-hungry case is that with media anchoring. With the most commonly used codec, G.711, a call consumes 197 kbps bandwidth in each direction.

To determine the maximum bandwidth needed calculate the product of maximum number of parallel calls by the bandwidth specific to the codec in use, 197 kbps if it is G.711.

Public IP Address Space

How many public IP addresses need to be allocated for an SBC Cluster?

The minimum number is one shared VIP address for every active/standby pair.

IT Integration

An SBC is rarely a standalone component. More often it integrates with other components for the sake of connecting to external policy logic, network monitoring, server naming and others. This section lists typical integration options you may need to consider for your deployment.

RESTful interface

Does the SBC need to consult an external server for its decision making?

If so, the ABC SBC built-in RESTful query allows to ask an external server how a session shall be handled. This query possibility allows to integrate external and complicated logic in the SBC which is customer-specific or for other reasons difficult to integrate with the SBC directly.

See more in the Section RESTful Interface.

Recording

For various reasons, audio recording may need to be configured. What needs to be integrated is access to the recorded files. The easiest way is none: the recorded files are stored on local storage and accessed through the events web-page. Uploading to HTTP may also be used. In either case, some deletion and retention policy must be created, otherwise the local storage will be soon full.

See more in the Section Audio Recording.

Monitoring

Do you need to see how the SBC is doing?

Of course you do. We suggest use of the optional ABC Monitor as it provides ABC SBC administrators with full history of users and analytical tools to audit it.

You can also use SNMP at a third-party management console to inspect health of your ABC SBC and the networks. It is possible to define your own custom counters. See more in the Section Measurements and Monitoring.

Mass Provisioning

Do you need to provision the SBC with lot of repetitive data such as thousands of Least-Cost-Routing Entries?

Then you certainly do not want to provision it rule by rule. Instead you devise one rule and fill it with data. The actual data can comes over a web interface REST API or RPC.

See more in the Section Provisioned Tables.

Call Detail Record (CDR) Exports

Do you need to access CDRs for sake of charging and reconciliation?

Then you must access the internally produced CDRs.

See the Section Call Data Records (CDRs) for more about CDR location, format and access.

DNS Naming

How do I make the SBCs known to their counter-parts?

While it is possibly to communicate with peer SIP-devices only using IP addresses we recommend that every single SBC has a DNS name which is communicated as the point-of-contact to its peers. If nothing else, it makes IP renumbering much easier should it occur.

DNS map entries for SIP servers follow the SRV DNS extension as described in **RFC 3263**.

4.3.2 Planning Checklists

This section provides you with a summary of questions raised in the previous section. We urge that you diligently check all the items before you proceed with commencing an installation.

Topology

- Have you identified all Call Agents present in your network?
- Have you specified additional processing rules for these Call Agents, such as network limits?
- Have you grouped all Call Agents in Realms present in your network?
- Have you specified additional processing rules for these Call Agents such as network limits?
- Have you specified all physical interfaces (Ethernet cards)?
- Have you specified all IP addresses, port ranges, and VLANs to be used on these interfaces?
- If there are firewalls in front of the SBC, have you verified that all needed ports are open?
- Have you verified that IP rules on the SBC restrict traffic to privileged ports from trusted IP addresses only?

SBC Logic

- Have you devised the SIP routing criteria used in your network? How many routing rules do you anticipate?
- Have you devised the routing flows between Realms and CAs?
- Does any of the conditions mentioned necessitate use of media relays?
- If you need to restrict codecs, which codecs shall be permitted and which codecs shall be restricted?
- Do you need to force use of a codec unsupported by a CA using transcoding? Which codec?
- If your SIP traffic includes REGISTER messages, will you enable registrar cache? If so, what will be the registration interval?
- Does presence of clients behind NATs necessitate use of media-relay, symmetric SIP and registrar-cache?
- Do you plan to set up high-available SBC pair(s)?
- If you use the high-available (HA) SBC pair, do you plan to use an administrative network for the HA protocol?
- If you need to handle downstream failover, have you devised appropriate DNS maps?
- Does the SBC accept traffic using different dialing conventions? If so, how will you translate between them?

Security

- What conditions do you devise to drop illegitimate traffic? Will you configure IP-based and/or URI-based blacklists?
- Will you introduce traffic shaping limits: call-rate, call-length, parallel calls and maximum call-length?
- Will all or only registered SIP devices be permitted to make phone calls?
- Will the need to troubleshoot your network easily or the need to hide topology prevail?

Dimensioning

- How many SBCs do you need?
- How many network cards shall each SBC have?
- How many IP addresses do you need?
- How much bandwidth do you need in each direction?

Integration

- Do you plan to use the management components over a dedicated administrative network?
- Is external session decision-making logic using RESTful interface needed? If so, what are the parameters passed from and to the RESTful server?
- Do you need to mass-provision some configuration data? What is the structure of the data and what size of tables do you anticipate?
- Do you need to record audio and access it? What is your deletion and retention policy for the stored audio files?
- Do you need to export CDRs?
- Have you devised appropriate DNS SRV and A entries for all IP addresses?

4.3.3 A Typical SBC Configuration Example

Many SBC deployments, especially in smaller networks, follow a simple schema which is given through the network structure. In this typical network, the SBC bridges between an internal network, where the home proxies, PBXs and other servers like conference and application servers are located, and the public network, where the user agents reside. Typically, in such a network the main motivations for deploying an SBC are

- network separation for security reasons
- foolproof and always-working NAT handling
- protection of the core network from high registration load
- protection against fraud by enforcing call limits
- possibility for monitoring and tracing for troubleshooting

This chapter presents step by step how to address these network aspects using the ABC SBC. It assumes an SBC "sitting" between two networks, a public one with user telephones and a private protected one with operator's infrastructure.

Identifying Network topology

Simple as it is in this case, the network topology is shown in Sample network Topology.

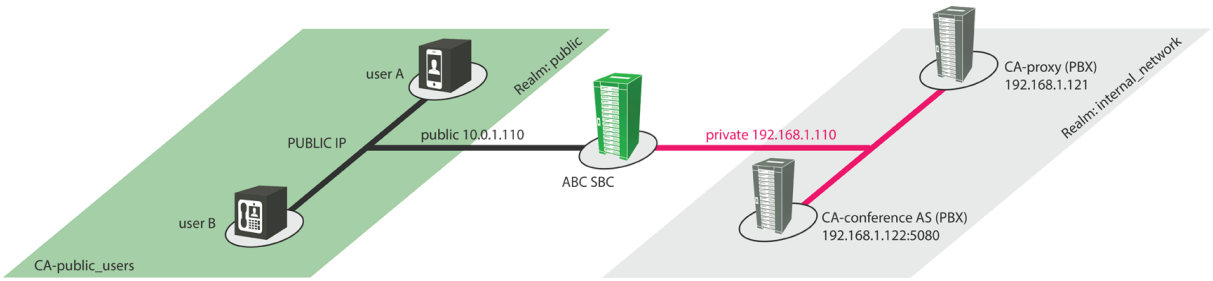

Figure 11: Sample network Topology

What administrator needs to do in this step is configuration of the physical network interfaces and of the SBC-level interfaces.

The ABC SBC has two physical interfaces, one *public* connecting to the public networks, here with IP address 10.0.1.110, and one *private* connecting to the private network, here with IP address 192.168.1.110. The physical interfaces are configured using procedures described in Section Physical and System Interfaces.

User agents are located in the public network and have IP addresses from any network, and they are configured to use the public interface of the SBC with the address 10.0.1.110 as proxy (in a real world deployment, this address would not be a private RFC1918 address, but a public one).

A proxy (or PBX) and a conference (or other application) server are located in the internal network. The ABC SBC can communicate with the entities in the internal network through its interface in the private network which has the IP address 192.168.1.110.

The detailed procedure for setting up SBC interfaces is described in Section SBC Interfaces. It links media processing, signaling and administration with physical interfaces, IP addresses and port ranges.

Describing ABC SBC Realms and Call Agents

The network topology is described in the ABC SBC configuration by Realms and Call Agents. Call Agents are typically consumer or operator SIP devices identified by their IP addresses or DNS names. They are grouped in networks called Realms whose processing rules they share.

In our example two Realms are created in the SBC: *public* and *internal_network*.

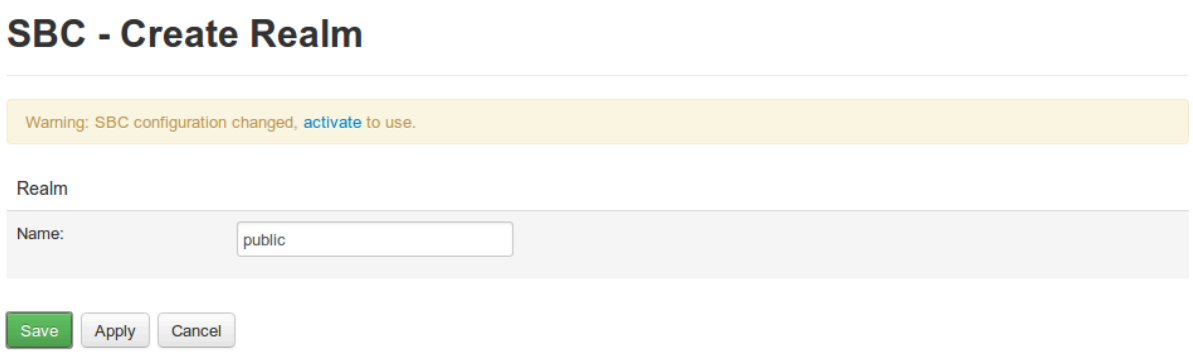

Figure 12: Creation of Realm

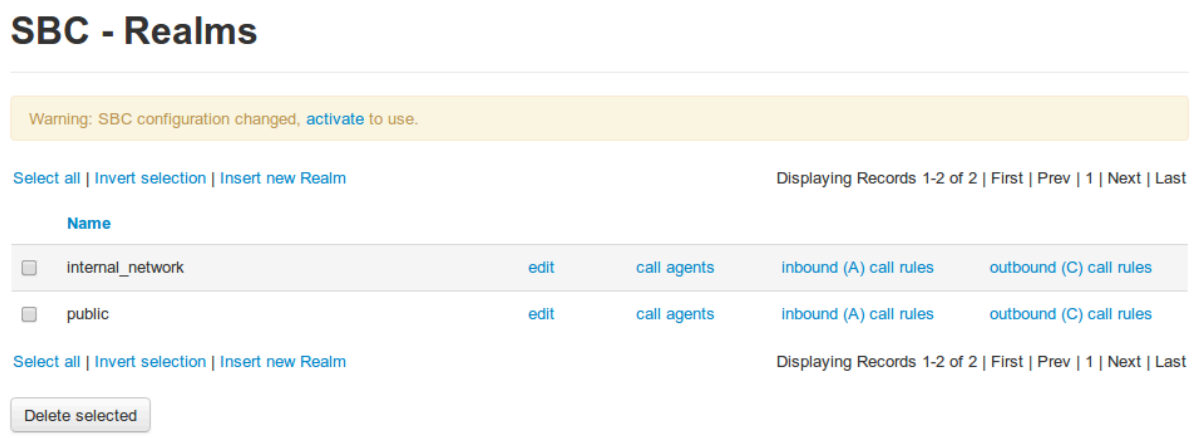

Figure 13: Public and private Realms

In the Realm *public*, the call agent *public_users* is created with IP address 0.0.0.0/0, which means that public_users can have any IP address, or: requests received from any IP address on the public interface will be identified as coming from the Call Agent *public_users*. The address list can include multiple addresses that are used for routing (See section Determination of the IP destination and Next-hop Load-Balancing). Also a backup call agent can be defined here which can be used as alternate destination if forwarding to the primary destination fails. The CA definition further specifies interfaces used for sending and receiving signaling and media and availability management information – see Section IP Blacklisting: Adaptive Availability Management for more information.

The call agents could be assigned to SBC nodes and/or config groups. This assignment basically specify what SBC nodes is the call agent known to.

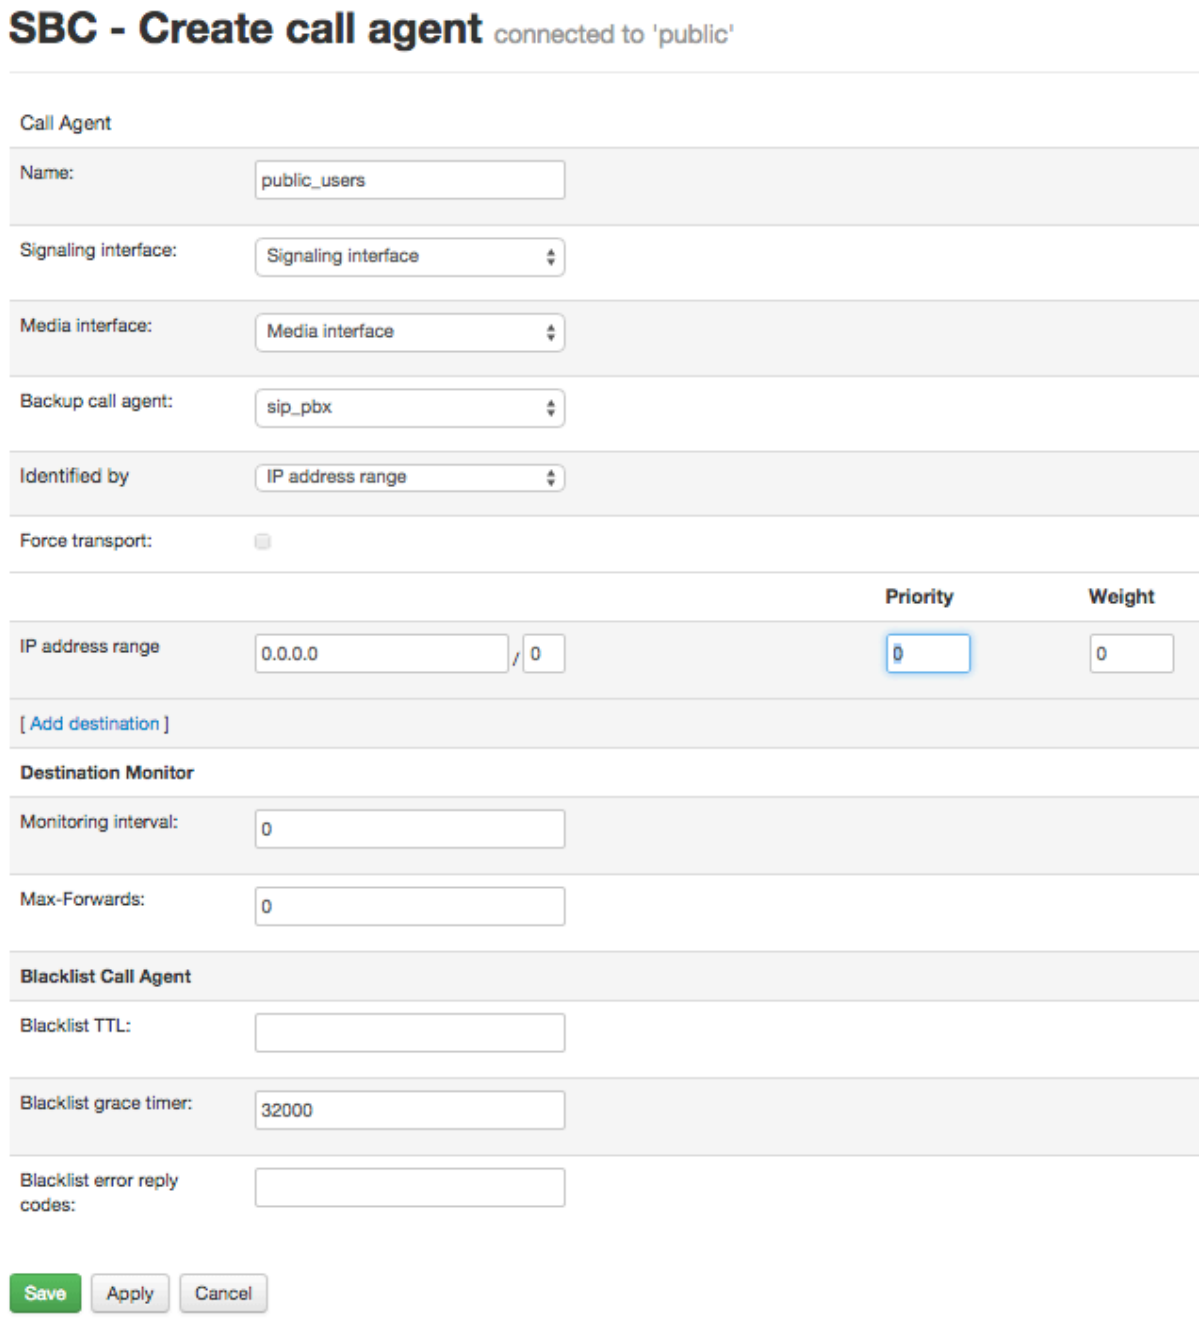

Figure 14: Create public-users Call Agent

As we have neither defined a specific IP:port for the Call Agent nor a hostname, requests can be routed to that Call Agent only by Request URI, or by setting the destination IP explicitly in the routing rule.

Figure 15: Public Call Agents list

For the internal realm, the call agents *proxy* and *conference* are created with IP addresses 192.168.1.121 and 192.168.1.122:5080 respectively.

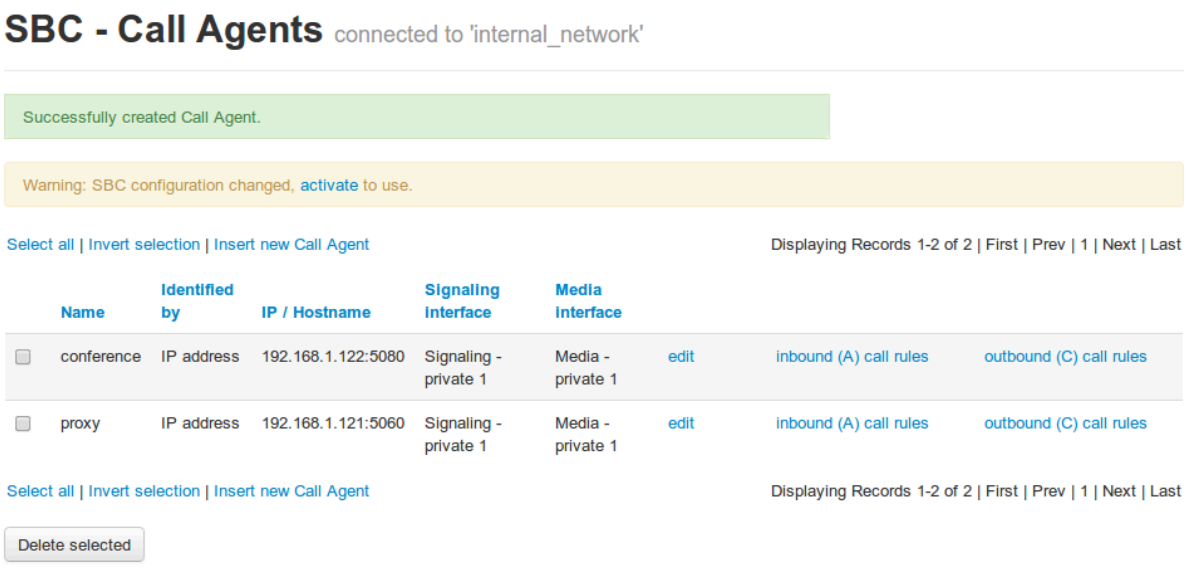

Figure 16: Internal Call Agents list

Provisioning Call Agents Using RPC

It is also possible to provision Call Agents using XML-RPC interface.

The following RPC commands for Call Agent provisioning are available:

- cagents.fetch() - fetch all the Call Agents
- cagents.insert($payload) - insert Call Agent
- cagents.update($payload) - update Call Agent
- cagents.delete($realm_name, $cagent_name) - delete Call Agent

The $payload parameter of the *insert* and *update* functions is structure of following format:

```
{
    "realm": "Some_Realm_Name",
    "name": "Some_Call_Agent_Name",
    "interface": "public_signaling",
    "media_interface": "public_media",
    "target": SEE_BELLOW
    "transport": "UDP",
    "backup_ca": {
        "ca":     "Backup_Call_Agent_Name",
        "realm":  "Name_of_Realm_Backup_Call_Agent_Belongs_To"
    },
    "backup2_ca": {
        "ca":     "2nd_Backup_Call_Agent_Name",
        "realm":  "Name_of_Realm_2nd_Backup_Call_Agent_Belongs_To"
    },
    "config_groups": ["Config_Group_Name"]
    "attrs": {
        "name_of_attribute_1": "value_of_attribute_1",
        "name_of_attribute_2": "value_of_attribute_2"
    }
}
```

For call agents identified by subnet, the target should look like this:

```
"target": {
    "subnet": ["0.0.0.0/32"]
}
```

For call agents identified by IP address or DNS name, the target should look like this:

```
"target": {
    "hosts": [
        { "addr": "192.168.1.1:5080", "weight": 10, "priority": 10 },
        { "addr": "192.168.1.2:5080", "weight": 10, "priority": 20 }
    ]
}
```

When updating Call Agent only the *realm* and *name* fields are mandatory, They identify the Call Agent to be updated. If any of the other fields is not specified it is not changed by the update action.

Provisioning Call Agents Using REST API

See details in the description of CCM REST API in API reference.

Configuring Registration Cache and Throttling

REGISTER processing accommodates several goals: off-loading servers behind the SBC, enforcing frequent re-registration load to keep NAT bindings alive and dealing with REGISTER avalanches caused by different sorts of outages.

For REGISTER requests coming from the "public side", the ABC SBC is configured to cache the registrations using the **Enable REGISTER caching** action. The cache works as follows:

- For every new registration, it creates an *alias*, a special unique one-time identifier.
- It saves the original contact along with the alias in the local registrar cache.
- To facilitate NAT traversal, it also saves the IP address, port and transport with which the REGISTER was received.
- It may re-adjust re-registration period so that it is frequent towards client for NAT keep-alives and less frequent downstream for better performance.
- It replaces the Contact in the REGISTER with a combination of the alias and the SBCs IP address: alias@SBC_IP:SBC_PORT.

This way, the "aliased" contact propagated downstream hides details of NAT-related address translation performed at the SBC and manipulates re-registration period as needed. The cache entry becomes effective once the REGISTER request is positively confirmed by the downstream SIP element.

Thus, when the REGISTER request is then routed to the registrar (the home proxy, here Call Agent *proxy*), the alias@SBC_IP:SBC_PORT is saved as he registered contact address of the user at the registrar.

We define this rule in the A rules of the *public* Realm, so that it is executed for REGISTER requests coming from any user agent defined under the Realm.

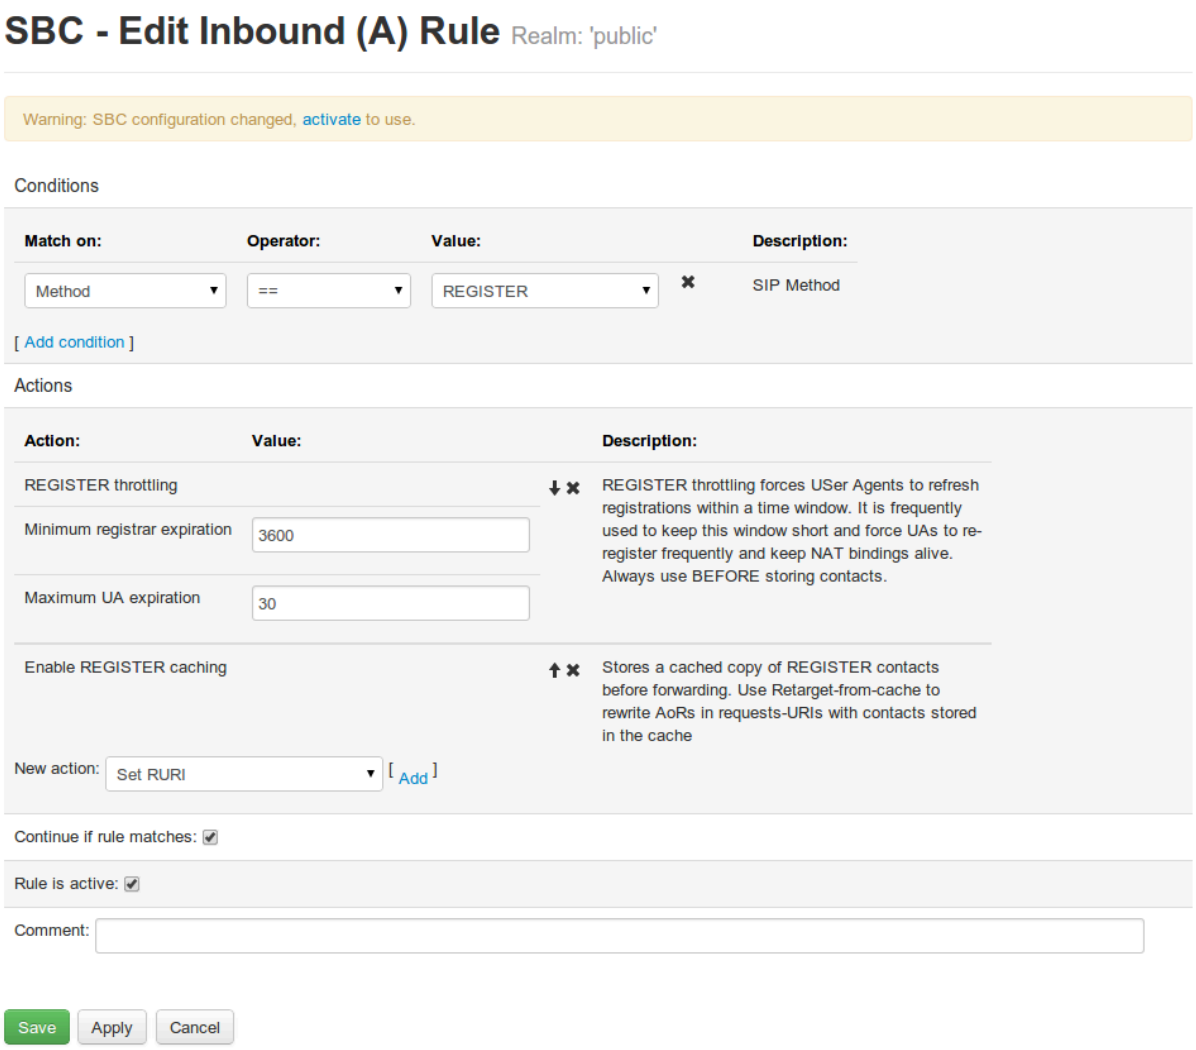

Figure 17: Rule A Definition for caching REGISTERs coming from public realm

In order to protect the home proxy from the bulk of the registration load, the action **REGISTER throttling** is enabled with a **Minimum registrar expiration**, i.e., the re-register interval used upstream to the home proxy, set to the default of 3600 (one hour), while the **Maximum UA expiration**, i.e., the re-register period for the user agents, is set to 30 seconds.

In cases where the call agent for the registrar have two destination addresses configured to work in a "round-robin" fashion (e.g. same priority), it may be desired to force the subsequent re-registers to the same destination. In order to achieve that, a rule similar to the following can be configured:

- A condition "Method Is -> REGISTER",
- a condition "Register Cache -> Is Registered",
- a rule "Fetch home-proxy IP"

Figure Register throttling destination binding shows this configuration on GUI.

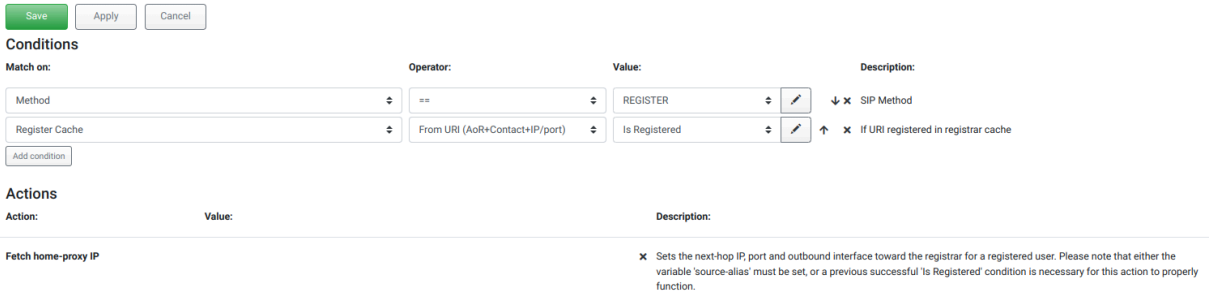

Figure 18: Register throttling destination binding

SIP Routing

The SIP routing tables (B tables) define to which Call Agent a call is forwarded. In our example, there are two cases: calls from the UAs towards the proxy server and calls from the internal network towards the UAs.

Calls from the User Agents are routed towards the proxy with a simple rule. Here we route all calls from the public realm to the proxy - we might also set a filter on Source Call Agent, which would be equivalent in our case. We route by setting the next_hop (the destination IP address) directly.

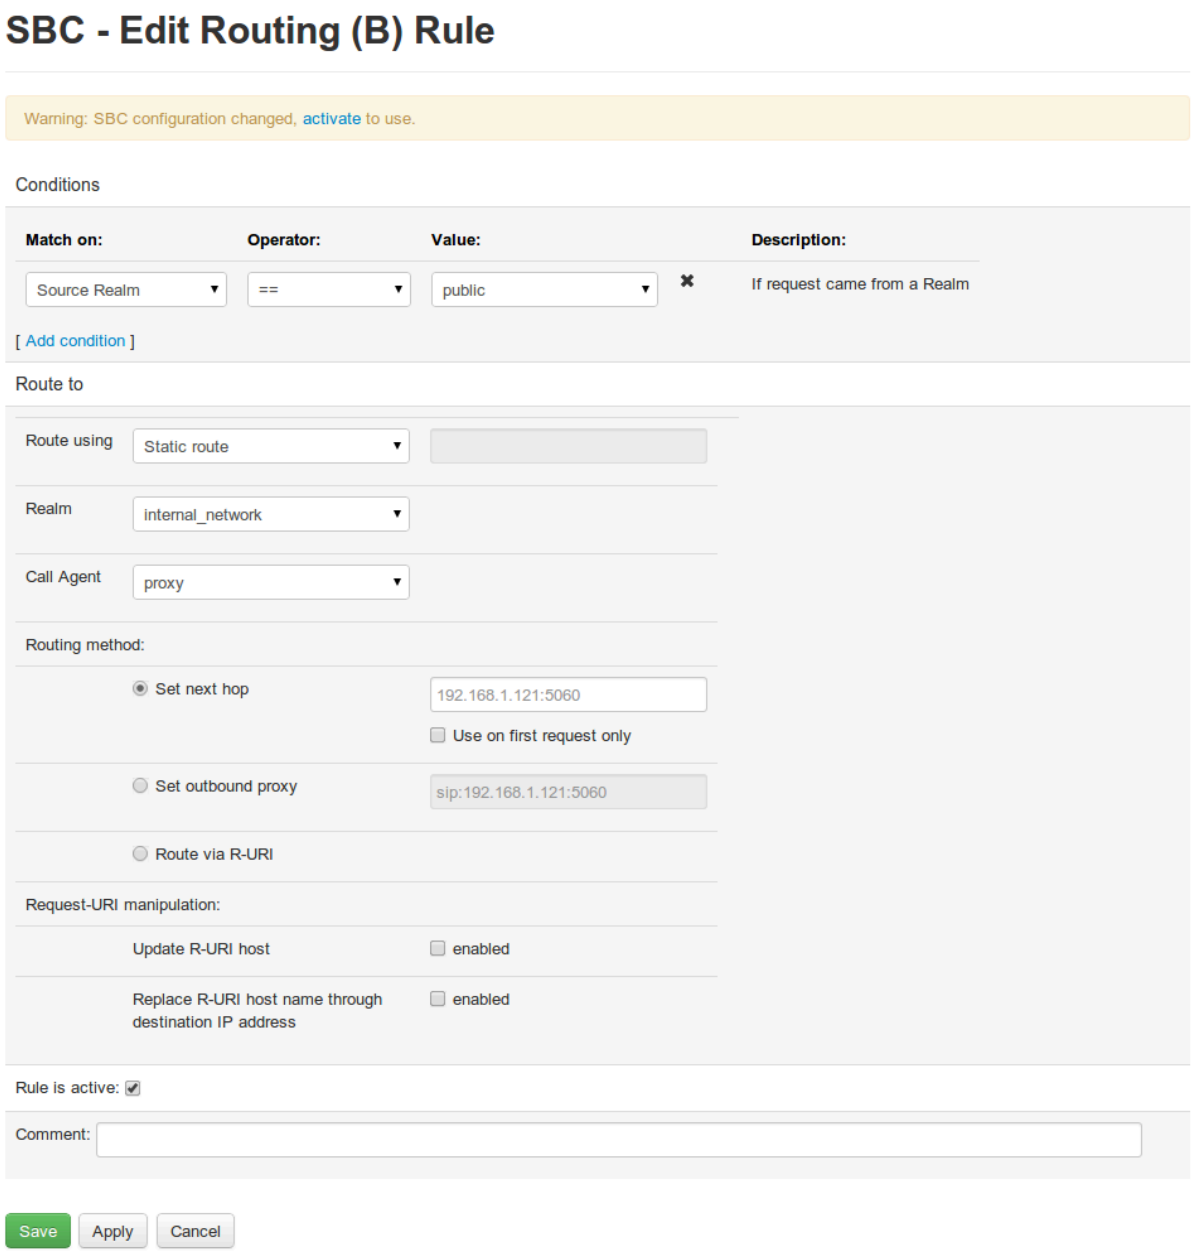

Figure 19: Rule B Definition for the sample network

The next rule specifies routing of all calls from the internal network towards the registered UAs. If the home proxy wants to send a call to a user, it finds in its registrar database the alias@SBC_IP:SBC_PORT as contact for the user, thus it sends the call to the SBC with the alias in the request URI like this: INVITE sip:alias@SBC_IP:SBC_PORT.

In the SBC, we use the action **Retarget R-URI from cache (alias)** to look up the UAs IP and port values and set the request-URI to it. We also use the **Enable NAT handling** and **Enable sticky transport** options to handle NATs properly. Using these options the SBC will send the request to the IP and port where the REGISTER request was received from and using the same transport protocol it was received on.

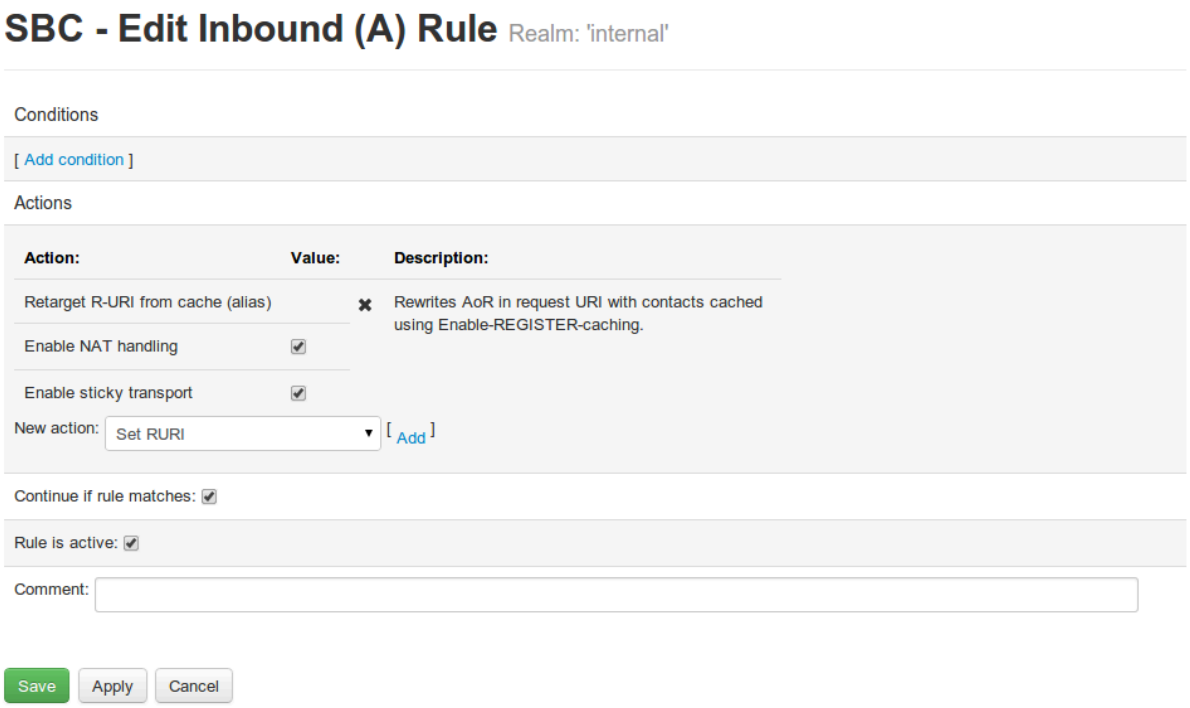

Figure 20: Rule A Definition for internal CAs

We can then use the R-URI to determine request's destination. For simplicity, in this example we define a catch-all routing rule for the complete internal network, which includes all call agents defined there. (We may also define special routing rules for the different call agents in the internal network if they would have to be treated separately, e.g. if some calls need to be sent to a peering partner.)

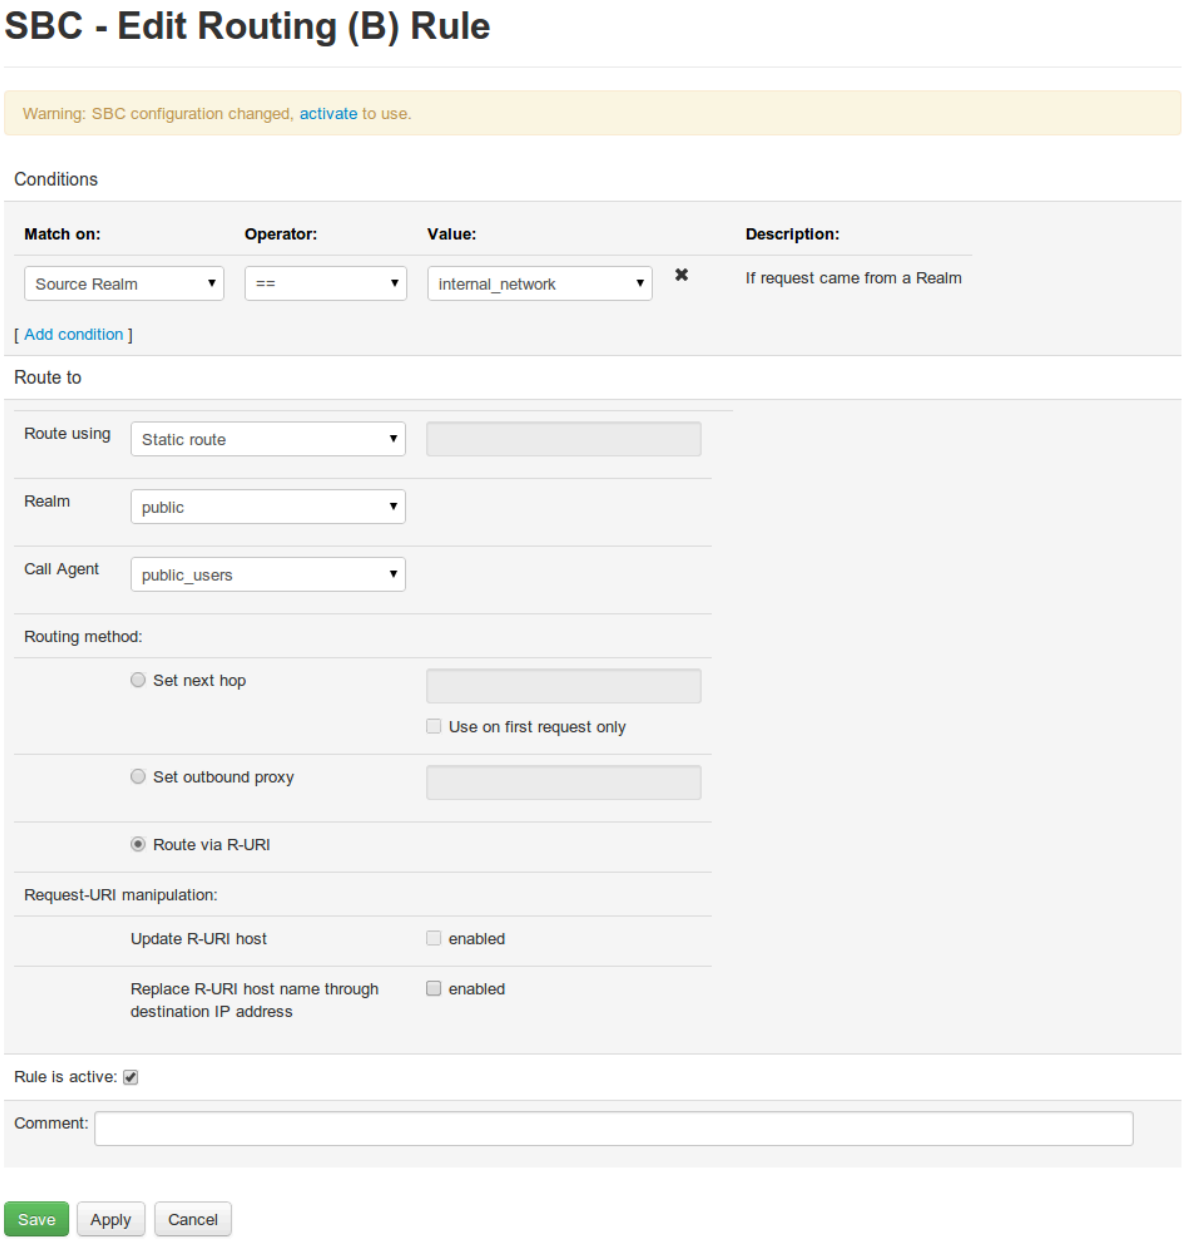

Figure 21: Rule B Definition for internal-network Realm

Configuring NAT Handling and Media Anchoring

We have already used the NAT option in the **Retarget R-URI from cache (alias)** action above. In order to route in-dialog requests to the caller properly even if the UA is behind NAT, we use the **Enable dialog NAT handling** action. This will make the SBC remember the source address of the caller for the dialog and use that to send in-dialog requests.

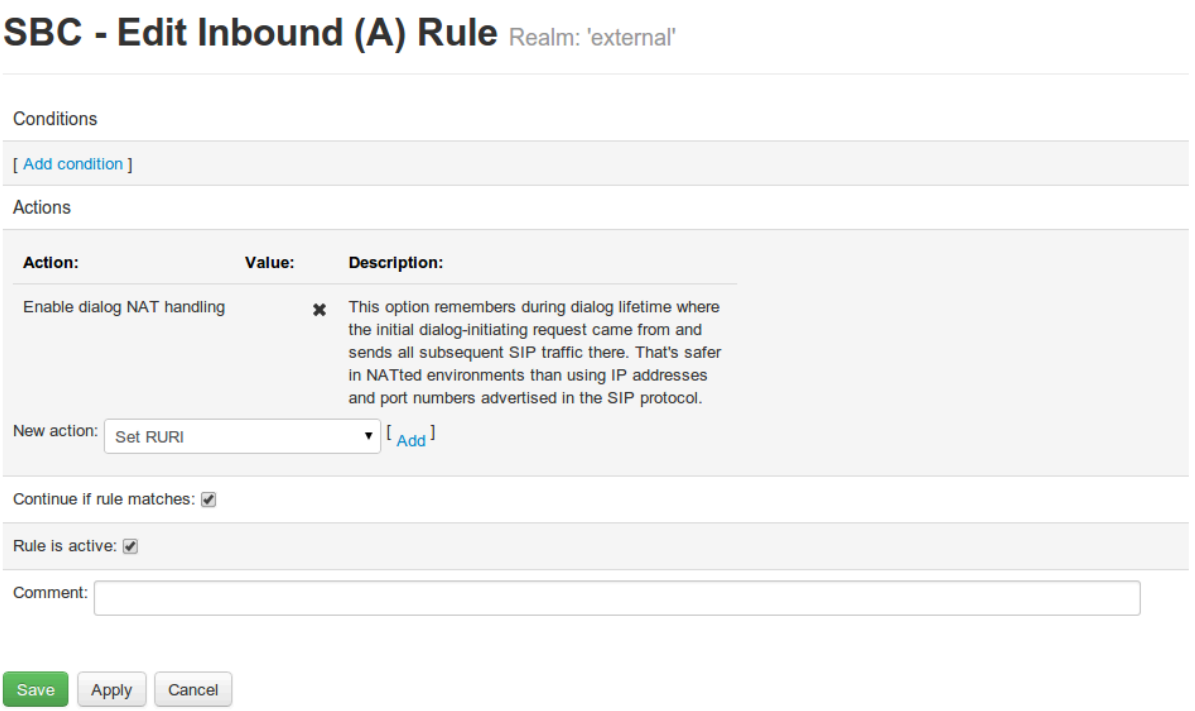

Figure 22: Rule A Definition for NAT handling

For the RTP to flow properly through different NATed users - and also from the internal network to the public network for calls to conference bridge server - we **Enable RTP anchoring** with the **Media far end NAT traversal for UAC** option enabled. To anchor the RTP of all calls at the SBC, we leave the **Enable intelligent relay** option unchecked; if we want to reduce bandwidth consumption and latency (total mouth-to-ear delay), we can also enable the intelligent relay option if we are sure that no users are behind double NATs. We enable this for calls in both directions - from and to the UAs.

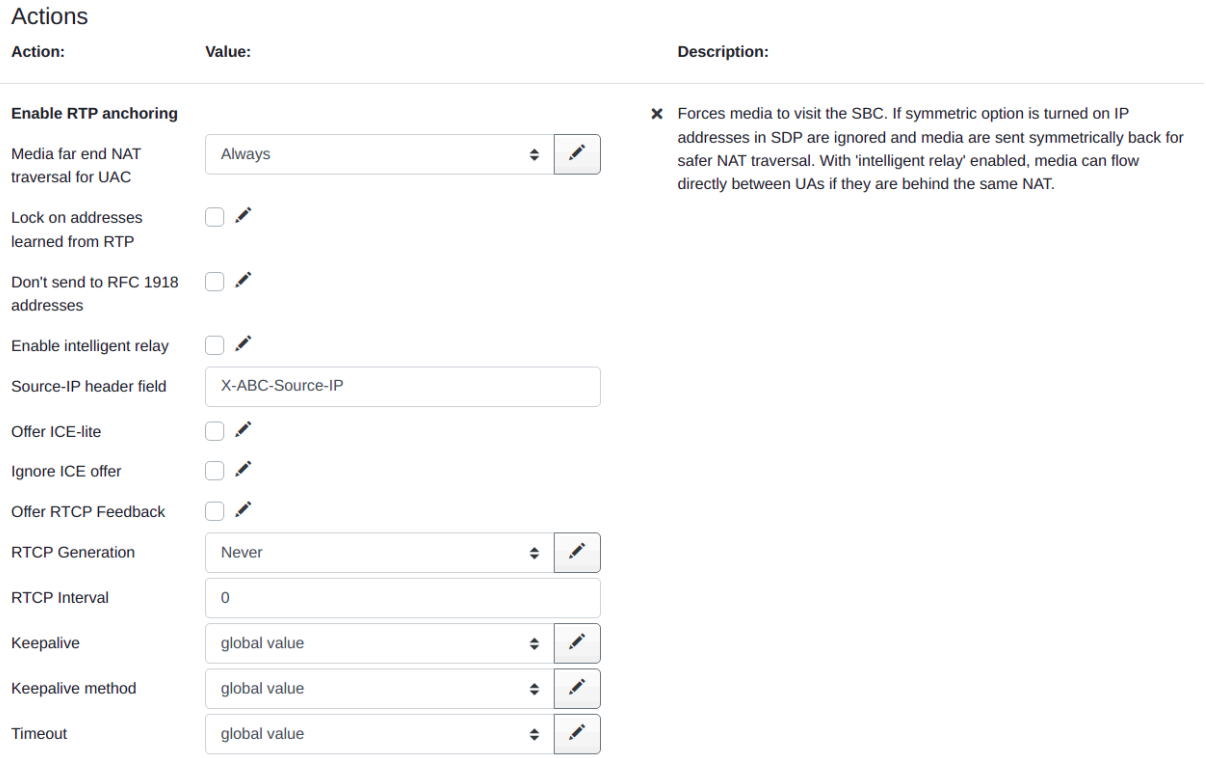

Figure 23: RTP Anchoring Rule Definition

Note well: it is important to realize that enabling **Media far end NAT traversal for UAC** will open a security weakness subjecting the call to a so called **RTP Bleed** attack. It can be mitigated partially by using the **Lock on addresses learned from RTP** option. Forcing usage of **Secured RTP** will effectively mitigate this attack as the SRTP packets will be authenticated prior to the address learning step.

Configuring transparent dialog IDs

If we want to enable call transfers through the SBC, and to simplify troubleshooting, we can **Enable transparent dialog IDs**.

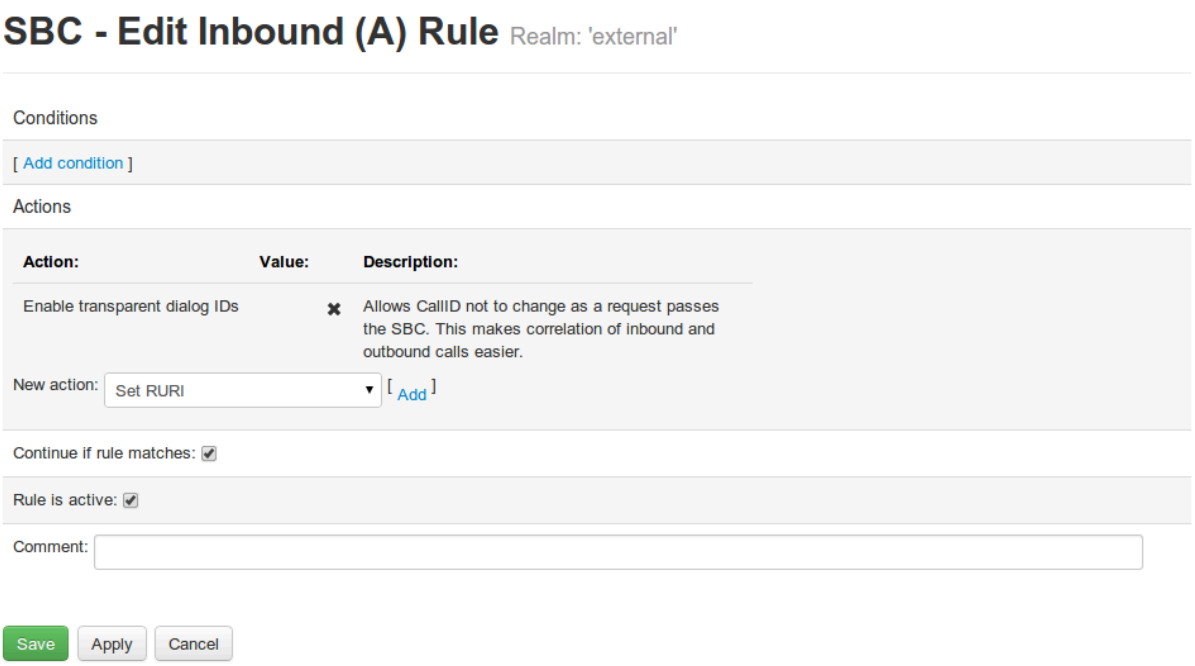

Figure 24: Transparent Dialog Rule Definition (A)

Setting up tracing

In the testing phase, we can enable tracing for calls with the **Log received traffic** action.

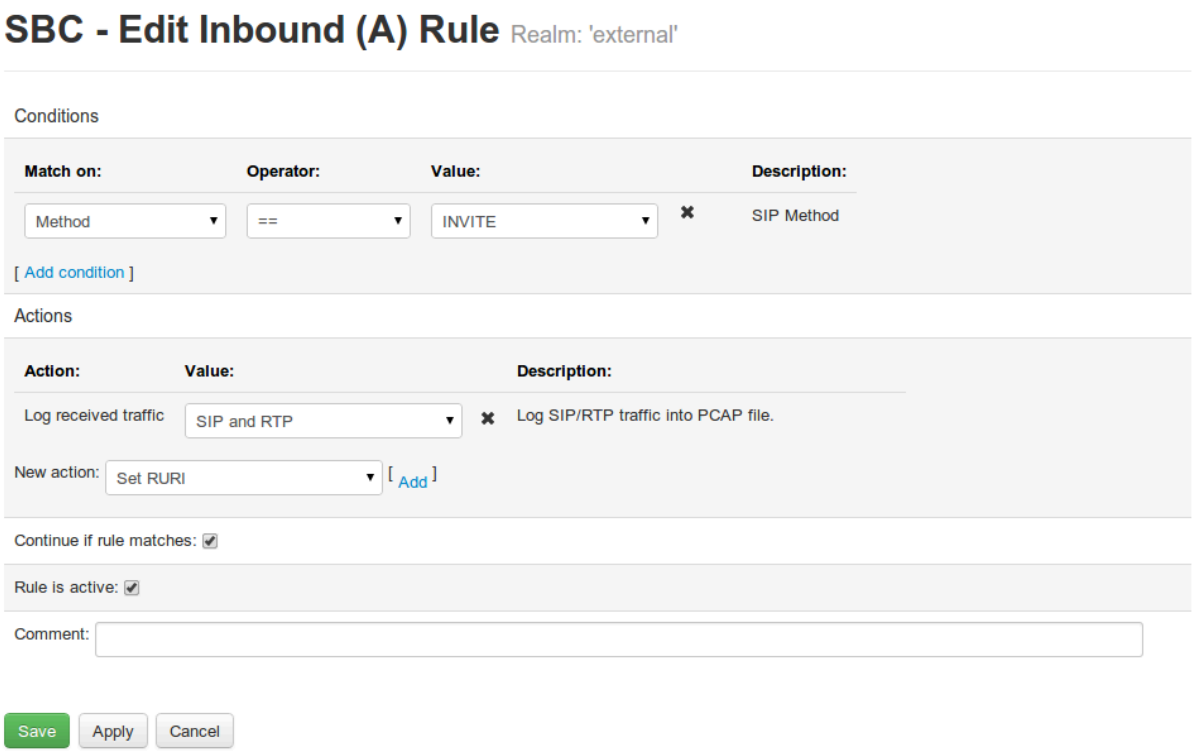

Figure 25: Tracing Rule Definition (A)

In production use we should not forget to disable or remove this rule to protect the privacy of the users and to reduce processing power and disk space requirements at the SBC host.

Summary of rules

The rules we have created so far can be seen in the Overview screen. The rules implement so far routing from the external to the private network and vice versa, recording traffic in PCAP files, NAT handling and registration caching and throttling.

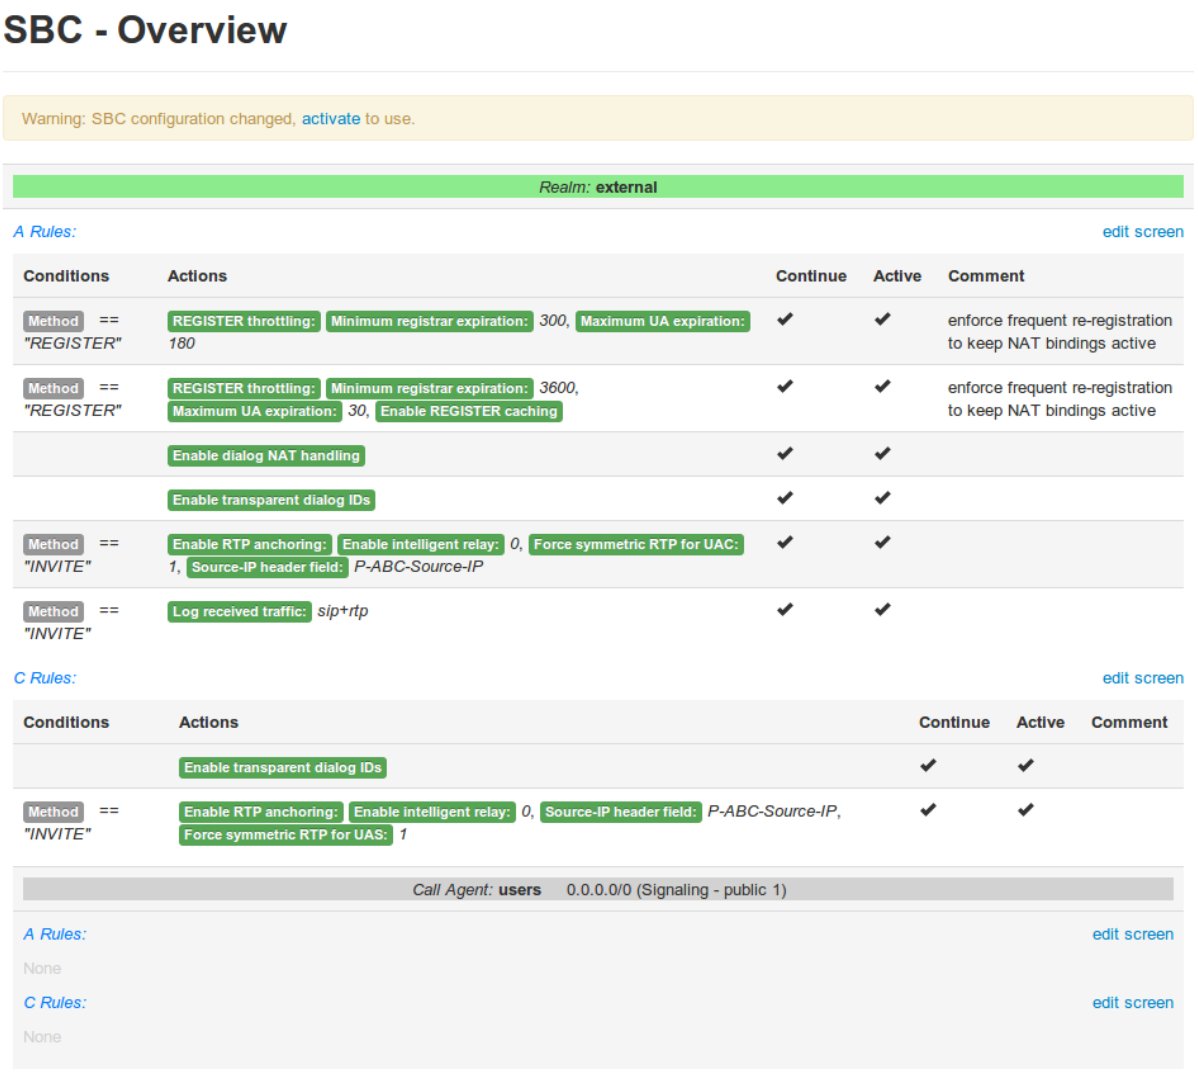

Figure 26: Rule list for sample network

Setting Call Limits

In order to reduce the risks of fraud, we can set some limits on traffic coming from the external network as shown in Figure Limiting calls and traffic:

- a parallel call limit of 10 for calls coming to the realm from the same source IP address (**$si**)
- a limit of 5 calls coming to the realm from the same user (**$fU**)
- a limit of call attempts per second (CAPS) of 10 for the calls coming to the realm from the same source IP address (**$si**)
- and a limit of 120 kbit/s for every single call coming to the realm - sufficient bandwidth for audio calls only. For video calls you might want to use a higher value.

For the limits per source IP address, it has to be noted that the limits may apply to a group of users if they are behind the same NAT. If for example there are enterprise users, we may group them into a separate Realm with a different, higher limit and/or group by a combination of IP address and domain name.

We set these limits in calls from the external realm.

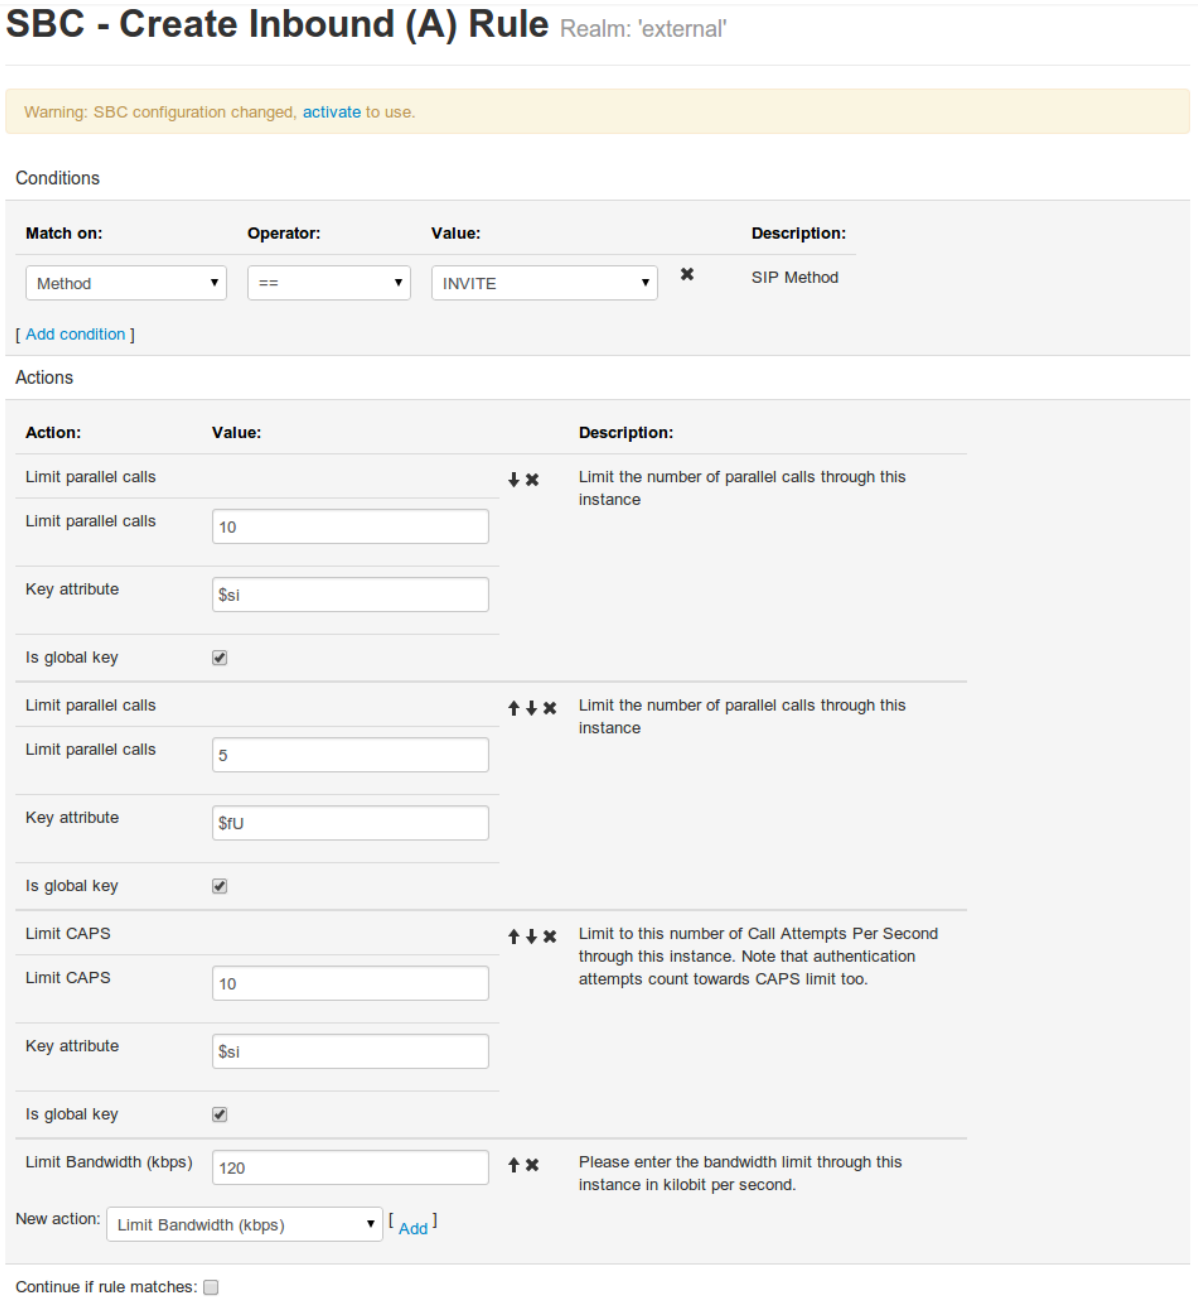

Figure 27: Limiting calls and traffic

Blacklisting specific IPs and User Agents

We can use a rule to block calls from a specific IP address.

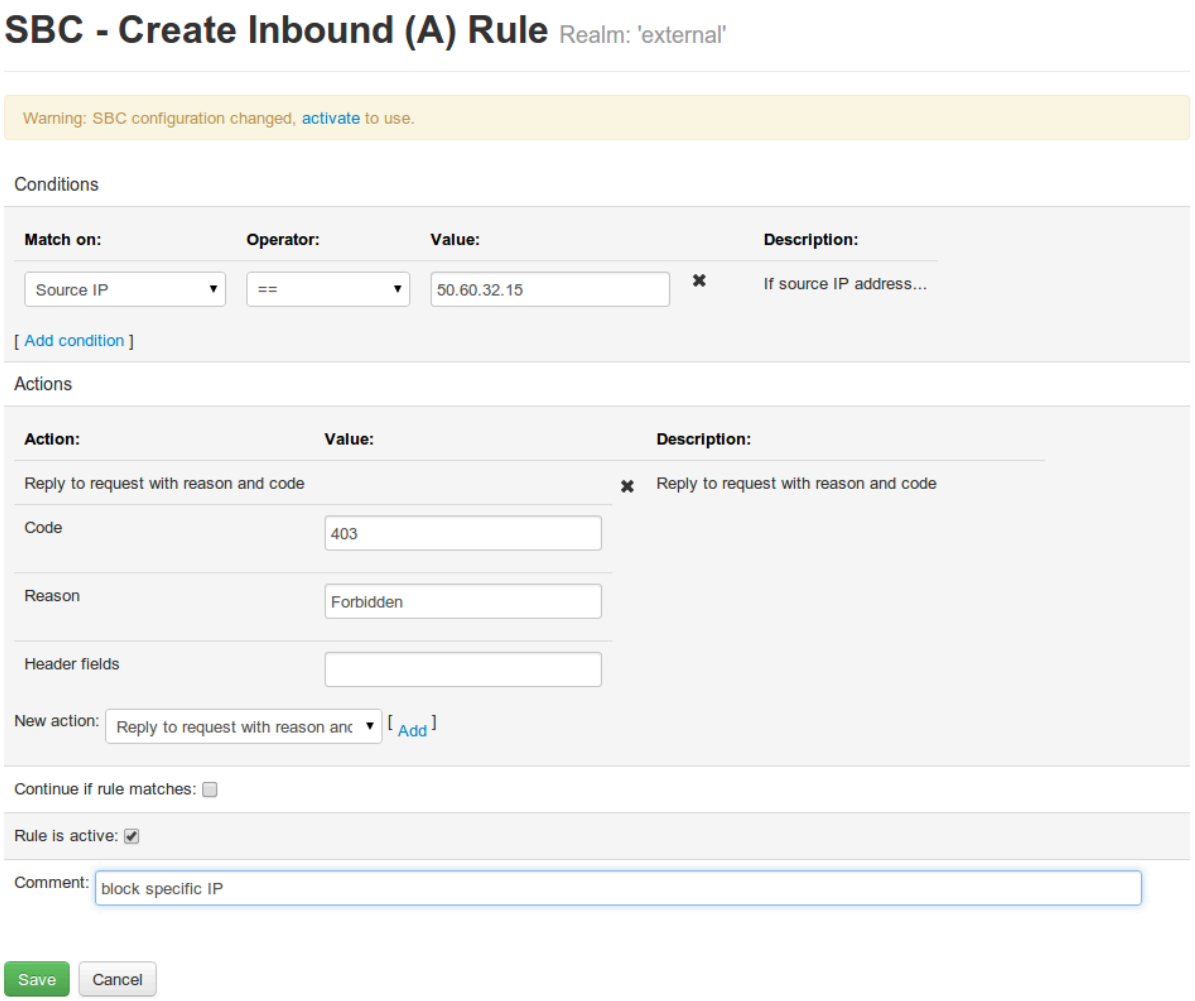

Figure 28: Blacklisting IP addresses

And also specific User Agent types, for example SIP scans from sipvicious which works if the User Agent header string is unchanged.

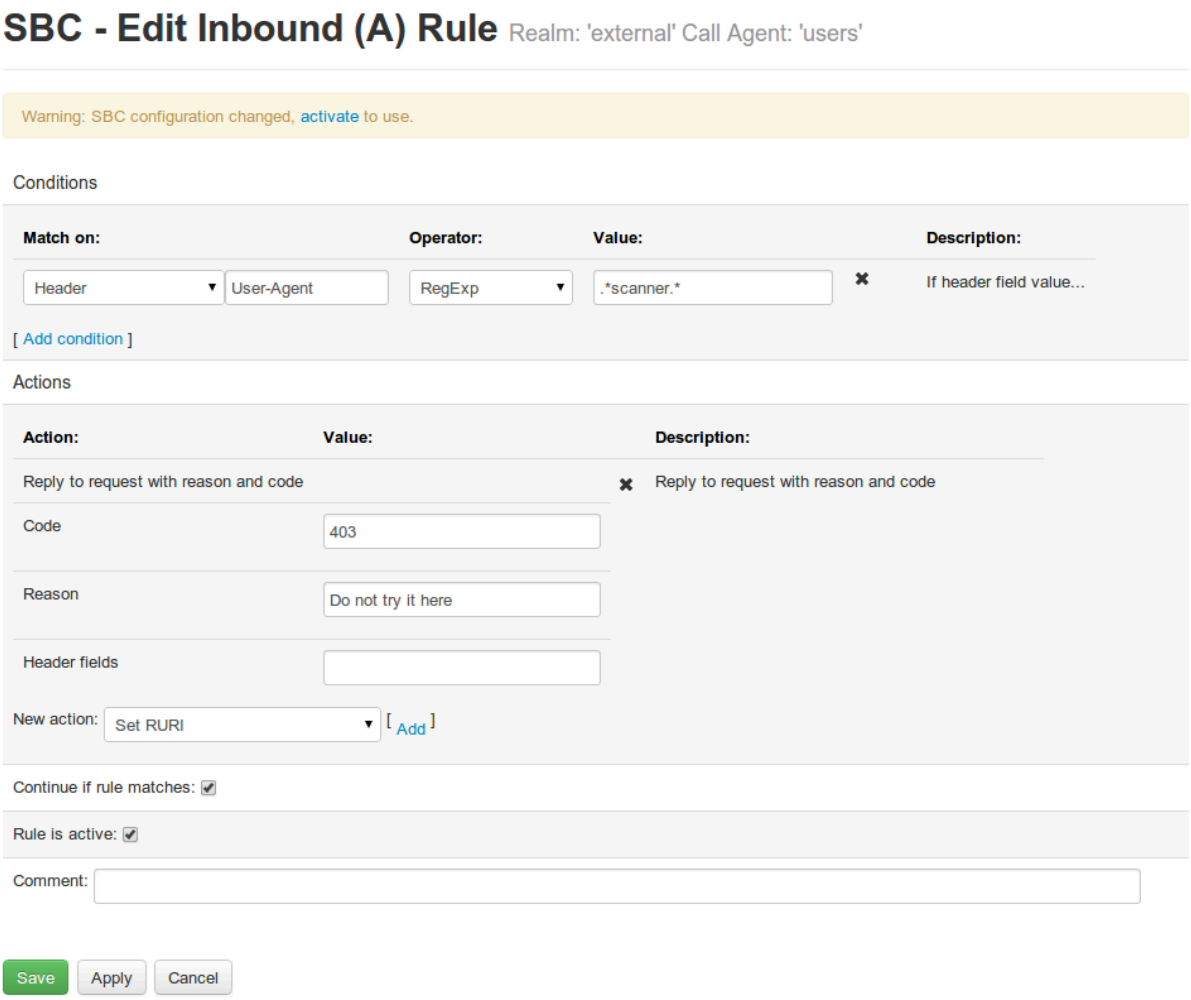

Figure 29: Rejecting calls from certain user agents

Handling P-Asserted-Identity

The *P-Asserted-Identity* header is usually used within a network to signal the caller, if the identity is asserted, e.g. if it is signaled from a trusted source.

The *P-Asserted-Identity* header should usually only be trusted if it was set by some element in the internal network, e.g. by the home proxy after authentication. Hence, for requests coming from an external network it is recommended to remove the *P-Asserted-Identity* header*.

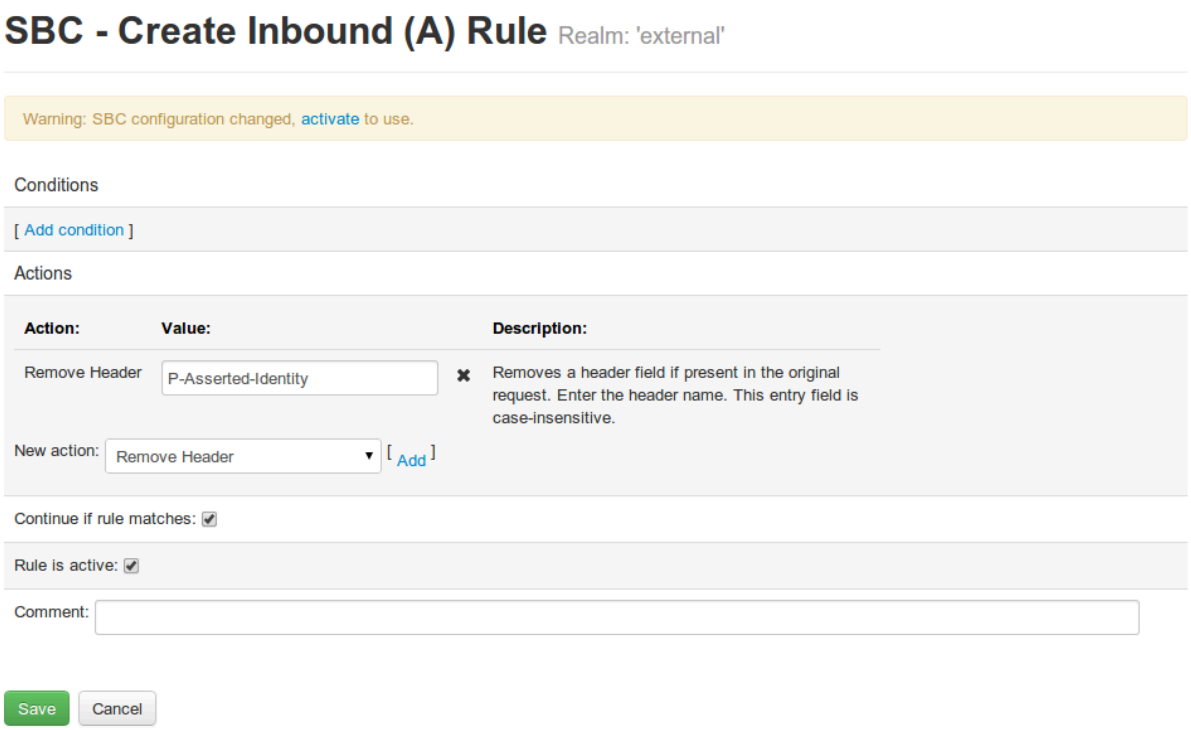

Figure 30: Remove P-Asserted-Identity header from untrusted requests

Where to go from here

This section described a typical initial configuration for a simple use case and a simple network topology.

Going further from here, various use cases that are solved with the ABC SBC are explained in various sections of this document:

- Interworking with various types of PBXs requires often very specific SIP mediation actions which can be implemented using special rule sets, see Defining Rules and SIP Mediation.
- Quality of calls with the Enterprise trunking use case can be improved by using intelligent RTP relay handling, see Media Handling for more details.
- Mobile clients may benefit from specific codec handling and transcoding. See Media Handling for more details.
- For more security mechanisms, refer to chapter Securing SIP Networks using ABC SBC and ABC Monitor (optional).
- Least cost routing can be implemented using Provisioned Tables. See Provisioned Tables for more details.
- For billing, the SBC can generate call data records (CDR). See Call Data Records (CDRs) for more details on how to use the CDRs and customize them.
- Both usage and the SBC host itself can be monitored through SNMP, see Measurements and Monitoring.
- System administration tasks like backup, maintenance and upgrades are explained in chapter ABC SBC System administration.

4.4 Initial Configuration

4.4.1 SBC Interfaces Overview

The ABC SBC uses five types of logical interfaces for management, signaling and media processing:

- IMI - Internal Management Interface - used for inter-node communication (in HA pair or between CCM and SBC node)
- SI - Signaling Interface - for SIP signaling (multiple SI interfaces can be configured)
- MI - Media Interface - for media (RTP/RTCP) processing (multiple MI interfaces can be configured)
- WS - Websocket Signaling - for SIP signaling over Websockets.
- CI - Custom Interface - for different applications specified by admin

Important: Before the following initial configuration, it is important to have all physical interfaces used by the SBC's logical interfaces configured and working (IP addresses and IP routing). Also, the hostnames of the machines have to be set, as they are used in initialization scripts and for distinguishing SBC nodes.

4.4.2 Web GUI Configuration (Cluster Config Master)

ABC SBC has two parts, installed as two separate containers: the configuration master aka Cluster Config Manager (CCM) which provides configuration GUI web interface, and one or more SBC nodes. Depending on the config synchronization mode (pull or push), the SBC nodes will either :

- automatically pull new configuration from the CCM
- receive a new configuration request from the CCM

The centrally configured configuration elements include ABC rules, Interfaces, Global Config Provisioned Tables, Realms, Call Agents, SNMP configuration data and Firewall Rules.

The CCM will ask to create username and password for configuration GUI admin access on first login into GUI. New GUI user will be created and added to "SBCadmins" GUI group. Minimum password length is 8 characters and it must not contain spaces.

It will also ask for setting username and password that will be used to authenticate the nodes to the configuration master when performing configuration pull or node status push. Note that this is different username / password than for the GUI access. The password must not contain spaces.

All ABC SBC nodes need to know which server is CCM in order to pull automatically new configuration from it. The recommended way to provide this information is to set environment variables on the ABC SBC container at startup — typically `MASTER` (CCM address), `CONFIG_USER` / `CONFIG_PASS` (the pullconf credentials configured on the CCM in the previous step), and optionally `CONFIG_MODE`, `NODE_UUID`, and TLS-related variables. The full list and their semantics are documented in Configuring the SBC container. With these variables in place the ABC SBC container comes up unattended and immediately pulls (or waits for) its configuration from the Cluster Config Manager.

By default the container comes with no root user password set, which allows access to the container shell only directly from the host system. SSH is also disabled by default. See the installation guide for details on container access.

> **Info**
>
> In earlier releases, ABC SBC nodes had to be bootstrapped by running the interactive `sbc-init-config` script from inside the container. That path is no longer the preferred method; use the environment variables described in Configuring the SBC container instead. `sbc-init-config` is still shipped as an alternative and prompts for the same settings (CCM address, administrative domain, node UUID, client certificate and key, CA certificate, synchronization mode, pullconf username/password).

Configuration synchronization in pull mode

If `pull` was selected for the SBC node configuration synchronization, the SBC node should try to pull from the CCM new configuration every 15 seconds. Please note, the configuration needs to be Activated first before being pulled by the node.

The process is handled by the `sbc-pullconf` service.

Note that it's the `sbc-status-checker` service that reports the configuration synchronization mode to the CCM.

Configuration synchronization in push mode

If `push` was selected for the SBC node configuration synchronization, it is up to admin to "manually push" a new configuration once it is activated:

- edit the configuration via the CCM web interface
- click the "Activate" button
- once the configuration is activated, the GUI redirects to the config push screen. This screen is also accessible from System –> Config push
- select one or more nodes to which you would like to push a new configuration, and click on "Push to selected" button
- if a new configuration should be pushed to all SBC servers then click the "Push to all" button

The process is handled by the `sbc-gopi` service.

Note that it's the `status-checker` service that reports the configuration synchronization mode to the CCM.

In case mutual TLS authentication shall be used, the *CCM pullconf certificate* must contain the *TLS Web Client Authentication* extended key usage attribute. It is also advisable to include the *TLS Web Server Authentication* extended key usage attribute.

4.5 Setting Up Web Interface Access and User Accounts

ABC SBC web interface is available at the IP address of CCM (config master) interface and can be accessed using https URL on port 443 like this: https://192.168.178.178/

For the configuration of ABC SBC please access the IP address of the CCM node. The ABC SBC GUI uses local browser's time to display all times and timestamps.

Further information about managing administrative users can be found in the Section User Management.

When user login attempt fails several times, the user account is locked for certain time period. For details please check Login Parameters. To unlock the account just wait for the configured *Blocking period* or use following CLI command from command line:

```
sbc-user-passwd -u <USERNAME>
```

4.5.1 Default User Accounts

The initial username and password for user with admin rights for GUI access is created on first CCM GUI login. Then the GUI users can be managed via GUI - new users can be added or assigned to groups, as described in the section User Management.

Group membership defines privileges of the respective users. The following groups come preconfigured:

- **ABCMonitorUsers** Access to ABC Monitor
- **SBCadmins** SBC administrators having access to all configuration
- **SBCrevisor** Read-only access to everything
- **SBCrest** Access to REST API interface. **Note:** using this group standalone is useless. You should use it together with other group specifying which REST API resources the user have access to.
- **SBCrpc** Access to XML-RPC interface. **Note:** using this group standalone is useless. You should use it together with other group specifying which XML-RPC resources the user have access to.
- **SBCusers** access to SBC related configuration (no rights to system configuration - networking, users, firewall etc.)

4.6 ABC SBC License

By default, the FRAFOS ABC SBC is installed in a demo version, which is limited to 90 seconds call duration, does not include support for replication, high availability and extension packages. Enabling these features requires a license file. FRAFOS issues license files according to the agreement between FRAFOS and the customer. The license file enables features as shown in the table bellow:

Licensing Package	*Feature*
transcoding	action: “Activate Transcoding”
recording	action: “Activate Audio Recording”
RTC	interface: “websocket”
media server	action: “refuse call with audio prompt”
high-availability	background active/standby replication
monitoring_enabled	gathering monitoring information for use in ABC Monitor

In the demo version without proper license set, the respective features are not executed. When the number of maximum calls is reached, the ABC SBC returns a SIP response “503 Server overload” and, if monitoring is enabled, issues a “limit” event with reason “licensed session limit reached”. When the maximum duration is reached, the server terminates the call by sending BYE request to both parties, and if monitoring is enabled, issues a “call-end” event with originator field set to “internal-disconnect”.

The license file has to be imported to the SBC using the ‘System ▸ License’ link. Using the ‘Insert new license’ button, the administrator should give a name and selects the proper license file from the local disk by clicking ‘Browse’ button. After applying the changes, the license file is automatically uploaded to the server and loaded.

On Amazon Web Services, the paid-AMI instances download their license files when they start and no additional license configuration is required.

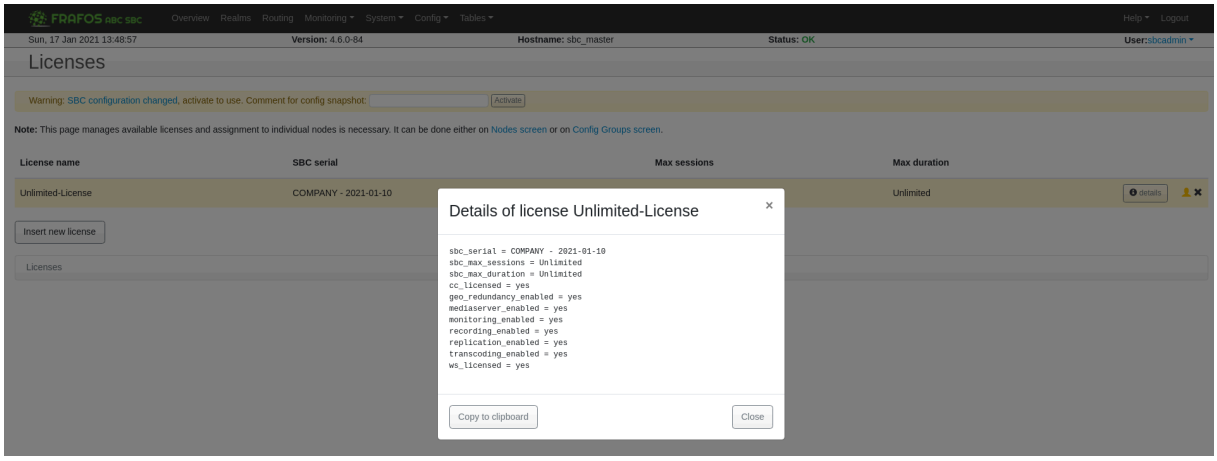

Important: For HA or a cluster deployment, the license file has to be imported on all nodes.

4.7 Interface Configuration

- Physical and System Interfaces
- SBC Interfaces
- Retro Compatibility

4.7.1 Physical and System Interfaces

System (network) interfaces inside the container can be seen either using the same names as on host, or using different name, depending on host or macvlan network mode used. If host mode is used, the interface cannot be configured inside the container, and uses IP address as configured on host. If the macvlan mode is used, the interface has to be configured inside the container, which can be done by adding a network interface configuration file in /data/interfaces.d/ directory. The DNS server address can be also configured there. See “man interfaces” for format details.

Several types, like simple system network interfaces (e.g. eth1), VLAN tagged interfaces (e.g. eth1.100) or bonded interfaces (e.g. bond0) can be configured and used in the ABC SBC configuration.

SBC nodes

If ABC SBC is installed in HA (active-standby) or cluster mode, the main configuration node should know about all the SBC nodes. This is required specifically in case the SBC interfaces settings differ between the nodes - e.g. when the nodes differ in IP address or system interface name used for one interface of the same logical type.

By default, each node has unique node UUID created locally at first container start, and on configuration master side the node records are added automatically when the nodes pull configuration for the first time. The automatic adding of node records can be disabled under “Config ▸ Global Config ▸ Misc ▸ Automatically add new nodes”.

In case it is needed to add node records manually, either because automatic adding of node records is disabled or the records are needed to complete configuration even before the nodes try to pull configuration for first time, it can be done on the “System ▸ Nodes” GUI screen of the main configuration master node.

For each SBC node, you have to enter it’s node name and node UUID. The node name field is just informational, and e.g. node hostname can be placed there. The node UUID is either generated on the nodes on first start, or if specific node UUID is required it can be entered manually when doing the initial node configuration. The node UUID is used to match the node to node specific settings.

Since ABC SBC release 5.5, a node can have one of two roles: either the standard SBC role or the probe role. If the SBC role is selected, the deployment type can be set to either “standard” or “aws”. For typical installations, use the default “standard” value. Use “aws” for deployments on Amazon AWS, which is intended for using HA under AWS. For more information about probe see Probe management.

Configuring Virtual IP (VIP) Address (OPTIONAL: in HA mode only)

When deployed in an HA active/standby mode two instances of the ABC SBC nodes will share one or more Virtual IP addresses. Virtual IP addresses are assigned to the currently active node.

The HA is configured using the “System ▸ HA” screen.

For each pair of HA nodes a “HA group” can be created and the nodes assigned to it using “System ▸ Nodes”, or the HA group can be created also directly on the Nodes page while adding a new node. This HA group says which nodes will share the HA VIP IP addresses.

Note that the ABC SBC containers for the two HA nodes should run on two separate host servers, not on the same host. (Running both ABC SBC nodes on one host may bring unexpected problems related to network interface link being down, if the HA macvlan mode is used.)

It is mandatory for the nodes in HA group to have IMI interface defined, and they must not use “IP autoconfig” option on the IMI interface, as the two nodes in HA group use the IMI interface IP address for HA “heartbeat” between the two nodes.

Under the HA group, you can add one or more VIP - Virtual IP addresses. For the VIP enter the IP address and optionally (recommended) also a netmask of the IP address. The netmask has to be in CIDR notation (like “24”) or subnet mask (like 255.255.255.0). If netmask is empty, a mask “32” (meaning single host) will be used.

Optionally, also one or more HA routes can be added, which are bound to a particular VIP address. Such routing rules will be brought up and down together with the VIP address. For the HA route, the following data can be entered: Route destination - in form of subnet/netmask, like 192.168.0.0/24, this field is mandatory. Other fields are optional: Gateway - the routing gateway IP address, Source - the source address to prefer when sending, Table - table id if policy based routing is used.

The “VRRP router ID” option of the HA group specifies the VRRP protocol’s virtual router identifier, which is used in VRRP advertisement packets to distinguish multiple HA pairs when sharing the same network. It accepts values 1-255 and a random value is pre-filled when adding a new HA group. In case there would be conflict of the ID with some other equipment that also uses the VRRP protocol on the same network, the ID can be edited. Note: for ABC SBC releases up to 5.4, a hardcoded ID value of 66 was used.

Note: if the global config option “Use also signaling interfaces for vrrp adverts” is enabled, then a range of IDs is used, starting with the value configured under the HA group, to use unique value for each interface.

Optionally, also a “gateway heartbeating” can be enabled and configured under the HA group properties. When enabled, the gateway reachability will be periodically checked using “arping” command, and possibly a HA switchover will be initiated if gateway becomes unreachable on one of HA nodes. The options to configure this are:

- Gateway address: the IP address of the gateway to check using “arping”
- System interface: Name of network device where to send ARP REQUEST packets. Needs to be set only if the node cannot find interface to use based on system routing, and if used it implies also using source address 0.0.0.0 for the requests, to be able to ping gateway even if Sbc node has no other IP address than the VIP one on the subnet towards gateway. Please use only if needed specifically, in usual cases leave empty.
- Number of pings to fail: sets after how many failed pings the gateway will be considered being unreachable

- Number of pings to succeed: sets how many pings have to pass to consider the gateway reachable again
- Ping interval: sets ping interval in seconds
- Ping timeout: sets timeout in seconds, how long to wait for ping answer. (Typically, it takes a bit more than one second to detect unreachable address, so recommended default value for timeout is 2 seconds.)
- Weight to increase/decrease by: sets weight to modify HA node VRRP priority (default 250) if gateway is reachable, to make node with higher priority become new HA master. Using value of 0 for the Weight will make the node go into HA FAULT state if gateway is unreachable, instead of just modifying weight - which is not recommended, as can lead to both nodes going into FAULT state even if the gateway check would return false negative result. When modifying this value, please consider also the global config option for the HA priority, as the sum of priority and the increase value must not get above 254 or below 1.

Once the VIP address(es) are defined, it is possible to select to use VIP and choose particular VIP address when configuring ABC SBC interfaces using "System ▸ Interfaces" screen. The VIP address can assigned to the SBC signaling, websocket or media type of interfaces.

Note: in case a setup is reconfigured to remove HA group and move SBC nodes to use normal IP address instead of VIP, it may be needed to perform additional config activation for the node which was previously acting as BACKUP in HA, as the signaling process will come up only on a node that was previously acting as MASTER in HA setup.

Before proceeding with the CCM configuration, please ensure that the following requirements are met for the High Availability (HA) setup on the AWS platform:

- Under Global Config, AWS access keys must be configured as SBC requires an access to configure VIPs on AWS platform.
- One private IP should be available in SBC interface (VPC) network, this IP address should not be assigned to any existing interface as it will be used as the secondary IP (VIP)
- If the interface requires to communicate to public area, then an Elastic IP should be created but not allocated to any interface.
- The Elastic IP must be set also in the Public IP address field within the CCM interface configuration.
- On AWS side, metadata service of SBC instance must be enabled and IMDSv2 must be marked as required.

4.7.2 SBC Interfaces

For signaling and management the ABC SBC uses several types of "logical" interfaces:

- **IMI - Internal Management Interface** – The IMI is used for inter-node communication (in HA pair or between CCM and Sbc node) and for configuration transfer from configuration master to ABC SBC node(s). Only one IMI can be configured. Separate system interface using IP subnet not routed or accessible from outside should be used for IMI, unless there is a external firewall in front of ABC SBC. The port number accessible on IMI for the config pull from configuration master is 444. It is mandatory to create the IMI interface.

  > **Info**
  >
  > There are also several services providing API on Sbc side on IMI interface, to which the CCM node connects for getting local monitoring, webconference and logs data. The access to these API ports is limited to CCM node src IP address by Sbc firewall. It is important that there is no NAT involved on traffic between the CCM and Sbc nodes.

- **SI - Signaling Interface** – SI is used for SIP signaling. Multiple SI can be configured.
- **MI - Media Interface** – MI is used for media (RTP, UDPTL, ..) processing and relay. Multiple MI can be configured.

- **WS - Websocket Signaling** – WS is used for SIP signaling over Websockets. This is useful only if the ABC SBC is configured to act as RTC gateway as described in Section SIP-WebRTC Gateway.
- **CI - Custom Interface** – CI is used for different applications which can be used for specific purposes like SSH, Prometheus pull service, TURN, HTTP proxy and HTTP redirect.

Signaling and media interfaces can be configured in different combinations. All SI/MI can share the same system interface, can be configured on a "per Call Agent" basis where each Realm has its signaling and media interface, or can share one assigned IP address with different ports per SBC interface.

It is also possible to create separate signaling and media interfaces on the same system interface for different purposes. For example, one for a PSTN gateway and one for receiving calls from residential users. In this case, a different signaling port and media port range shall be used. A typical ABC SBC configuration is to have one separate IMI and one shared signaling and media IP address for each Realm.

When doing the initial ABC SBC configuration, add IMI interface. The IMI interface has to be defined always if HA or cluster mode is used (otherwise needed firewall rules would not be set).

Then add the interfaces for the SBC application: signaling (SI) and media (MI), optionally websockets signaling (WS).

If a specific application is needed, custom interface (CI) can be used with any port requested by admin.

When adding logical SBC interface, you first define it's name and options that are common to all SBC nodes using this interface, then you add records under the logical interface which map it to system interface for node(s) that will be using this logical interface. The list of records that map logical SBC interface to system interface on node(s) can be expanded or collapsed using the "+" or "-" icon before interface name. New mapping of logical SBC interface to system interface can be added using the "insert new system interface" button located at left hand side of the list.

In HA or cluster mode, if the interfaces differ between the nodes (use different IP address or system interface name), you have to create more separate logical to system interface mapping entries under the logical interface. Create a separate entry for each SBC node, set owner type to Node and select the node under Owner. Note: if records both for all nodes under a config group (owner type of config group selected) and for specific nodes are created, each node will use the record for all only if specific record for that particular node does not exist.

If the SBC interface settings do not differ between nodes, you can create just one logical to system interface mapping entry under each logical interface, set Owner type to config group and use the "default" config group.

If SBC interface is going to use VIP address (shared IP), the VIP address should be added before adding the interface.

SBC Interfaces are configured in the "System ▸ Interfaces" screen.

The following parameters can be defined for logical SBC interface:

- Interface name: a unique identifier of the logical interface - [a-z, A-Z, 0-9].
- Interface type: Signaling, Media, WebSocket Signaling, External management, Internal management, Custom.
- Interface description: description (alias) for the interface that is used in the GUI configuration.
- TLS profile. By default, the TLS profile is set to None, meaning no TLS will be used on the interface. If TLS is to be used, select the TLS profile to use on the interface. The TLS profiles can be edited under System / TLS profiles page. There is profile named "default" which is automatically created at ABC SBC installation and uses self-signed certificate.
- Applications (Apps): each logical interface can have one or more "Apps" enabled, which tune what service on which port will be listening on that interface, plus allow setting more specific option.

Please refer to - Reference Application Interface Options section for the Apps options details.

After creating entry for the logical SBC interface, add at least one logical to system interface mapping under it. The following parameters can be set for the mapping:

- Owner and Owner type: These list-boxes options set to which specific node the mapping of SBC logical interface applies. It can be assigned to a particular node as been pre-configured under **"System ▸ Nodes"**, or all nodes belonging to a particular config group ("default" by default). Note: currently SBC supports only one common config group named "default", which can be used if the mapping applies to all nodes.
- System interface: system interface name (eth1, eth1.123 - VLAN tagged, bond1 - bonded interface)
- Type of IP address: use "manual" to manually specify the IP address, which is the default. If "auto-config" is used, the first IP address from the corresponding system interface will be taken automatically. Use "VIP" to select one of VIP addresses, which can be configured in case of HA deployment mode under "System ▸ HA" screen. Note: when configuring IMI interface of a node belonging to a HA group, the IP address type has to be set to manual.
- IP address: you can specify the IP address of the interface.
- Type of public IP address: use "manual" to manually enter the public IP address in the following field, which is the default. Use "Amazon autoconfig" to autodetected the public IP address. Current options of autodetection include Amazon EC2 cluster method.
- Public IP address: this parameter is optional. It allows to configure an IP address that will be used instead of the real or virtual IP address in SIP signaling (in case of the signaling interface) or media description (SDP; in case of a media interface). This is very useful to support near end NATs, e.g. Amazon EC2. Please refer to Sec. Physical, System and SBC Interfaces more details on the topic.
- TLS profile. If any value is set there, it override the TLS profile value set for the logical SBC interface. Otherwise TLS profile set on logical SBC interface is used.

The fields: `System interface`, `IP address`, `Public IP address` and `TLS profile` supports cluster config parameters (values in format "`%param_name%`") so even single logical to system interface mapping record may result into different IP address or system interface used on different nodes.

Important: When the SBC interfaces are configured, a warning message with a button to activate the new SBC configuration is shown in the GUI. No SBC interface changes are applied until the "activate" button is used. When the configuration changes are applied, all services using network configuration are restarted (e.g. SIP and RTP processes, etc..). Note that this may cause service disruption.

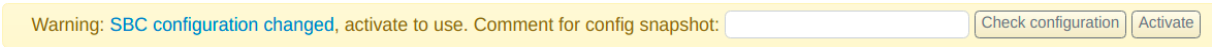

4.7.3 Retro Compatibility

Retro compatibility was introduced with the ABC SBC 4.5 because of increment of the JSON config version from `1.0` to `1.1`. The major change was the addition of interface applications, allowing a better transparency and tweaking of what is running where.

As the change brings some inconsistencies between the two config versions, a `retro-compatibility` module was developed. The purpose of that module is to, for every new config version to be deployed, check the current config (version 1.1) against the target node release, and when needed convert the JSON config to the older format.

While some basic changes are made under the hood (converting application interface options to older global config options - sshd port value for example), more complex changes are reported as an error.

The following table maps possible error situations with suggested solutions.

Common issues and fixes

Error message	Cause	Fix
`json retro-compatibility 1.1 to 1.0: [APP NAME] app not supported on older setup`	The application is enabled on an interface assigned to an node which does not support it (< ABC SBC 4.5)	See the list of unsupported applications in the section Applications. Disable the problematic application.
`json retro-compatibility 1.1 to 1.0: custom interface isn't retro compatible`	Custom interface is a new type of logical interface which was introduced in ABC SBC 4.5. This interface is not backward compatible.	Unlink the custom interface from the node or node's config group.
`json retro-compatibility 1.1 to 1.0: different value provided for the [PARAM] (app [APP_NAME]) Imported values: [VALUES]`	Different values assigned for options on a pre 4.5 node. Example: sshd port 22 on IMI, 23 on XMI	Use the same value for every application assigned to the faulty node. Example: set 23 for both IMI and XMI

Applications

Name from error	Description	Interface	GUI name
`pkapman`	Service which generates and serves pcap files on SBC. Available since: 4.3. Replaced by `gopi` since: 5.5.	Internal management interface (IMI)	`PCAP query service`
`gominstrator`	Perform administrator actions on a host. Please note this does not affect xmloredis service in pre 4.5 releases. Available since: 4.5. Replaced by `gopi` since: 5.5.	Internal management interface (IMI)	`Management for host`
`webconf-api`	Expose sems's webconf mgmt via RESTful json API. Available since: 4.6. Replaced by `gopi` since: 5.5.	Internal management interface (IMI)	`Local webconf API`
`turn`	Enable COTURN. Available only on 4.5 to 5.1.	Custom interface	`TURN server for websocket`
`http_proxy`	Allow custom nginx proxy. Available since: 4.6.	Custom interface	`HTTP proxy`
`http_redirect`	Allow custom nginx redirect. Available since: 4.6.	Custom interface	`HTTP redirect`
`goplog`	Provides access to logs on SBC node. Available since: 5.0. Replaced by `gopi` since: 5.5.	Internal management interface (IMI)	`Access to log files`
`conference_gui`	Simple Web GUI for Meet-me conference. Available since: 5.1.	Custom interface	`Simple Web GUI for Meet-me conference`

Name from error	Description	Interface	GUI name
`gopacla`	Provides firewall interaction for the node. Available since: 5.4. Replaced by `gopi` since: 5.5.	Internal management interface (IMI)	`Packet classifier API`
`gopi`	Provides various interaction with the node. Available since: 5.5.	Internal management interface (IMI)	`Unified SBC management server`

SBC 5.0 introduces possibility to hide GUI options which are not present / compatible with the selected version. This can be found in CCM –> CCM Config –> Misc. The default value is "None" which means everything is visible.

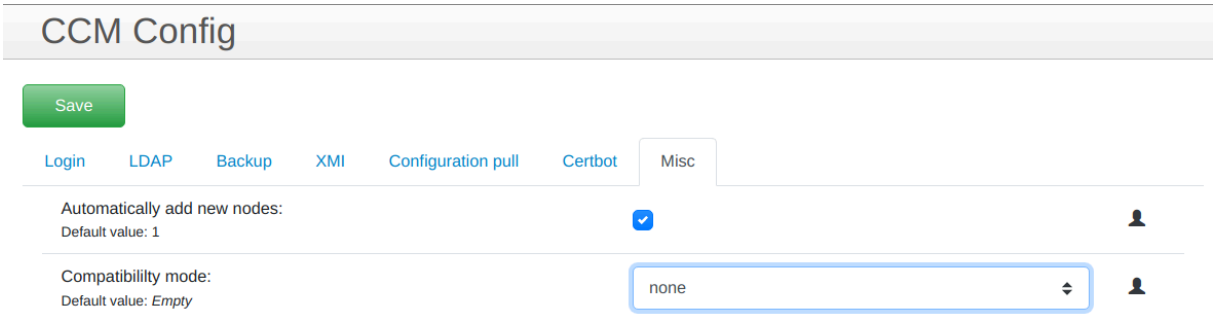

Figure 31: Retro compatibility mode selection

4.8 TLS profiles Configuration

The TLS profiles screen is used for manage TLS profiles used by SBC interfaces (see section Physical and System Interfaces). Each SBC interface can be configured to use different TLS profile.

The TLS profile takes effect only on those "Apps" enabled on the corresponding SBC interface, which support using of TLS.

4.8.1 TLS profile options

Name	Human friendly name, used for logging and others.
SSL certificate file	Select a file containing SSL certificate in PEM or PKCS#12 format Please note that, CA and possible intermediate CAs should be contained in the PEM file if all should be presented.
SSL private key file	Select a file containing key for SSL certificate in PEM or PKCS#12 format
Trusted CA certificates file	Select a file containing of trusted CAs in PEM or PKCS#12 format If provided, all certificates presented to the SBC will be checked against it, regardless of whether "Mandate peer certificate:" is set.
Mandate peer certificate	If checked, a peer certificate must be presented and will be checked against the trusted CA file.

Verify client hostname/ip	If checked, the SBC verifies whether hostname / IP address of the client match the one mentioned in the peer certificate. This applies only to Client signaling connections to the SBC (where the SBC is server).
Disable server hostname/ip verification	As of now, disable of verification of server hostname or IP works for signalling application only.
Allow wildcard certificate	If checked, the SBC accept wildcard certificates. Warning: this option shall not be used in most cases. You enable it on your own risk. This option takes effect for signalling application only.
Enable Let's Encrypt	If checked, no certificate, private key nor CA certificates are required. The ABC SBC will handle by himslef the completion of either an ACME HTTP01 or a DNS01 challenge against Let's Encrypt certificate authority. Refer to Let's encrypt gocertbot for more information about the requirements.
DNS	DNS domain associated to the node. The DNS is used to complete the ACME challenge on the let's encrypt side. Require **if** `Enable Let's Encrypt` is checked.
Challenge Type	Type of Let's Encrypt certificate authority challenge. Possible values are `http01` or `dns01`. Refer to the official home page or Let's encrypt gocertbot for more information. Require **if** `Enable Let's Encrypt` is checked.
DNS Provider	DNS provider furnishing the node's DNS. Require **if** `dns01` is selected.
Challenge Options	Set of settings specific to the selected DNS provider. Refer to the supported provider list for more information. Example: the following was used to test against namecheap's sandbox platforme: `{ "NAMECHEAP_PROPAGATION_TIMEOUT": "600", "NAMECHEAP_API_USER":"QQQ", "NAMECHEAP_API_KEY":"XXX", "NAMECHEAP_SANDBOX":"true" }` Require **if** `dns01` is selected.

Table 1: TLS profile options

4.8.2 Certificate requirements

For the TLS certificates to be used with ABC SBC, the following requirements have to be met:

- The IP address or hostname for which the certificate is issued needs to to be listed in it's SAN (Subject Alternative Name) field.

If IP address is used for access to CCM, the SAN field format should be like "IP:1.2.3.4", or if DNS name is used then "DNS:test.example.com".

- The "serverAuth" should not be set in "extendedKeyUsage" field of the certificate for SBC node (client) side.

- If the certificate is not of a “wildcard” type and was issued only for one IP address, it has to be carefully considered to which Sbc node or group the TLS profile is assigned under Interfaces. It can be e.g. used for two Sbc nodes that are used as HA pair and the IP address is used as VIP address.

Note that if using the Let’s encrypt certificates together with http challenge, each certificate issued by LE is for a single unique IP address (aka a single node’ interface).

4.8.3 Let’s encrypt gocertbot

If the “Enable Let’s Encrypt” option is selected, a set of TLS certificate, private key and CA bundle will be automatically acquired and renewed against Let’s Encrypt certificate authority challenges services.

Renewal

The certificate renewal will be attempted automatically 15 days before it’s expiration. On certificate obtention or renewal, a notification email is sent to the administrator, using the email address set under ‘Global Config > System Monitoring’ and a new configuration has to be activated from the Cluster Config Manager.

Settings example

We start by creating a dedicated TLS profile. Depending on the challenge type, we end with a configuration similar to one of those two :

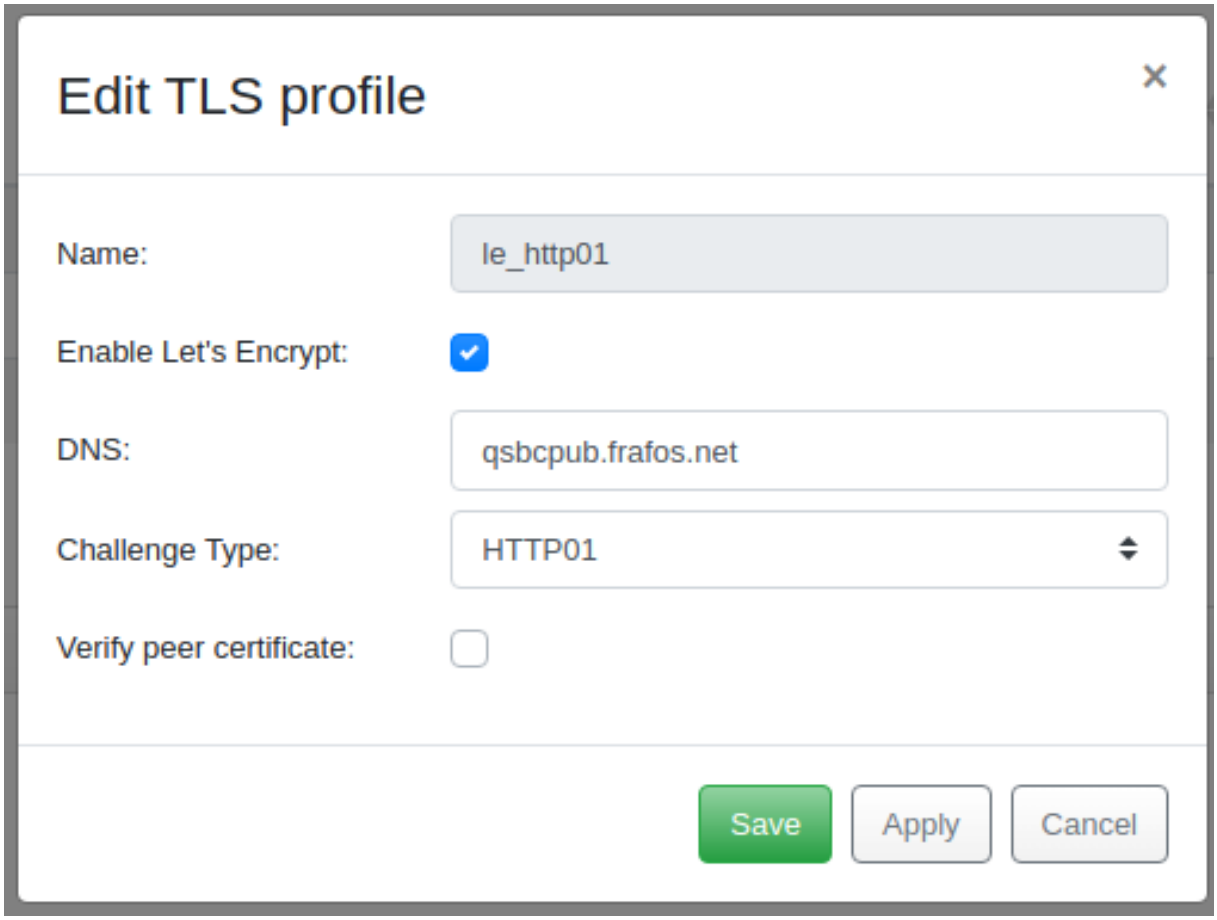

Figure 32: Profile using the http01 challenge

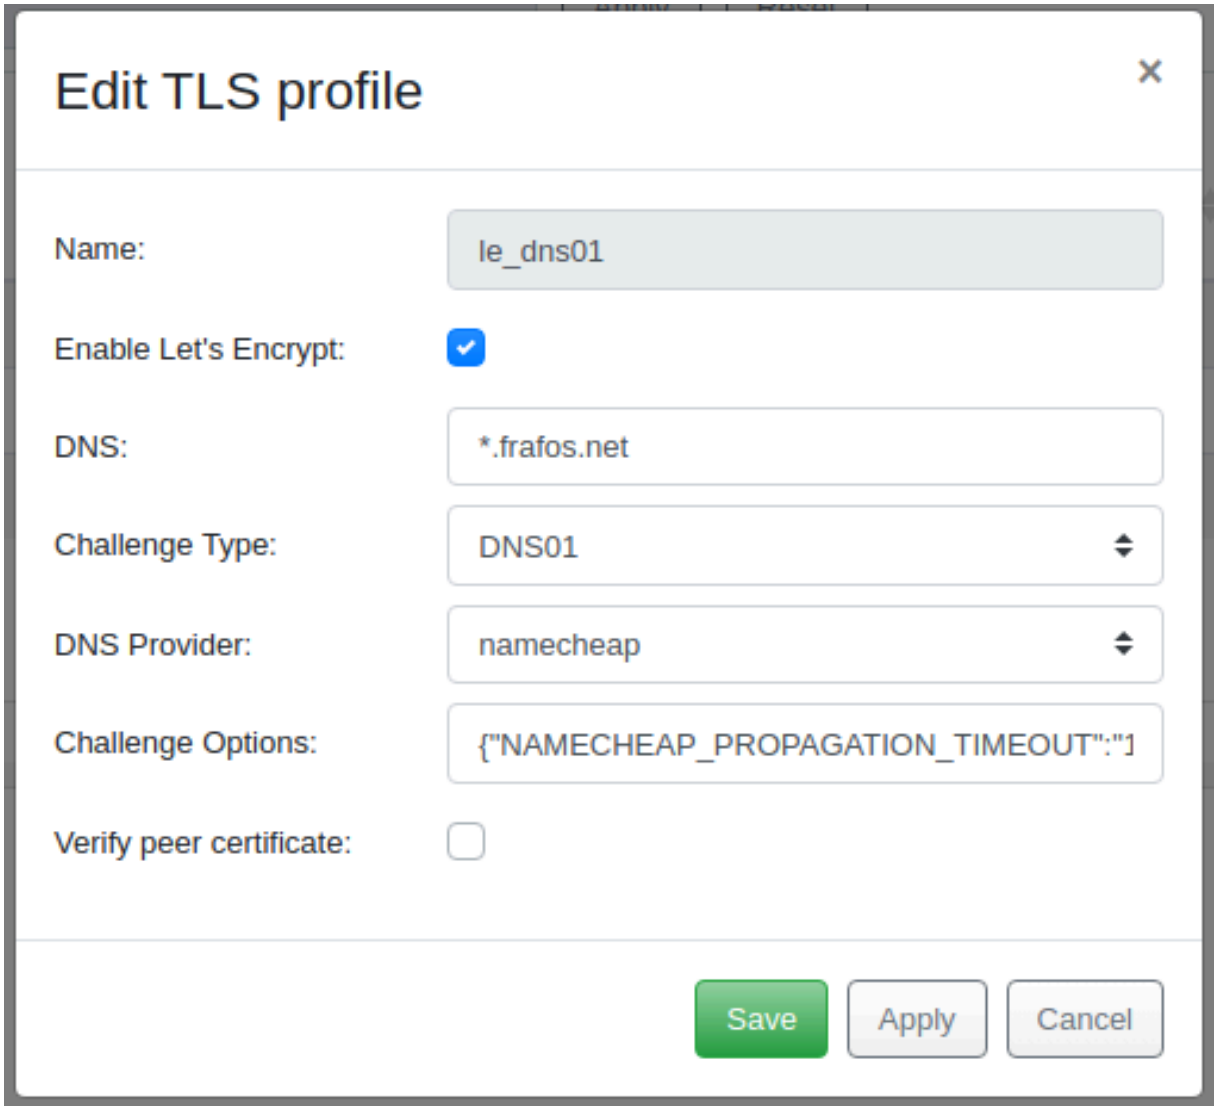

Figure 33: Profile using the dns01 challenge

Once the profile created, depending on the challenge type (refer to Limitations for more), we assign it to one or more node/config groups interfaces. For the Let's Encrypt certificate authority challenge to be attempted, we then trigger a new configuration deployement from the Cluster Config Manager.

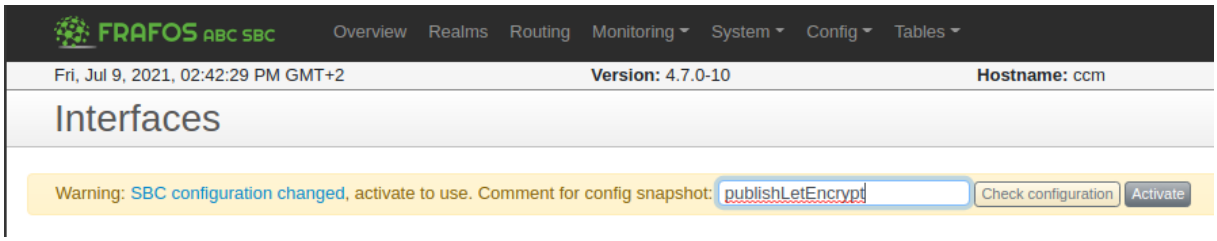

Figure 34: Triggering the publish of a new configuration from the Cluster Config Manager

Process

Once the Let's Encrypt certificate authority requirements deployed to the requested nodes, the Cluster Config Manager `gocertbot` will attempt to complete the challenges.

1. `gocertbot` ask Let's Encrypt certificate authority to complete the selected challenge for the given DNS
2. Let's Encrypt certificate authority asnwer with a set of secret token to present to complete the challenge

The following occure depending on the challenge type:

http01

1. `gocertbot` forward the challenge's token to the target node's `gopi` (port :4224)
2. `gopi` dump the token to the node filesystem (`/var/www/lets_encrypt/{TOKEN}`)
3. `gopi` start an nginx instance to serve the token trough the HTTP protocol
4. `gopi` allow the port 80 trough the SBCTEMP iptable rules
5. `gopi` give the green light to `gocertbot`, whom will forward it to Let's Encrypt certificate authority
6. Let's Encrypt certificate authority attempt the challenge by accessing `http://DNS/.well-known/acme-challenge/{TOKEN}`

dns01

1. `gocertbot` create a TXT record derived from that token and our account key
2. `gocertbot` request the DNS provider to put that record at `_acme-challenge.<DNS>`
3. `gocertbot` give the green light to Let's Encrypt certificate authority
4. Let's Encrypt certificate authority attempt the challenge by accessing `https://_acme-challenge.<DNS>`

Success

On success, a set of TLS certificate, private key and CA bundle certificate are delivered to the `gocertbot` process. Those values are persisted in database, which triggering a new "dirty" state warning.

To complete the process, we need once again to publish the new configuration from the Cluster Config Manager, which provoke a dump of the new certificates into the `/data/sbc/tls/` directory of the target nodes.

If doable (refer to Limitations for more), sems process hot reload the ssl certificate so active calls aren't interrupted.

Failure

In case of failure, a mail is sent to the configured mail address. Logs are accessible either via syslog, or in `/var/log/frafos/sbc.log`. Note that errors are also reported per certificate in `/var/log/frafos/certbot/[profile name]` and monitored by the `sbc-status-check` service.

Requirement

The profile's "DNS" field has to be set to a DNS name resolving to the public IP address of the ABC SBC target node where the corresponding TLS profile is to used. The challenge to verify ownership will be done automatically against it.

A valid email address need to be registered so it will be used to create an Let's Encrypt certificate authority account and receive email alerts (GlobalConfig > System Monitoring > email address). It's better for the `from email address` field to be set. Refer to System Monitoring Parameters for more.

Renewal

The certificate renewal will be attempted automatically 15 days before it's expiration. After certificate creation or renewal, a notification email will be sent to administrator email address set under Global config / System monitoring and new config has to be activated from ABC SBC GUI to propagate the new certificate to SBC nodes.

Limitations

- `http01` challenge profile cannot be assign to more than a single node system interface
- `http01` challenge **doesn't allow** wildcard certificate
- `dns01` challenge **allow** wildcard certificate, but the DNS provider must be supported (_provider list)

- `sems` isn't able to hot reload WS interface certificate - as so, in case of certificate renewal, the whole process is restarted

Debug

process	good for	logs
gopi	up nginx instance, present LE token	`systemctl status sbc-gopi` You can also set `"dev": true` in `/etc/frafos/sbc-gopi.conf`
gocertbot	request LE's for the challenge filter and persit the cert database	logs are outputed to the USER syslog facility, in /var/log/frafos/sbc.log

Table 2: Debug options

You may manually invoke the certbot, from within a Cluster Config Manager' shell by running the following:

```
% sbc-gocertbot -d
```

In case of testing, to avoid reaching LE' 168h rate limit, please remember to enable the "Query Let's encrypt staging environment" Cluster Config Manager' config options.

4.9 Hardware Specific Configurations

Depending on the hardware used for the ABC SBC deployment, there may be some fine-tuning needed to get maximum performance.

4.9.1 Network adapters

If the SBC is configured to work as an RTP media relay and a high number of concurrent calls is expected, a good choice of hardware is critical, specifically in terms of the used network adapter. RTP media traffic means high packet rate, with many small packets passing through. Some network adapters have suboptimal throughput under such conditions. Important things to consider when choosing a network adapter are:

- More receive and transmit packet queues are better. Each queue should be using separate interrupt.
- The adapter method of distributing packets to individual queues should include not only IP addresses into the "hash" calculation algorithm, but should include also IP packet port numbers. (Otherwise the traffic may end up in just one or two queues in case the SBC is communicating with only a small number of other devices on even just one IP address.)
- The adapter should be able to buffer packets received and issue interrupts only after some amount of the packets were received or some timeout. This can be usually configured using "coalesce" adapter options.

There is a global config section "Lowlevel" prepared to allow fine-tuning of settings related to network adapters. The settings are applied after the server is rebooted. The reference of the low-level configuration parameters can be found in Section Low-level Parameters.

System administrator can edit the settings depending on the particular hardware used. The settings are:

- Network interfaces on which a "receive packet steering" kernel feature should be enabled. Recommended setting is to enable it on network interfaces used as media interface.
- Ethernet adapter coalescing options and rx/tx ring parameters. These affect how many packets the adapter may buffer before issuing an interrupt. There is no recommended setting, as the values highly depend on the ethernet adapter used.

- Network interfaces on which the individual interrupts for receive and transmit queues should be statically bound to individual CPUs. If running on multi-CPU or multi-core platform, the recommended setting is to enable this option for all network interfaces used as media interface.
- Options to unload kernel modules for connection tracking or to disable connection tracking completely. Recommended setting is to stop connection tracking. However firewall rules used on the SBC have to be considered as those may need connection tracking active. Note: the default firewall rules that come with the SBC do not use connection tracking.
- Option to enable or disable automatic run of "mysqlcheck" command at end of server boot process. This command checks and repairs (if needed) MariaDB ABC SBC database tables. Default and recommended setting is to enable it.

4.9.2 Configuration of SBC Number of Threads

The major processes of the ABC SBC are running under the name of **sems**. The number of SBC "sems" process threads affects the overall performance in terms of the maximum number of concurrent calls or maximum rate of calls per second supported by the ABC SBC. The optimal settings depend quite a lot on the number of CPU cores of the server used and also on the type of traffic being processed. As a general rule, for high number of concurrent calls including RTP media with relatively low calls per second rate lower numbers of threads performs better, while for high rates of calls per second with SIP only and no RTP media higher number of threads performs better.

The default value for the number of threads is 16. The recommended settings are:
- for SIP+RTP traffic use a number of threads equal to the number of CPU cores multiplied by 4
- for SIP only traffic (no media) use a number of threads equal to the number of CPU cores multiplied by 16

The number of threads can be configured under "Config ▸ Global Config ▸ Lowlevel".

Figure 35: Configuration of SEMS threads

4.9.3 Configuration of sysctl settings

Tuning of some kernel sysctl settings can be considered too, for better performance. These settings need to be applied on the host where the container is running, as usually inside the container the values cannot be increased above the host side settings.

The kernel sysctl settings are typically configured by editing "/etc/sysctl.conf" file or by providing custom config file in "/etc/sysctl.d/" directory on the host system, and activating by running "sysctl -p", depending on the OS used there.

It is recommended to increase the socket receive and send buffer sizes, by setting these sysctl options:

```
net.core.rmem_max = 26214400
net.core.wmem_max = 26214400
```

If ABC SBC uses firewall and connection tracking is used, for high traffic case it is recommended to increase the maximum number of connection tracking entries:

```
net.nf_conntrack_max = 1000000
```

4.10 General ABC Configuration Guide

4.10.1 Physical, System and SBC Interfaces

In the ABC SBC we distinguish between physical, system and SBC interfaces, see the Figure ABC SBC Interface definition:

- A physical interface is one of the network interfaces (cards) physically available on the system.
- System interfaces is an interface mapped on one or more of the physical interfaces. A system interface can be a "simple" physical interface (e.g. "eth2"), a VLAN (e.g. "eth3.1") or a bonded interface that is bound to two physical interfaces (e.g. "bond0" created by bonding "eth0" and "eth1" physical interface). For active/hot standby high-availability mode, it is highly recommended to use bonded interfaces with each physical interface connected to separate L2 switch to ensure reliable physical connections.
- SBC interfaces: These are logical interfaces used by the ABC SBC in order to distinguish between management, signaling and media traffic.

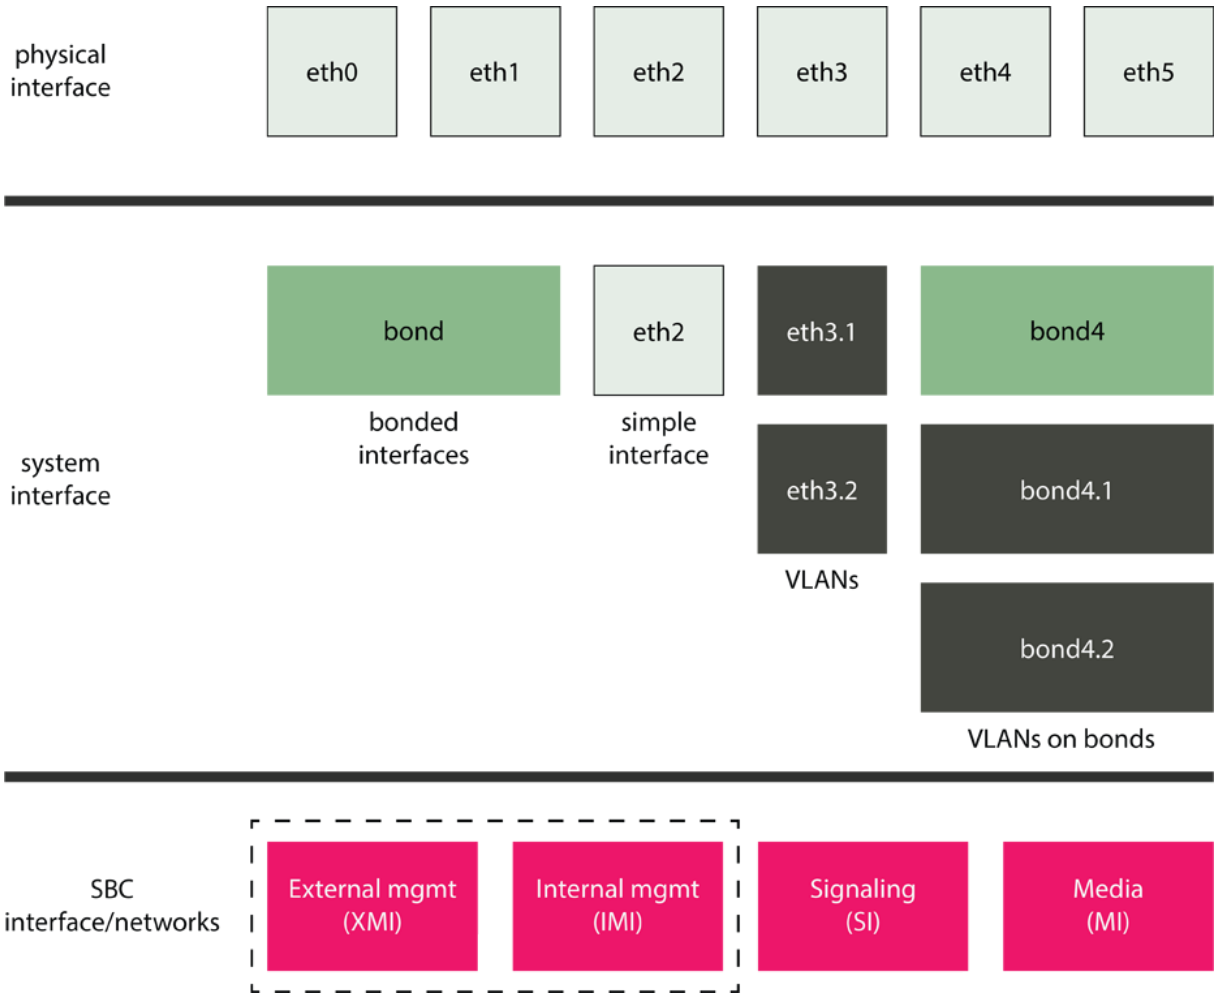

Figure 36: ABC SBC Interface definition

For the details on the configuring the interfaces see Section Interface Configuration.

4.10.2 Defining Rules

The ABC SBC's behavior is specified in form of rules, as explained in Section A-B-C rules. These rules consists of conditions and actions and are processed sequentially until a matching rule is found.

Each rule may have none, zero or multiple conditions. If no condition is specified, the rule always matches and all its actions are executed. If multiple conditions are associated with a rule, the rule matches only if all of the individual conditions match.

An example of such a rule is shown in the Figure Example Rule. It consists of two conditions that match if a call is intended for a telephone number beginning with 900 and the caller is not registered. If the condition applies, the call is rejected using the 403 SIP code.

	Conditions	Actions	Continue	Active	Comment				
☐	R-URI User begins with "900" AND Register Cache From URI (AoR+Contact+IP/port) "Is Not Registered"	Reply to request with reason and code: Code: 403, Reason: must be registered for 900 calls, Header fields:	✔	✔	requests to 900 numbers permitted only for registered users	edit	clone	up	down

Figure 37: Example Rule

Each individual rule condition consists of three parts: a condition type, an operator and a value. Subsequent subsections describe all these parts in details.

Condition Types

The type of a condition defines the left operand for the operation. The following table describes all available condition types and operators types that are applicable to the respective condition types. Operators “==”, “!=”, “RegExp”, “does not match RegExp”,”begins with”, “does not begin with”, are supported unless specified otherwise.

Condition type	Description
Source Call Agent	Check the source call agent. Only operators == and != are supported.
Source Realm	Check the source realm. Only operators == and != are supported.
Source IP	Check IP address the incoming request was sent from.
Source port	Check port number the incoming request was sent from.
Inbound interface	Check local interface the incoming request was received on. Value has to be chosen from a list of configured signaling interfaces. Only operators == and != are supported.
Source-IP CC (GeoIP)	Check country code against source IP’s geographic region. Note: the geoip database must exist, which needs providing geoip license created using customer account at MaxMind, via global config option under Config / Global Config / Misc tab.
Destination Call Agent	Check the destination call agent. Only operators == and != are supported and only available in realm C rules.
R-URI	Check current request URI.
R-URI User	Check user part of request URI
R-URI User Parameter	Check parameter in username part of request URI. For example in the R-URI “sip:106;name=franta@domain.com”, the parameter “name” can be checked for value “franta”
R-URI Domain	Check host part of request URI, can contain port number
R-URI URI Parameter	Check parameter of request URI.
From	Check From header field value.
From URI	Check value of From URI.
From User	Check user part of From URI.
From Domain	Check host part of From header URI, can contain port number
To	Check To header field value.
To URI	Check value of To URI.
To User	Check user part of To URI.
To Domain	Check host part of To header URI, can contain port number.
Supported header	Check content of Supported header (since version 4.5)
Require header	Check content of Require header (since version 4.5)
Header	Check value of given SIP header or test if a SIP header does not exist. This condition is a kind of “escape-code” for testing headers for which no other conditions exist. The following header-fields will not be processed using this condition: From, To, Call-ID, CSeq, RAck, RSeq, Route, Contact, Via, Max-Forwards, Record-Route, Content-Type, Content-Length
Codecs	Check presence/absence of codecs within SDP. Right operand specifies codec name. Only operators Contain , Contain RegExp and Do not contain are supported.

Condition type	Description
Media Types	Check presence/absence of media type within SDP. Right operand specifies media type name (e.g., audio,video) Only operators Contain, Contain RegExp and Do not contain are supported.
SRTP Types	Check presence/absence of SRTP type within SDP. Right operand specifies media type name (e.g. DTLS, SDES, none) Only operators `All media contain`, `Any media contain`, `All media do not contain` and `Any media do not contain` are supported.
Call Variable	Check call variable value using selected operator. The call variable has to be already defined by Set Call Variable action. Any condition referring to an undefined value returns FALSE as result.
Call Variable Existence	Tests if a variable exists or is undefined. This is useful for example when table lookups are used to discriminate accurately between non-existing and empty values.
Generic Text Match	Compare two generic text expressions, supporting replacements.
Generic Floating-point Operation	Compare two floating point values, supporting replacements. Only `>`, `<` operators are supported.
IP Address Operation	Filter based on IP addresses and network. For `==` and `!=` operators, both left and right operands must be an IP address. For `is in` and `is not in`, the left operand needs to be an IP address and the right operand needs to be a network in CIDR format. If the prefix part is missing in the CIDR, then it is assumed to be using full bits (i.e. `/32` for IPv4). Available since: 5.5
Method	Check SIP request method. Value has to be chosen from a list of allowed methods. Only operators == and != are supported.
Register cache	Check content of register cache. Operands From URI (AoR + Contact + IP/port), From URI (AoR + IP/port), Contact URI (Contact + IP/port), To URI(AoR),R-URI (Alias) are supported.
NAT	Check first Via address whether the sender is or is not behind NAT. This compares the SIP message source IP with the first Via address and works only if the UA directly communicates with the SBC.
Read Call Variables	Trigger a predefined query in provisioned tables by a specified key value. The condition returns true if the lookup was successful, false otherwise.
Last Action Result	Returns true if the last action completed successfully, false otherwise (analogical to shell $? variable).
Blacklist	Checks if a call-agent is on a black-list (or not). A call-agent is blacklisted when it is not reachable to make sure that no futile attempts to send traffic to it are undertaken.

Condition type	Description
Date and Time	Checks whether a Datetime, Date or Time value is before or after now plus an optional offset. The value may be in the form '2016-05-25 13:10:41', '2016-05-25' or '13:10:41'. The offset can be years, days, hours, minutes or seconds, e.g. '2y', or '30d', or '12h', or '5m', or '1800s'. Note that this changes behavior based on the timezone of the SBC container. The inputs are assumed to be provided in the local timezone of the SBC container. If the SBC container does not have the timezone set, then the inputs are assumed to be in UTC.
Parallel Call Count	Tests if the number of parallel calls is below or above a threshold. The number refers to the specific place in rules execution flow from which the condition was evoked. It does not refer to a global number of calls.
Parallel Call Count (global key)	Tests if the number of parallel calls for a global key is below or equal/above a threshold.
CAPS count (global key)	Tests if the number of calls attempted per second (CAPS) for a global key is below or equal/above a threshold. The global key works in the same way as for the "Parallel Call Count (global key)" condition: a shared counter is maintained across rules for each distinct key value (e.g. a trunk name).
Parallel Call Rule Hit Count	Tests if the number of parallel calls that has the current rule applied is below a threshold. The number refers to the specific place in rules execution flow from which the condition was evoked. It does not refer to a global number of calls. I.e. if this condition is used in a rule on A-rules of a realm, it will increment its counter as long as the rule is successful (i.e. all conditions of it evaluate to `true`). Once it reaches its threshold, the counter will not be incremented any more and the condition will evaluate `false`, until one of the calls that were previously created by successfully applying the rule is terminated. In the case of the routing rules, the outcome of the routing operation is also taken into account in addition to the conditions in the rule. I.e. if no matching dst CA is found with a "call agent based on r-uri" routing rule, the counter is not incremented.
Request source	Tests whether request being currently processed is generated by the SBC itself, alternatively as a result of unattended call transfer.

Table 3: Condition Types

Condition Operators

Operators supported within general conditions:

Operator	Description
==	left operand equals given value
!=	left operand does not equal given value
RegExp	left operand matches given regular expression
does not match RegExp	left operand does not match given regular expression
begins with	left operand starts with given string
does not begin with	left operand does not start with given string
Contain	right operand is contained in

Operator	Description
Contain RegExp	sample described by right operand is contained in
Do not contain	right operand is not contained in

Table 4: Condition Operators

It is important to know that if a mediation action (Section SIP Mediation) changes content of SIP message, the condition will refer to the value **after** modification. E.g., if you apply the rule action "SetFrom(sip:new@from.com)", the "From URI" operator will return sip:new@from.com!

Some conditions types take special operators and/or values. Particularly the "Register Cache" condition tests if a registration can be found in SBC's cache. The condition uses a specific operator that determines which URIs are used for the test.

Supported operators for "Register Cache" are:

Operator	Description
From URI (AoR + Contact + IP/port)	the user with given From URI and Contact is registered from given IP:port
From URI (AoR + IP/port)	the user with given From URI is registered with any Contact from given IP:port
Contact URI (Contact + IP/port)	a user with given Contact is registered from given IP:port
To URI (AoR)	the user with given To URI is registered
R-URI (Alias)	the user with given request-URI is registered

Table 5: "Register Cache" Operators

The value for the "Register cache" condition allows to refine the test. It can be one of the following:

Condition	Description
Is Registered	true if registered using built-in registrar or cache
Is Not Registered	true if not registered at all
Is Registered Locally	true if registered using built-in registrar using the action "Save REGISTER contact"
Is Not Registered Locally	true if not registered using built-in registrar
Is Registered Remotely	true if URI cached using "Enable REGISTER caching"
Is Not Registered Remotely	true if URI not cached

Table 6: "Register Cache" Conditions

Supported operators for "Date and Time" are:

Operator	Description
is after now plus	The left operand is checked for being after the current time plus an offset
is before now plus	The left operand is checked for whether it is before now plus an offset
is after now minus	The left operand is checked whether it is after the current time minus an offset
is before now minus	The left operand is checked for whether it is before the current time minus an offset

Table 7: "Date and Time" Operators

Note that the offset is optional, and it is always added or subtracted to the current time before the comparison.

Available operators for “Supported header” and “Require header” conditions are:

Condition	Description
contains	Checks for presence of given option tag in Supported or Require header field
does not contain	Checks for absence of given option tag in Supported or Require header field

Table 8: “Date and Time” and “Require header” Conditions

Supported operators for “Request source” are:

Operator	Description
is	Matches if request being currently processed was generated in accordance with right operand.
is not	Matches if request being currently processed was NOT generated in accordance with right operand.

Table 9: “Request source” Operators

Supported right operands for “Request source” are:

Operand	Description
local	request locally generated by SBC
call transfer	request locally generated by SBC as result of unattended call transfer

Table 10: “Request source” right side operands

Supported extra operators for “Header” are:

Operator	Description
RegExp All Of	left operand matches given regular expression for all values of headers that can have multiple values, such as Contact header (RFC3261#7.3). Available since: 5.3
RegExp Any Of	left operand matches given regular expression for any value of headers that can have multiple values, such as Contact header (RFC3261#7.3). Available since: 5.3

Table 11: “Header” Operators

Condition Values and Regular Expressions

Values in a condition may be of several kinds. They are interpreted in the following descending order.

- */SBC/ Escape Codes*. These are characters prefixed by backslash (\) that are supposed to be interpreted literally. These are normally used only for special characters. For example, \\ stands for backslash and \$ stands for the dollar character.
- */SBC/ Replacements*. These are variables that refer to different parts of SIP messages or internal variables. They are referred to by $ character followed by variable name and replaced with value of the variable. The variables that can be used are listed in Section Using Replacements in Rules.

- *regular expressions.* Regular expressions are expected if one of the regular-expression matching operators is used. ABC SBC uses the "extended POSIX regular expression" syntax. That means, among others, that a section enclosed in parenthesis can be referred to from back referencing expressions in actions' parameters (see Section Using Regular Expression Backreferences in Rules), the special characters * (star: zero to any), + (plus: one to any), ? (question mark: none or one), and {a,b} (curly brackets: from a to b) specify the number of occurrences, . (dot) stands for wildcard, ^ (caret) stands for beginning of a string, $ (dollar) stands for end of a string, | (pipe) stands for alternation and square brackets are used for character sets (^ as leading character means negation).
- *literals.* This is the simplest case: a value is used for condition "as is" without further interpretation. For example, in condition "R-URI User == foo", the word foo is matched against the value of userpart of request URI.

Note that this interpretation order determines the condition result. If a regular-expression includes the "end-of-string" character, $, it must be preceded by backslash. Otherwise it will be interpreted in the previous step as an attempt to use a replacement. For example, the "empty string" regular expression must be denoted as "^\$". Another more tricky example is "telephone numbers consisting of a star and two four-digit number blocks". To make sure that a regular expression matches the whole userpart of a URI and not just a part of it, it must begin with "^" and end with "$". Because star has a special meaning in regular expression language, it must be preceded with a backslash. And because the backslash may have special meaning in the ABC SBC GUI, it must appear twice. The resulting expression looks like this

```
^\\*([0-9]{4,4})([0-9]{4,4})\$
```

Also note that an expression in the right operand can contain replacements, but can not contain backreferences as described in Section Using Regular Expression Backreferences in Rules. These are only available as action parameters.

Actions

Actions define how a request shall be treated. There are many kinds of, described in the following sections of this guide as well as in the Section Reference of Actions.

The key functionality available through the actions covers the following aspects of VoIP processing:

Action Group	Purpose
SIP Mediation	Manipulation of identity and URIs, header fields, and response codes. See Section SIP Mediation.
SDP Mediation	Manipulation of codec and early media negotiation. See Section SDP Mediation.
Management and Monitoring	Logging traffic and reporting SNMP statistics (Measurements and Monitoring).
Traffic Shaping	Putting quota on SIP and RTP traffic and reporting violations. See Section Traffic Limiting and Shaping.
Media Processing	Handling RTP traffic: RTP anchoring, RTP/SRTP conversion, RTP inactivity detection, audio recording and transcoding. See Section Media Handling.
Identity	Verifying a 2FA PIN number via DTMF and enrolling a user for 2FA number verification.
SIP dropping	Eliminating non-compliant traffic, silently or with a SIP response. See Section Manual SIP Traffic Blocking.

Action Group	Purpose
Scripting	Processing of internal variables that are used to link multiple actions together using intermediate results stored in variables. See Section Binding Rules together with Call Variables.
Register Processing	REGISTER caching and uncaching, registrar, throttling. See Section Registration Caching and Handling.
External Interaction	Queries to external servers by REST or ENUM or internal pre-provisioned database. See section Advanced Use Cases with Provisioned Data.
NAT Handling	Fixing SIP to facilitate NAT traversal in a safer way than by the SIP specification. See Section NAT Traversal.
Other	Some other actions.

Table 12: Actions

Additional rule properties

It is possible to set some additional properties of the rules. Mostly for documenting and maintenance purposes.

Property	Description
Rule is active	Allows to temporarily deactivate the rule.
Comment	For documentation purposes.
Color	Allows to color the background of rules, so they can be categorized in a way (e.g. normalization, security rules, functional, adaptations, etc…)

Table 13: Additional rule properties

4.10.3 Using Replacements in Rules

In many cases, the conditions values and parameters of actions are not known in advance: they depend on elements of processed SIP messages and results of the message processing. Therefore, it is possible to compose the parameters of special strings that refer to SIP processing status. These strings are called "replacements" and are denoted by a dollar ("$") sign followed by an identifier. Each instance of a replacement is replaced by its value when the rule is evaluated.

For example, **$aU** is a replacement for the User part of the *P-Asserted-Identity* header; **$th** is a replacement for the host part of the *To* header. The action Set R-URI with the parameter set to **sip:$aU@$th** combines mentioned parts of *P-Asserted-Identity* and *To* headers of the incoming request and puts them into the request URI of the outgoing request.

All supported replacements are listed in the table below.

Note that these special characters should be backslash-escaped as follows:

- \ → \\
- $ → \$

Note that where replacement expressions are supported, it is possible to use **\r**, **\n** and **\t** to input carriage-return, line-feed and tab, respectively. This can possibly be used to i.e. insert multiple headers but it is likely to break functionality and should be avoided unless absolutely necessary.

It is important to know that if a mediation action (Section SIP Mediation) changes content of SIP message, the substitution expression will refer to the value **after** modification. E.g., if you apply the rule action "SetFrom(sip:new@from.com)", $fu will return new@from.com!

Repl. group	Replacements	Description
$r	$r.	request-URI; note that the expression refers to current request URI which may be changed during the course of request processing
	$ru	user@host[:port] part of request URI
	$rU	R-URI User
	$rd	R-URI Domain (host:port)
	$rh	R-URI Host
	$rp	R-URI Port
	$rP	R-URI Parameters
$f	$f.	From header
	$fu	user@host[:port] part of From URI
	$fU	From User
	$fd	From Domain (host:port)
	$fh	From Host
	$fp	From Port
	$fn	From Display name
	$fP	From Parameters
	$ft	From Tag
	$fH	From header Headers
$t	$t.	To header
	$tu	user@host[:port] part of To URI
	$tU	To User
	$td	To Domain (host:port)
	$th	To Host
	$tp	To Port
	$tn	To Display name
	$tP	To Parameters
	$tt	To Tag
	$tH	To header Headers
$a	$a.	P-Asserted-Identity header
	$au	user@host[:port] part of P-Asserted-Identity URI
	$aU	P-Asserted-Identity User
	$ad	P-Asserted-Identity Domain (host:port)
	$ah	P-Asserted-Identity Host
	$ap	P-Asserted-Identity Port
	$aP	P-Asserted-Identity Parameters
	$aH	P-Asserted-Identity Headers
$p	$p.	P-Preferred-Identity header
	$pu	user@host[:port] part of P-Preferred-Identity URI
	$pU	P-Preferred-Identity User

Repl. group	Replacements	Description
	$pd	P-Preferred-Identity Domain (host:port)
	$ph	P-Preferred-Identity Host
	$pp	P-Preferred-Identity Port
	$pP	P-Preferred-Identity Parameters
	$pH	P-Preferred-Identity Headers
$c	$ci	Call-ID
$C	$C.	complete Contact-HF
	$Ci	user@host[:port], port is included if present in Contact-HF
	$Cx	x' is anything supported for other URIs
$s	$si	Source (remote) IP address
	$sp	Source (remote) port number
$d	$di	expected destination host
	$dp	expected destination port
$o	$o.	Outbound source IP address:port (local). Only works in C-rules and port is the non-TLS listening port configured on the interface (which could be different than the actual used port for outbound TCP/TLS transport).
	$oi	Outbound source IP address (local). Only works in C-rules.
	$op	Outbound source port (local). Only works in C-rules and port is always the non-TLS listening port configured on the interface (which could be different than the actual used port for outbound TCP/TLS transport).
$R	$Ri	Destination (local/received) IP address
$R	$RI	Destination IP address – like above but when a public IP is configured on the receiving interface, its value is used instead.
	$Rp	Destination (local/received) port number
	$Rf	local/received interface id (0=default)
	$Rn	local/received interface name (SBC interface name)
	$RI	local/received interface public IP
	$Rt	local/received transport protocol one of: tcp, tls, udp, ws (WebSocket), wss (secure WebSocket)
$H	$H(headername)	value of header with the name headername; not applicable to from/to/ruri/contact for which specific replacements must be used
	$HU(headername)	header headername (as URI) User
	$Hd(headername)	header headername (as URI) domain (host:port)
	$Hu(headername)	header headername (as URI) URI

Repl. group	Replacements	Description
	$Hh(headername)	header headername (as URI) host
	$Hp(headername)	header headername (as URI) port
	$Hn(headername)	header headername (as URI) display name
	$Hp(headername)	header headername (as URI) parameters
	$HH(headername)	header headername (as URI) headers
$O Available since 5.5	$O(index)	Value of route header at given index. Negative values can be used to access from the end, $O(-1) is the last route header.
	$OU(index)	Route (as URI) User
	$Od(index)	Route (as URI) domain (host:port)
	$Ou(index)	Route (as URI) URI
	$Oh(index)	Route (as URI) host
	$Op(index)	Route (as URI) port
	$On(index)	Route (as URI) display name
	$Op(index)	Route (as URI) parameters
	$OH(index)	Route (as URI) headers
	$Oc	Number of route headers
$m	$m	request method
$V	$V(gui.varname)	value of Call Variable varname
$B	$B(cnum.rnum)	value of backreference with *rnum* number from the condition with *cnum* number
$U	$Ua	register cache: originating AoR
	$UA	register cache: originating alias
$_	$_u(value)	value to uppercase
	$_l(value)	value to lowercase
	$_s(value)	length of value (size)
	$_5(value)	MD5 of value
	$_g(value)	Computes a reproducible unique string (UUIDv5) based on the provided value (i.e. `$_g($ci)`)
	$_r(value)	random number 0..value, e.g. $_r(5) gives 0, 1, 2, 3 or 4
	$_UU(value)	value (as URI) User
	$_Ud(value)	value (as URI) domain (host:port)
	$_Uu(value)	value (as URI) URI
	$_Uh(value)	value (as URI) host
	$_Up(value)	value (as URI) port
	$_Un(value)	value (as URI) display name
	$_Up(value)	value (as URI) parameters
	$_UH(value)	value (as URI) headers
$#	$#(value)	value URL-encoded

Repl. group	Replacements	Description
$time	$time(value)	time format as described in the libc strftime() function. ie: $time(%m-%d-%y-%H-%M)
$attr	$attr(value)	value of the given global attribute
$cntr	$cntr(value)	value of the given counter defined by User Defined Counters
$e164	$e164(number, country_code)	Convert the number parameters to the e164 format of the country code. ie: $e164(0635215099, FR) = +33635215099
$T	$T(number, country_code)	Return the phone number type as a numeric string, given a phone number and a country code. Uses libphonenumber internally. Possible values: 0=FIXED_LINE, 1=MOBILE, 2=FIXED_LINE_OR_MOBILE, 3=TOLL_FREE, 4=PREMIUM_RATE, 5=SHARED_COST, 6=VOIP, 7=PERSONAL_NUMBER, 8=PAGER, 9=UAN, 10=VOICEMAIL, 11=UNKNOWN. Example: $T(0635215099, FR) = 1 (MOBILE)
$rc2cc	$rc2cc(region_code)	Return the country code of the region. ie: $rc2cc(FR) = 33
$cc2rc	$cc2rc(country_code)	Return the region code of the country. ie: $rc2cc(33) = FR
$tls_	$tls_subject	TLS Subject Name formatted as in RFC2253
	$tls_subject_cn	TLS Subject Common Name
	$tls_issuer	TLS Issuer Name formatted as in RFC2253
	$tls_issuer_cn	TLS Issuer Common Name
	$tls_peername	Return the verified peer name
$MOS Available since 5.5	$MOSavgCA(ca_uuid, s/m/l, default_value)	Return the RTP MOS QoS average value. Second argument is either `s`, `m` or `l` which returns the short-term, medium-term and long-term average value. Coefficients for these can be configured in global configuration. If no MOS information exists (i.e. no calls on the CA yet), default value is returned.
	$MOSavgRealmCAName(realm_name, ca_name, s/m/l, default_value)	Same as `$MOSavgCA` except that it takes realm name and CA name instead of CA UUID.
	$tls_peername	Return the verified peer name
	$tls_peername	Return the verified peer name

Table 14: Replacements

Example Use of Replacement Expressions

In the following example, see Fig. Using Replacements, we set up the outgoing INVITE request as follows:

- set Request URI of the outgoing INVITE request to the user part of the *P-Asserted-Identity* header ($aU) combined with the host part of the *To* header ($th) of the incoming INVITE request
- set host part of the *To* header to the value of the *P-NextHop-IP* header ($H(P-NextHop-IP)) of the incoming INVITE request (the user part will not be changed)
- convert the user part and the host part of the *From* header into lower case (<sip:$_l($fU)@_l($fh)>).

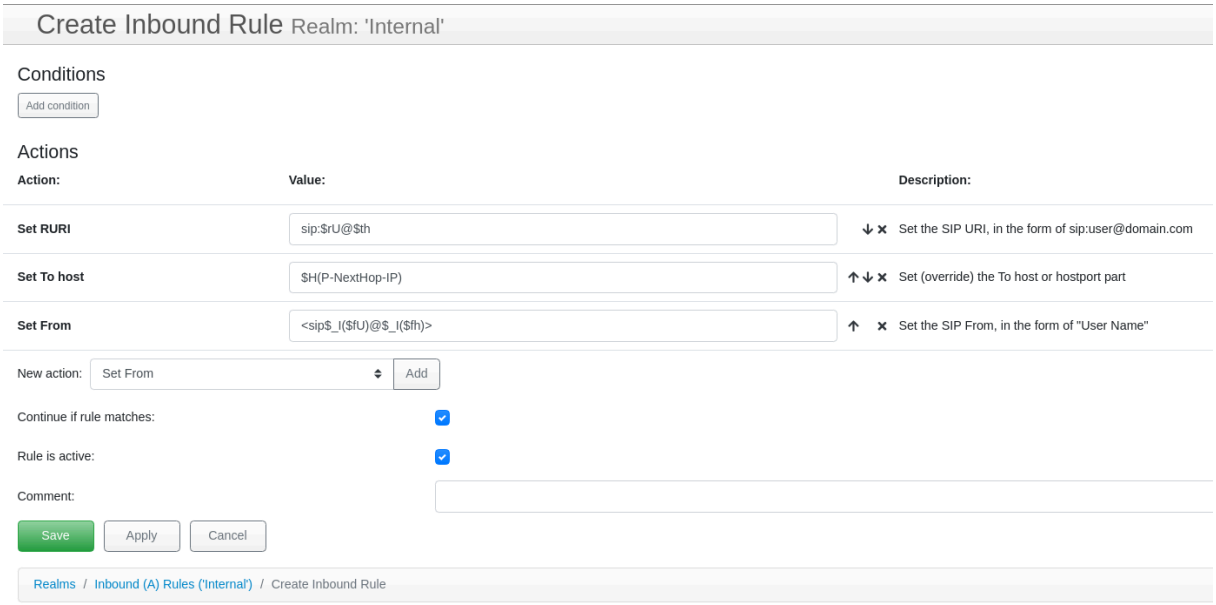

Figure 38: Using Replacements

The effects of this transformation on a SIP message is depicted in Fig. Effects of using replacements:

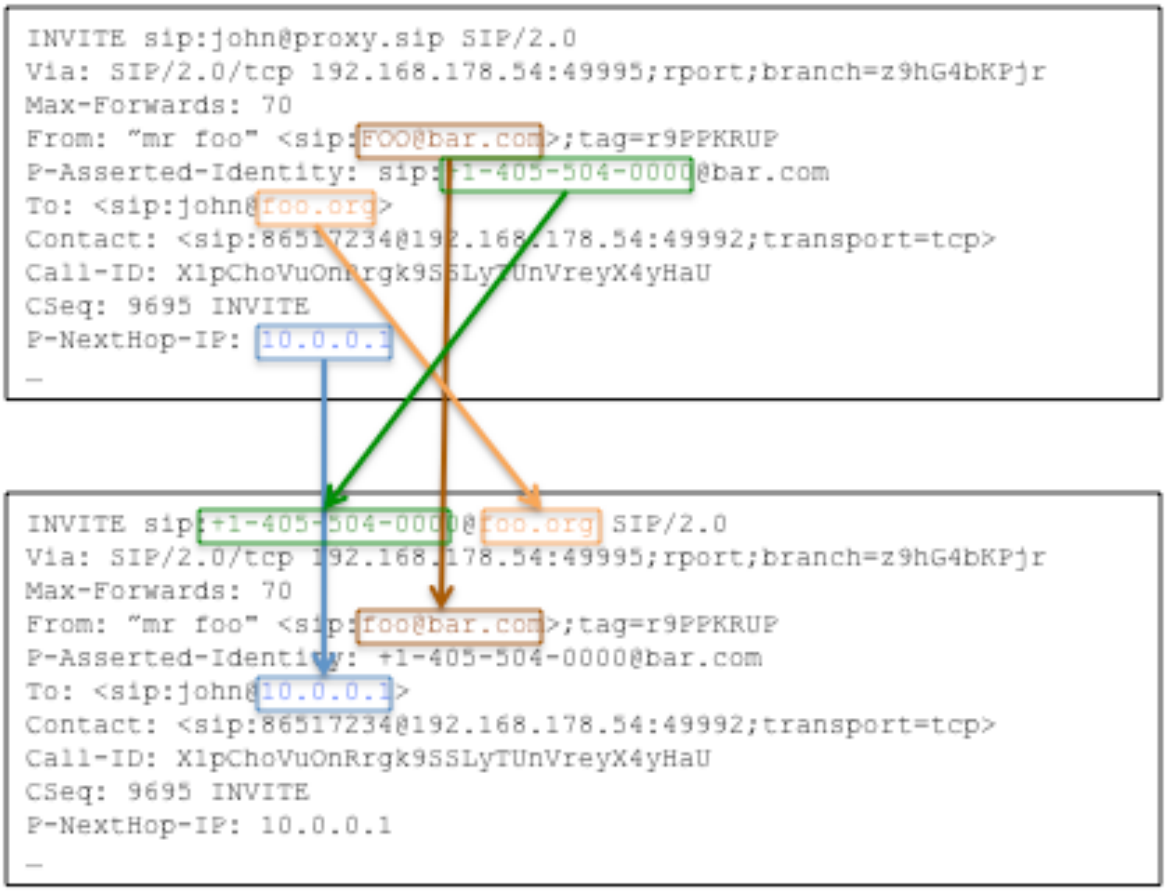

Figure 39: Effects of using replacements

4.10.4 Using Regular Expression Backreferences in Rules

Whenever a regular expression match is executed in a rule condition, the matched substrings can be later used in subsequent actions or conditions. The matched result is referred to by so called "backreferences".

Backreferences are used by the replacement **$B(c.r)**. The first index in the backreference, **c**, denotes the index of the condition, where the first condition has the index 1, the second condition the index 2, and so forth. The second index, **r**, denotes the index of the substring selection in the regular expression, where the first selection has the index 1, the second the index 2, an so forth.

In the following example, see Fig. Using backreferences, we use backreferences to separate protocol discriminator ("sip" or "tel") from the rest of request URI. These two parts are matched in the regular expression in the 2nd condition and are therefore referred to as $B(2.1) and $B(2.2). Particularly, the example saves the protocol discriminator from the request URI in an INVITE request to a call variable called **uri_scheme**. Further it enforces the "sip" scheme for the R-URI of the outgoing INVITE request.

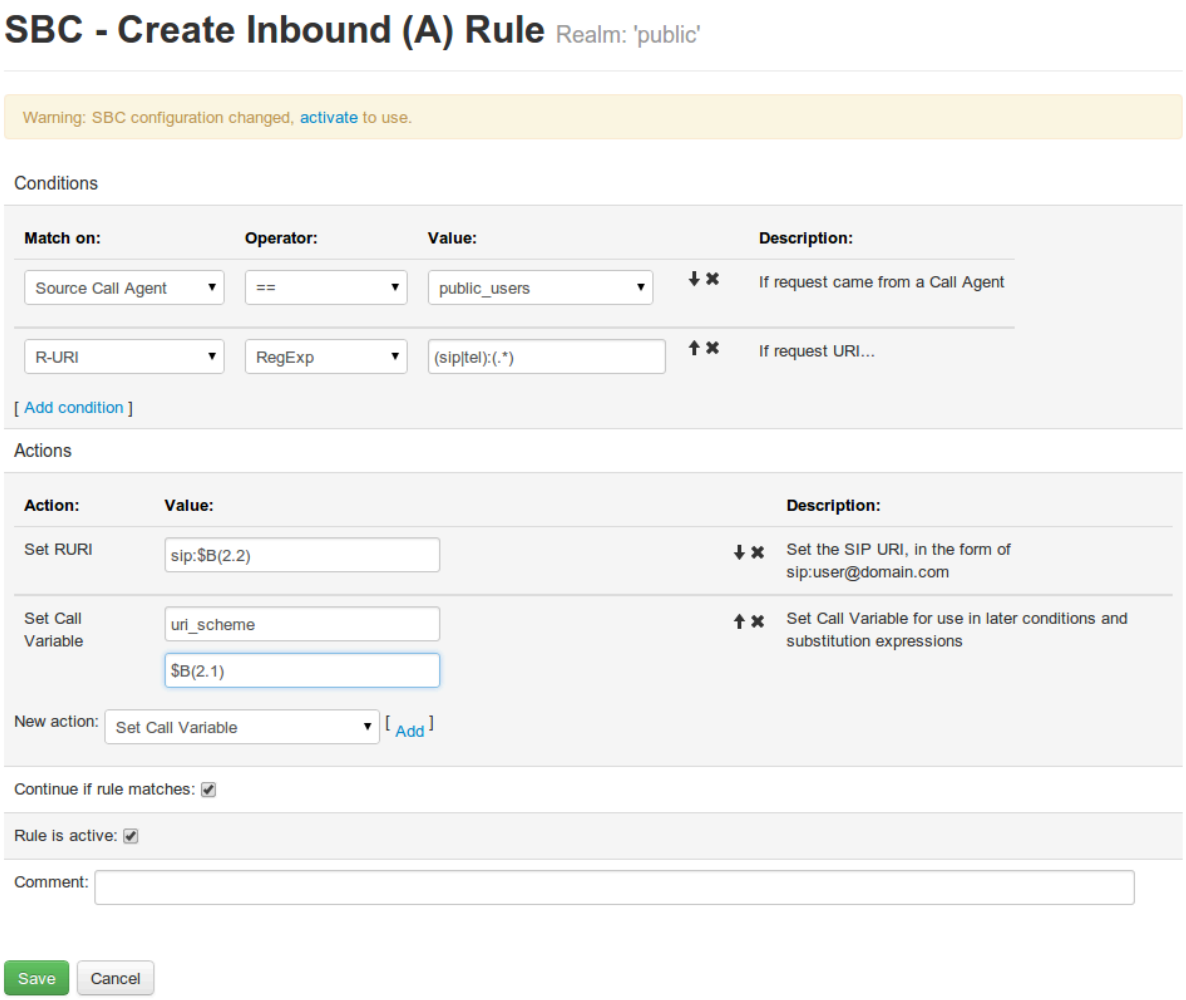

Figure 40: Using backreferences

4.10.5 Binding Rules together with Call Variables

Call Variables are a very powerful tool in the ABC SBC, because they can bind together different rules or rule sets. Call Variables can be set by rules to any value using the **Set Call Variable** action. This variable will the persist during the lifetime of the call. They can be set to a different value by a subsequent rule, again with the **Set Call Variable** action.

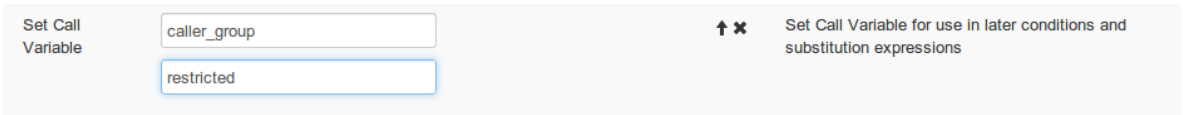

Figure 41: Setting Call Variables

Values of Call variables can be tested with the **Call Variable** condition using several operators: ==, !=, "RegExp", "does not match RegExp", "begins with" and "doesn't begin with". Operands may be literal strings, regular expression if the "Regexp" operators are used, and they may contain Replacements (see Section Using Replacements in Rules).

An additional condition, "Call Variable Existence", allows to test if a variable exists and accurately discriminate it from the case when it is empty-valued. This may be particularly handy when table lookups are used as described later in Section Provisioned Tables. Otherwise reference to undefined variables always returns empty string.

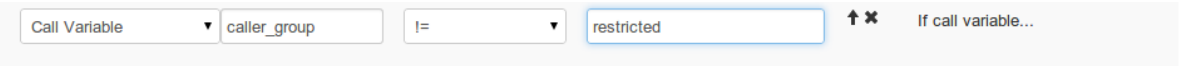

Figure 42: Testing Call Variables

They can also be used in other actions using the replacement expression $V(gui.varname), where *varname* refers to the name of the variable, e.g. $V(gui.caller_group).

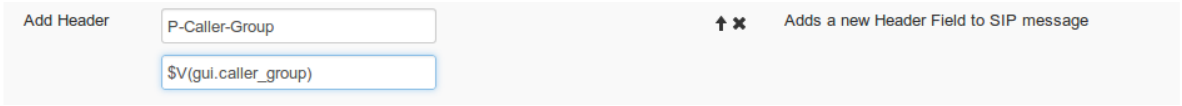

Figure 43: Using Call Variables

The following example shows how a variable is assigned a value using the "Set Call Variable" action (see Figure Example for using Call Variables), tested for a specific value "restricted" (see Figure Testing Call Variables) and referred to from an action for adding a new *Reason* Header field using the $V replacement (see Figure Using Call Variables).

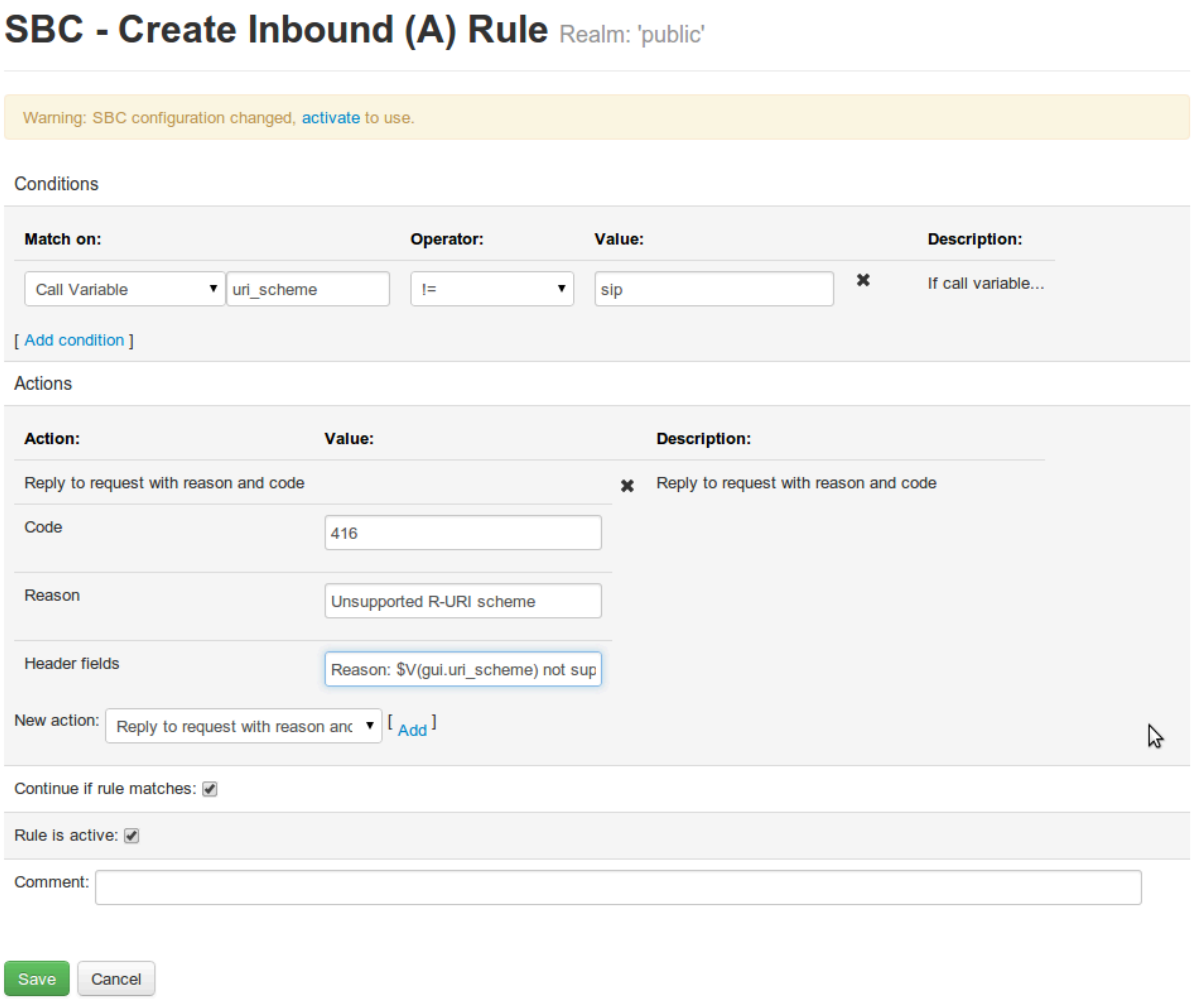

Figure 44: Example for using Call Variables

4.10.6 SIP Routing

The key functionality of the ABC SBC is that of SIP routing: based on criteria chosen by the administrator, the SIP destination for a SIP dialog is chosen. The routing decision is the step "B" in the A-B-C process: After A-rules are applied based on who sent the SIP traffic to the ABC SBC, the destination

Call Agent is chosen in the B-rules. The final step is processing of the outbound C-rules, that are specific to the Call Agent chosen in the step B.

The routing decision is also an important part of the network reliability concept: it has implications to the way how traffic is re-routed if downstream destinations are unavailable or overloaded.

Unlike steps A and C, the routing step B is global: it is executed for every combination of inbound and outbound call agents and realms. It can be seen like the wiring board between these. Also unlike A and B rules, matched rules can have only one action: selection of the destination.

The routing rules are processed sequentially until one is found that matches. Repetitive rules, such as least-cost-routing tables, can also be managed by provisioned tables as described in the Section Provisioned Tables. If no route matches, the ABC SBC stops processing the SIP request and returns a 404 SIP response.

The outcome of the routing process is unique determination of the destination Call Agent. This decision determines the following aspects:

- which C-rules are executed,
- which backup Call-Agent is used, if forwarding to the chosen Call Agent fails,
- which interfaces are used for forwarding,
- which IP address or addresses are used as next hop for forwarding.

Note that the routing process is applied only to dialog-initiating or out-of-dialog SIP requests. During the dialog life-time, routing of in-dialog SIP requests follows a fixed path established in the process of the dialog initiation with the peering SIP devices. The path may or may not be the same as that of dialog-initiating transaction and is formed using Record-Route and Contact header-fields as governed by the **RFC 3261** specification. Only if the "use on first request only" option is turned off, or "Dialog NAT handling" is enabled, the ABC SBC routes subsequent requests in a sticky way to the same hop as the initial one.

The following subsections describe how routing rules are organized and how to use the three types of routing rules: "static" for well-known next-hop Call Agents, "table-based dynamic" for a massive amount of static routes, and "request-URI based" for destinations identified in request URI. Eventually we show how the abstract destination is translated into next-hop IP addresses for request forwarding.

Routing Rules (B)

The routing rules are an ordered list of routes, which are processed one by one after completion of A rules processing. When the first rule condition matches, the destination call agent is chosen and route processing stops.

The configured Routing (B) rules can be viewed when clicking on the "Routing" menu entry.

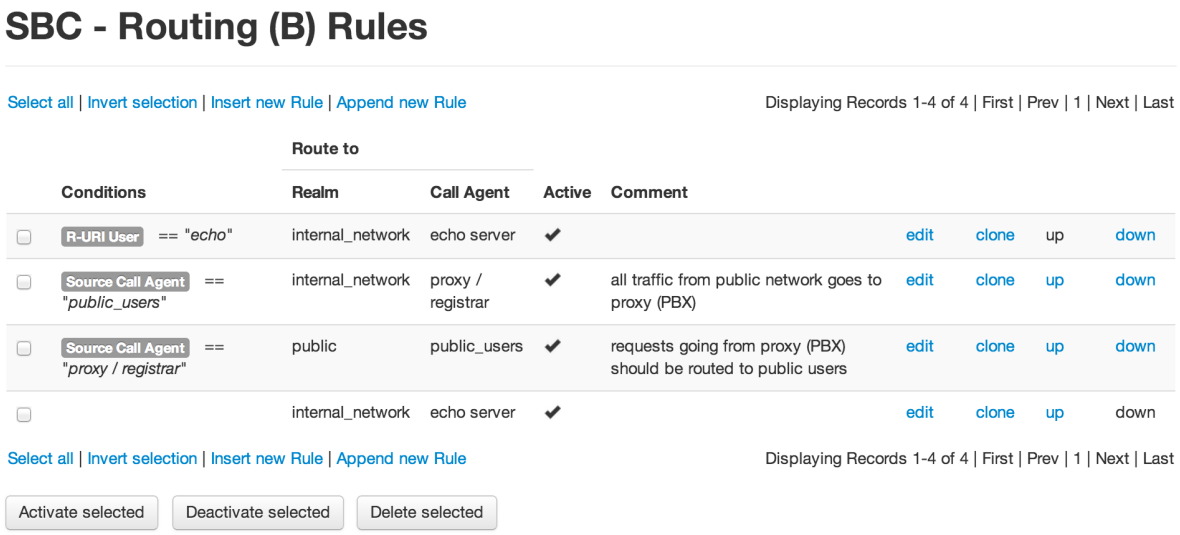

Figure 45: List or routing rules

A new routing rule can be created either at the beginning of the list ("Insert new Rule") or at the end ("Append new Rule").

When creating a new Routing (B) rule, one or more conditions can be entered, see Fig. Define matching conditions for routing, similar to the other types of rules (inbound (A) and outbound (C) rules).

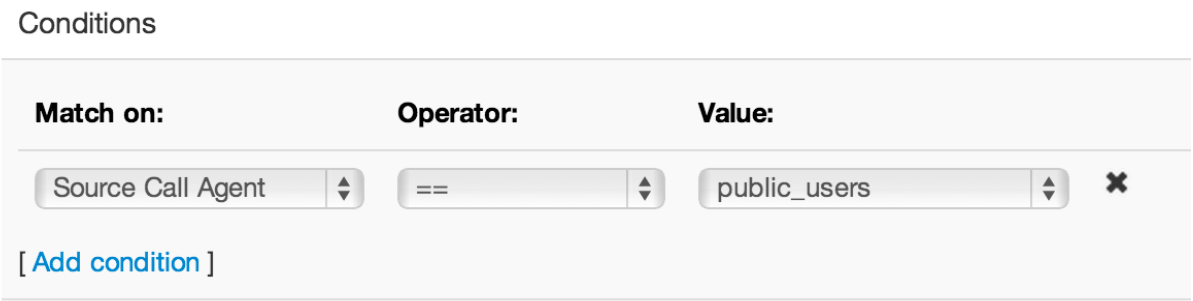

Figure 46: Define matching conditions for routing

It is also possible to define a default routing rule by omitting the condition(s). In this case, the rule will always match, and thus finish the Routing rules evaluation. In case there is no such default rule and no other rule matches, the ABC SBC answers the request with a *404* response.

There are several types of routing rules, that can be combined with each other and processed in sequential order:

- **Static route** – In a static route, all data describing the routing decision must be entered manually. If the static route matches, routing finishes, otherwise it proceeds to the next rule.
- **Dynamic route** – if multiple repetitive rules, such as with Least Cost Routing, are configured, they are better placed in a provisioned table as described in the Section Provisioned Tables. To make a lookup in the table, select the table name in the "Route using" drop down list and indicate by what key the lookup shall be performed. Like with static routes, if a matching entry is found, routing finishes, otherwise it proceeds to the next rule.
- **Call Agent based on R-URI** – the ABC SBC tries to find a Call Agent that matches host in request URI. This can be particularly useful if A-rules change hostname in request URI, for example by ENUM lookup. If the lookup yields an address of a valid Call Agent, it is used for routing and routing finishes, otherwise it proceeds to the next rule.

These route types are described in more detail in the following sections.

Validity of routing rule can be limited to some nodes only. In this case the routing rule is not executed on other nodes. Routing rule can be assigned to only those nodes which belongs to any of the config groups assigned to the particular routing table.

Static Routes

Static route is the simplest type of routing rule: the administrator explicitly chooses the destination Call Agent. If no further specific treatment is desired, that's all. The Call Agent is chosen, subsequently its C-rules are executed and eventually signaling is forwarded through the interface associated with the Call Agent. This routing method is applicable to all Call Agent types but those identified by a subnet address – these are used primarily for matching of incoming traffic and do not uniquely identify a destination forwarding address.

The choice of Call Agent is accompanied by several other options. The most important is that of routing method which specifies how the next-hop IP address is determined. Either it is determined from request URI or from pre-provisioned information. Note that whichever method is chosen to determine the next-hop IP address, Call-Agent does not change and its C-rules are used for request processing. Both methods may yield multiple IP addresses, in which case the ABC SBC load-balances among them by their respective priorities.

The "Route via R-URI" method uses the request URI to find out the next-hop IP address. That is particularly useful when A-rules altered the request URI using actions like reverse registration cache or ENUM lookup. If the host part of request URI includes a DNS name that resolves to multiple destinations per **RFC 3263**, the ABC SBC load-balances among the respective destinations by their priorities.

If "Set Next Hop" (also known as "outbound proxy") is used instead, the next-hop IP address is determined using pre-provisioned information. Either the IP address (or addresses) associated with the Call Agent is taken, or these are explicitly overridden using the option "Use another destination instead of CAs' destination(s)". Please note that, if "Use another destination instead of CAs' destination(s)" is used, backup CA will not fully work (C rules from the CA **will** be applied to outgoing request). In that case, if the first CA fail, all calls are sent to the same IP, located in "Use another destination instead". Further method-specific options include:

- "Use on first request only" – this option changes default behavior for forwarding subsequent in-dialog requests. By default when turned off, all subsequent outbound requests will follow exactly the same the path of the previous dialog-initiating request. If however this option is turned on, the next-hop logic for subsequent requests is governed only by the SIP standard procedures. Particularly, if the next hop in the INVITE path was a non-record-routing proxy, it will not be included in request's path.
- "Update R-URI Host". This option rewrites host part of request URI with the address of the next hop. By default it is turned off and the request URI remains untouched when forwarding.
- "Add Route HF". This option is also known as "preloaded Route". It prints the next-hop destination in Route Header-field. Use only if downstream SIP hop is known to require such behavior.

The following advanced options can be also used with both methods:

- "Replace DNS name in R-URI through the resolved IP address" makes sure that if DNS names appears in request URI, it is rewritten to IP address before forwarding.
- "Force transport". Allows to override transport protocol to be used for the next hop. One of the following protocols can be chosen: UDP, TCP, TLS.
- "Enable redirect handling". If this option is turned on, incoming 302 are not passed upstream. Instead, the ABC SBC takes content of Contact header field and uses it as another next-hop for forwarding the original request. Particularly the Contact URI in the 302 response is used to rewrite request URI, determine the next-hop IP address and look up a Call Agent whose C-rules are processed. Note that an error occurs and a 500 response is sent upstream if none or more than one matching Call-Agents are found, or the Contacts include DNS names. Use this option with care only for trusted destinations since the 3xx responses may greatly impact where and how requests are forwarded.

An example is shown in Figure Static Routing destination: the Call Agent "users" is chosen so that its C-rules will be processed. There is no additional IP address included in the "set-next-hop" choice of routing, so that the dialog-initiating request is forwarded to IP addresses associated with the particular Call Agent.

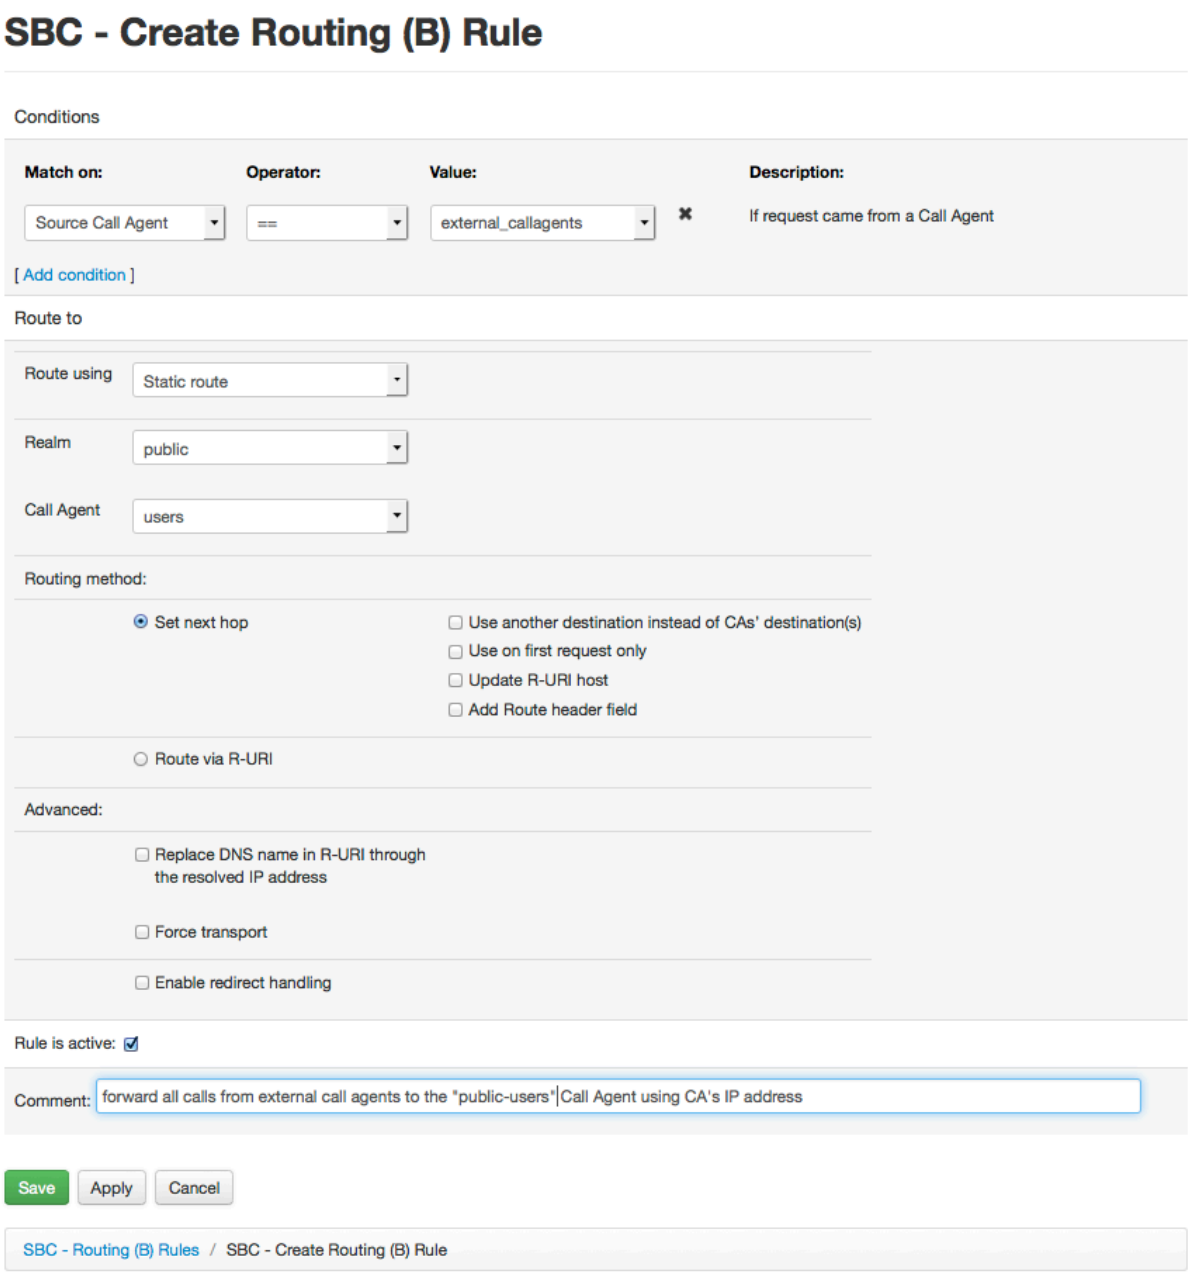

Figure 47: Static Routing destination

Table-based Dynamic Routes

Long repetitive routing rule sets can be better managed as tables. All other aspects of the routing logic remain the same as with statically defined rules.

To deploy dynamic rules the following steps must be performed:

- definition of a routing table (see Section Configuring Tables)
- definition of routing lookup performed against the table in B-rules (see example in Figure Configure a route lookup in a provisioned routing table)

- filling in the routing table with routing data (see example in Figure Adding a new entry to routing table)

The lookup definition requires two parameters: name of the table defined in the first step, and the value used to lookup a matching table row defined in the second step. The value is defined in form of a replacement expression. For example, $rU can be used to trigger lookups by user part of request URI.

The table than includes rows identified by unique keys. In the screenshots bellow, the user part of request URI ($rU) is compared against a row with key 911. If a match occurs, the call agent "external_callagents" is used for call forwarding.

When the table lookup is performed and the value matches no key, routing proceeds to the next entry in the routing table. If there is no more such, routing fails and the SIP request is declined using the *404* SIP response.

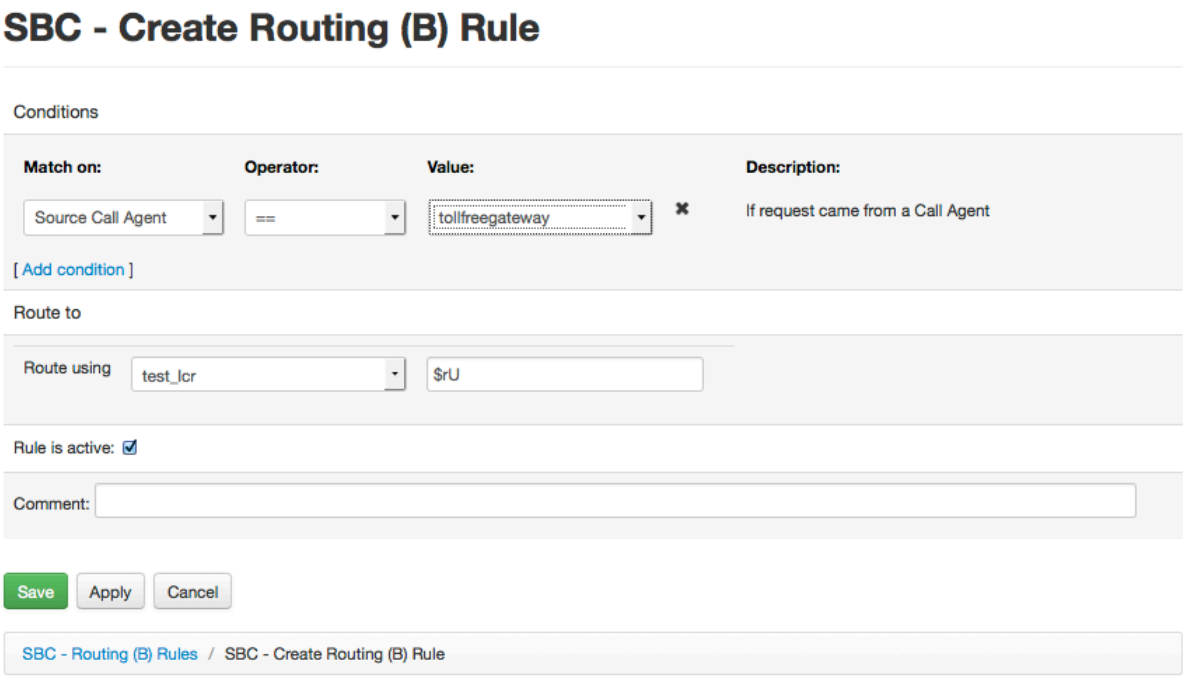

Figure 48: Configure a route lookup in a provisioned routing table

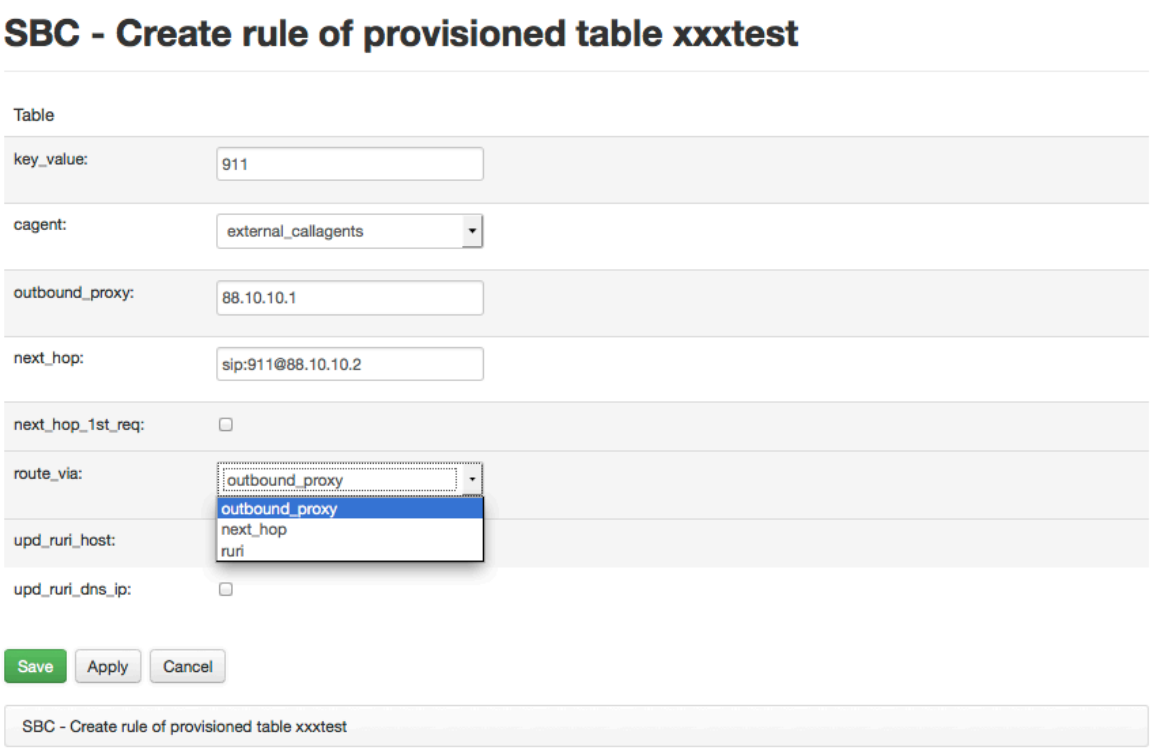

Figure 49: Adding a new entry to routing table

Request-URI Based Routes

In some scenarios, the next-hop Call Agent is not exactly known at the time of devising a routing policy. Instead it is known that a request URI identifies the Call Agent. This is often the case if the request URI is rewritten by an external query, such as ENUM or REST. There would be little point in formulating rules like "if a CA's IP address present in R-URI, route to the CA" for every single CA.

Therefore there is the "route via R-URI" routing type, which finds a Call Agent based on address in request URI and if found, routes to it.

Note that this is different from the "route via R-URI" option, which is only used to override the transport destination but does not determine the Call Agent with its rules.

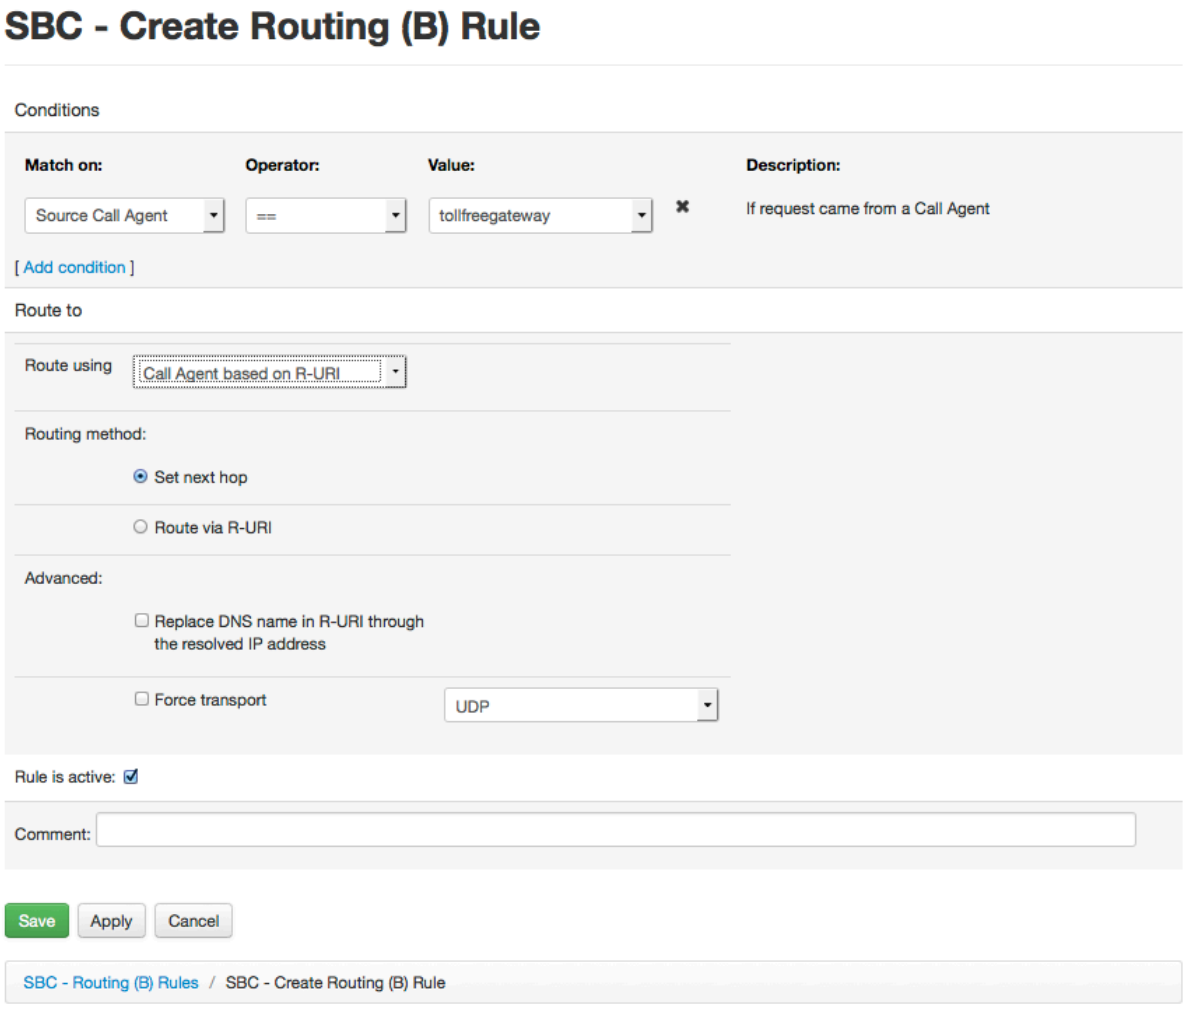

Figure 50: Route by Request URI

Like with Static Routes, there are two routing methods for determining the next IP-hop: Either it is taken from request-URI (the "Route via R-URI" method) or it is taken from Call Agent's profile. The difference is subtle because by use of the lookup an IP address gained from the request URI must match an IP address of a Call Agent. A difference may occur when some other IP addresses linked with the DNS name are different from those linked with the Call Agent's profile. Also if an IP address comes in request URI and the "route-via-r-uri" method is used, alternate destinations associated with the Call Agent will not be used.

Determination of the IP destination and Next-hop Load-Balancing

When the destination Call Agent is selected, one or multiple IP addresses are chosen for forwarding. These may come from Call Agent definition, explicit addresses in the route or from request URI. Capability to choose more than one IP address is important for load-balancing downstream hosts and for dealing with their unavailability. If there are multiple IP addresses (so called "destination set") the ABC SBC "hunts" through them based on their priorities to find one that is responsive.

The destination set is formed depending on the choice of routing method described in previous sections. It works the same way for static, dynamic and request-URI based types and it can be one of the following:

- if the "Set-next-hop" routing method is chosen without the "Use another destination instead of CAs' destination(s)" option, the addresses specified in Call Agent's profile are used

- if the "Set-next-hop" routing method is chosen with the "Use another destination instead of CAs' destination(s)" option, the addresses specified in this option are used. Addresses associated with the Call Agent are not used for forwarding.
- if "Route via R-URI" is chosen, the address is taken from the request URI.

If an address in the destination set is a DNS name, it is resolved to IP address(es) using procedures specified in **RFC 3263** before further processing.

If the resulting destination set includes multiple entries they are attempted in successive order. An 8-second timer is used to try up to 4 destinations, so that the hunting attempts complete before standard SIP transaction timeout of 32 seconds. A 503 response makes the ABC SBC to attempt the next destination in the set immediately.

The hunting order is determined by priorities specified in DNS, "Use another destination" option or CA profile. The way priorities are set complies to the **RFC 2782**: the ABC SBC initially contacts hosts with lowest-number priorities. If there are multiple hosts with the same priority they are tried by probability as defined in their weight field. The weight field specifies a relative weight, larger weights are given a higher probability of selection.

When no responsive destination is found, the ABC SBC will check if there is a backup Call Agent defined in the current Call Agent's profile. If so, it will undo previous mediation changes, process backup Call Agent's C-rules and retry the IP calculation process for the backup Call Agent.

The whole process is shown in Figure Flowchart of Process for Determination of the Next-hop IP Address.

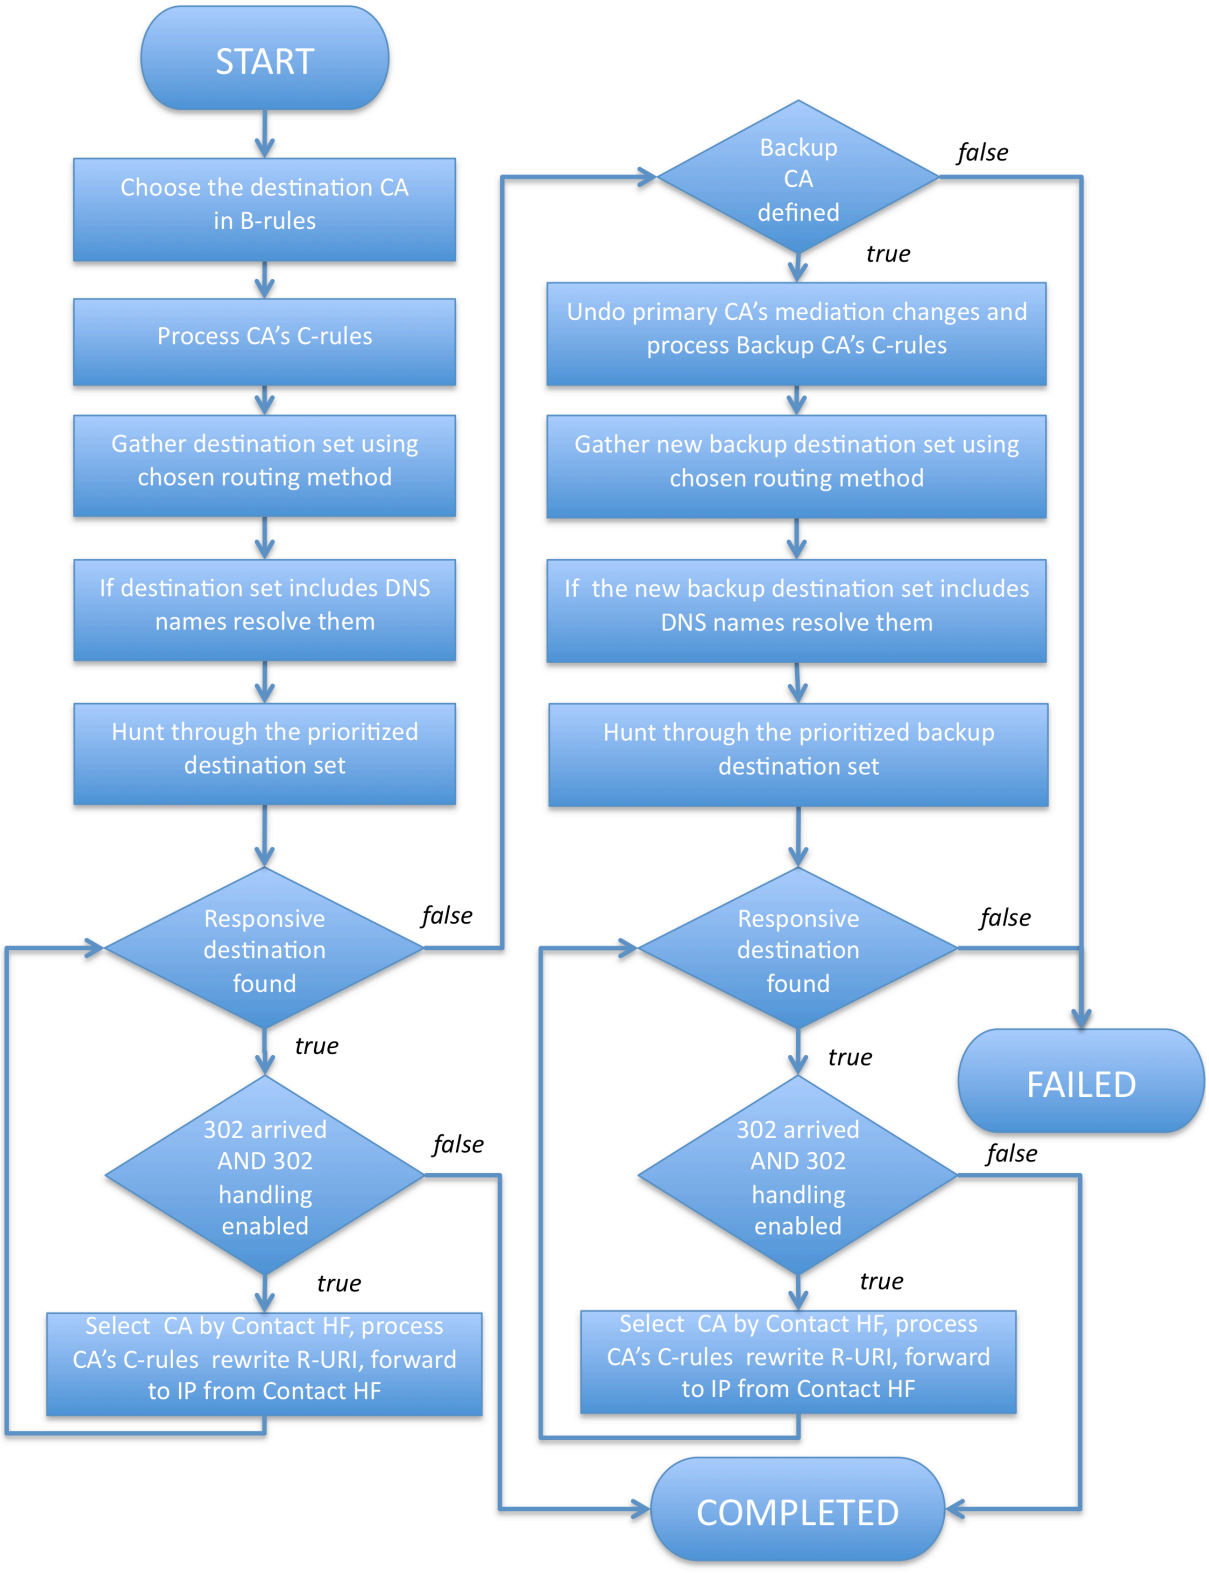

Figure 51: Flowchart of Process for Determination of the Next-hop IP Address

IP Blacklisting: Adaptive Availability Management

Attempts to forward traffic to IP addresses known to be unavailable would be futile and impair call setup time. Therefore the ABC SBC keeps a "destination blacklist" of IP addresses that were detected

as unresponsive. The ABC SBC dispatches no traffic to such destinations until the blacklisting time-to-live expires and the destination is removed from the blacklist.

Blacklisting is done when a normal SIP request to a destination fails. Additionally the ABC SBC can proactively probe destinations so that their unavailability is detected even before real traffic reaches them. Similarly their renewed availability is detected earlier thanks to the probes even while they are on the blacklist. This is called "Destination Monitor" or "OPTIONS monitoring".

OPTIONS monitoring can be enabled for any SIP Call Agents that are identified by IP addresses or DNS name. To turn it on, the "Monitoring Interval" under Call Agent's "Destination Monitor" options must be set to a non-zero value. The OPTIONS request are then sent in this interval periodically and have the following form:

```
OPTIONS sip:10.0.0.234 SIP/2.0
Via: SIP/2.0/UDP 10.0.0.155;branch=z9hG4bKo8lw1a70;rport
From: <sip:10.0.0.155:5060>;tag=b280210db5678d3c77dfc06c07acaac3
To: <sip:10.0.0.234>
CSeq: 32603 OPTIONS
Call-ID: 5F1BBCB9-57149447000B9232-0FF75700
Max-Forwards: 0
Content-Length: 0
```

On error, the destination address is placed on blacklist. If it is already there and the OPTIONS transaction completes successfully, the destination address is taken off the blacklist immediately.

Note that this type of blacklisting is different from that used in the context of security policies as described in the section Manual SIP Traffic Blocking.

To turn IP blacklisting on, set the time-to-live blacklisting period to a positive value under global options in "Config ▸ Global Config ▸ Default Destination Blacklist TTL" or under Call Agent properties as shown in Figure Configuration of Destination Blacklisting under Call Agent Properties.

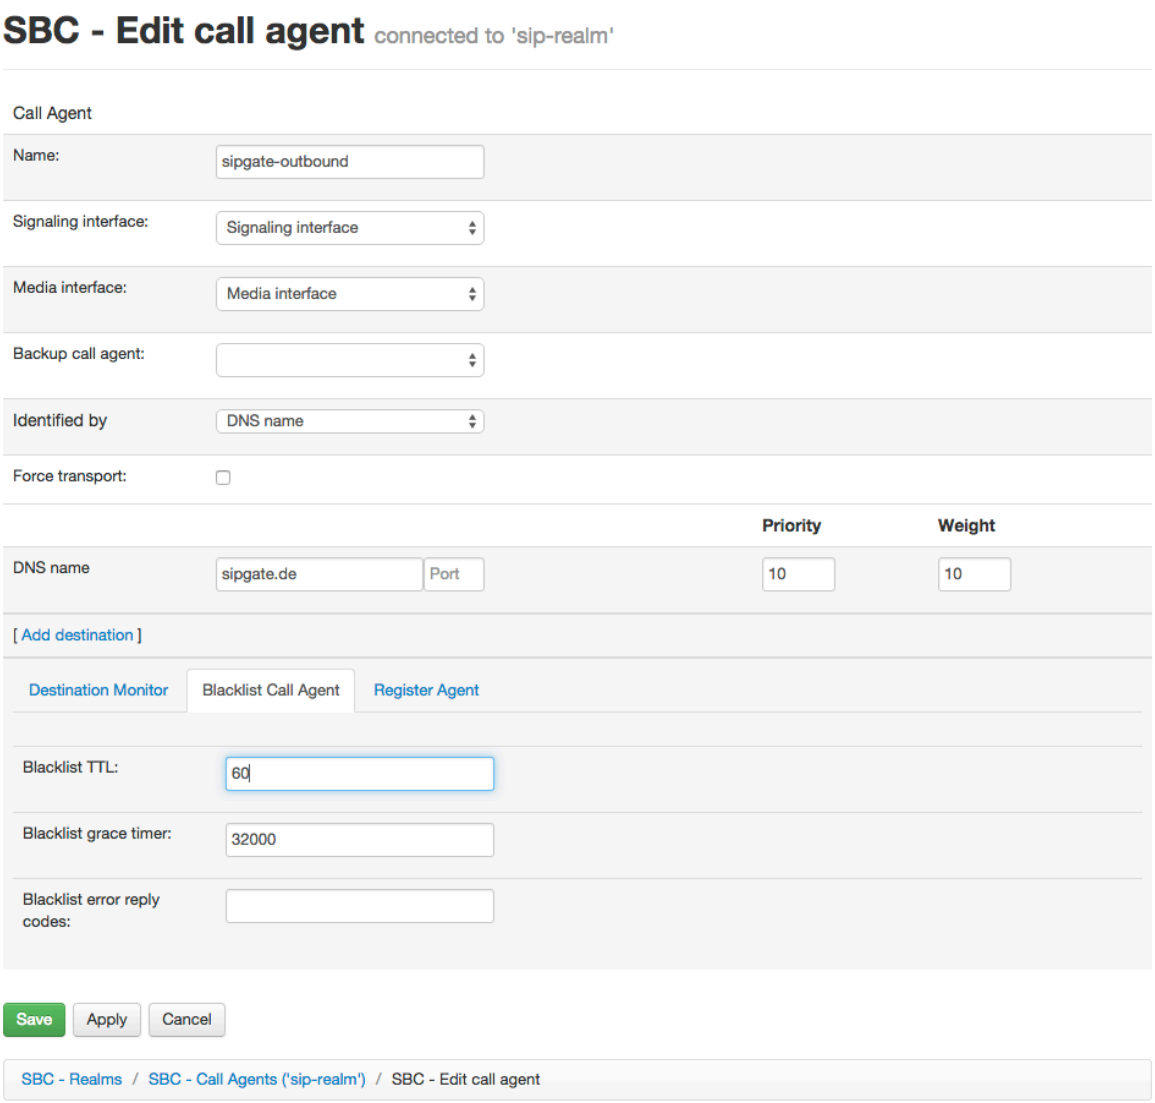

Figure 52: Configuration of Destination Blacklisting under Call Agent Properties

IP blacklisting occurs in an almost automated way and does typically require minimum administrative attention. Addresses are added to the blacklist once they are identified as unavailable and held on the list for a predefined period of time, known as "time-to-live". The following procedures may still be of use to an administrator:

- If ABC Monitor is used along with the ABC SBC, the history and status of the monitored Call Agents can be tracked in the "Connectivity CA" Dashboard.
- Monitoring blacklisted addresses. It is possible to inspect the addresses which are currently blacklisted. The list is available from the main menu under "Monitoring ▸ Destination Blacklist". (See Section Destination Blacklists)
- Manual blacklisting. The administrator may add a new address to the blacklist from the main menu under "Monitoring ▸ Blacklist ▸ New Destination/Save".
- Testing presence on blacklist in rules. Rule conditions may include a test if a Call Agent is present on a blacklist using the "Blacklist" condition type. The condition returns true of all Call Agent's IP addresses are blacklisted.
- Changing Time-to-Live (TTL). The addresses are held on blacklist for period of time specified under "Config ▸ Global Config ▸ Default Destination Blacklist TTL". This value is used for newly blacklisted destinations, unless a CA-specific TTL takes precedence. If TTL is set to zero, no blacklisting takes place.

- Configuring Call Agent specific handling. There are the following options available under Call Agent profile:
  - Destination Blacklist TTL (seconds). This value overrides the globally specified time to live.
  - Blacklist grace timer (ms). Normally, a destination is blacklisted when the transaction timer expires. This value provides some extra time before a downstream element is blacklisted after the transaction timeout. If the destination responds before the grace timer expires, then it is not blacklisted. That is especially useful when there is a proxy server between the ABC SBC and an unresponsive User Agent Server. A too aggressive blacklisting process would otherwise blacklist the proxy before it times out and sends a 408 message and make the proxy and elements behind it unreachable.
  - Blacklist error reply codes. This feature allows to blacklist destinations that answer to monitoring requests using these codes. When left empty, blacklisting happens when the reply code is 503. When set, blacklisting happens if the status code of a reply matches one of the codes provided in this parameter. To activate the feature, include a comma-separated list of response codes that lead to inclusion of a destination on a black-list. Blacklist error reply codes also controls whether to failover to backup CA. When blacklist error reply codes are left empty, failover happens:
    - When the destination responds with 503
    - when the reply is internally generated by the SBC (i.e. unable to resolve the destination address) and the generated reply code is not 483, 488, 400

    When blacklist error reply codes are set, failover happens:
    - When the destination responds with one of the reply codes in the list.

    When the reply is internally generated by the SBC and the code is 408, failover always happens regardless of the blacklist error reply codes field being set or not. Having reply codes 300, 301 or 302 in the list will not be effective as a means to failover to backup CA when redirect handling is active and reply includes redirect destinations. Blacklist error reply codes are also respected for failover during processing multiple ENUM query results.
  - Destination Blacklist for in-dialog requests. This feature allows to blacklist destinations during in-dialog requests. This can be used to allow in-dialog failover to another destination if the currently used destination becomes unavailable during a dialog and thus lands on the destination blacklist.
  - Monitoring interval (seconds). If set to a non-zero value, the ABC SBC tests availability of the destination by sending test message (OPTIONS). This allows to detect unavailable destinations even before a real call hits it. It is recommended to use a value shorter than the blacklisting TTL: if the monitoring period was longer, unresponsive destinations would be considered healthy in the time-window after removing from blacklist before the next monitoring check.

Whenever an address is added to the IP blacklist, an event of type ‘notice’ is generated. The same occurs when the TTL expires or the destination becomes responsive and the address is removed from the blacklist. The monitoring status is regularly reported to the ABC Monitor using “dest_monit” events.

SIP Routing by Example

One important configuration step is the definition of routing rules, i.e. to which IP address shall incoming traffic be forwarded. If a proper routing entry is not set up, the ABC SBC does not know where to forward traffic and returns the SIP reply “404 Not Found”.

In our example, the routing configuration is very simple: What comes from the external Realm is routed to the internal Realm and vice versa. That takes two rules defined from the “Routing” section of the web user interface: Traffic coming from the Call Agent “Proxy” is routed to the Call Agent “Users” in the external realm and traffic coming from the “external” Realm will be routed to the “proxy” Call Agent in internal Realm. The resulting configuration is depicted in Fig. Example Routing Rules.

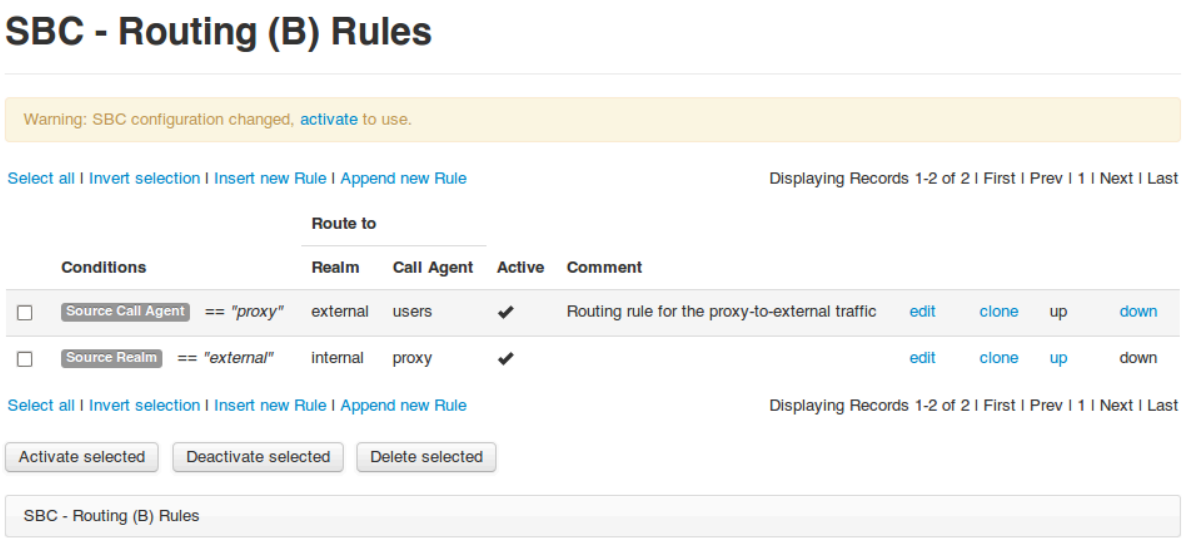

Figure 53: Example Routing Rules

When you define the respective routing destinations, specifying the abstract Call Agent may not be enough. You may additionally need to help the SBC to determine to which IP address to forward the SIP message and which hostname to use in the Request URI. For example, traffic leaving the SIP proxy carries the final destination in the Request URI. You must configure the ABC SBC to use the IP address from the request-URI as the next hop. The particular configuration is called **Route via R-URI**. On the other hand, all traffic from the public Internet goes to the same proxy server. The appropriate configuration choice is **Set Next Hop**. You may specify the IP address explicitly, if you do not do so, the IP address is taken from definition of the Call Agent.

The routing rule for the proxy-to-external traffic flow is shown in the Figure Routing Rule for internal to external traffic, whereas the rule for the opposite direction is shown in the Fig. Routing Rule for external to internal traffic. That's it. We now have the routing policy which specifies that traffic from the external Realm shall be forwarded to the internal Realm and vice versa. This simple policy can in fact serve an incredible number of use-cases.

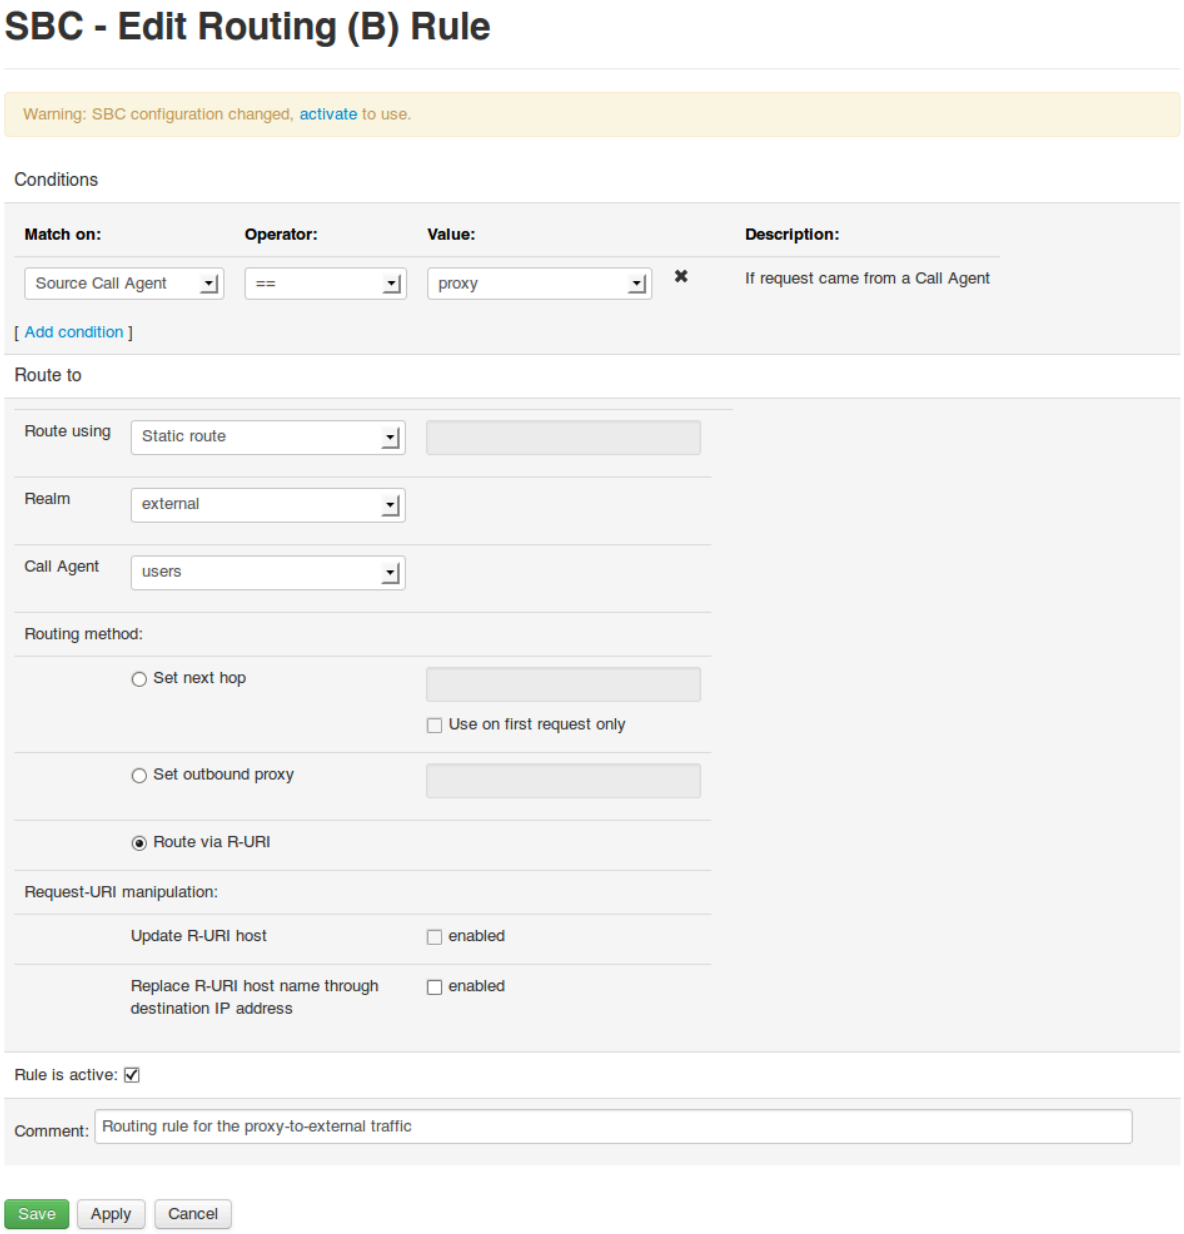

Figure 54: Routing Rule for internal to external traffic

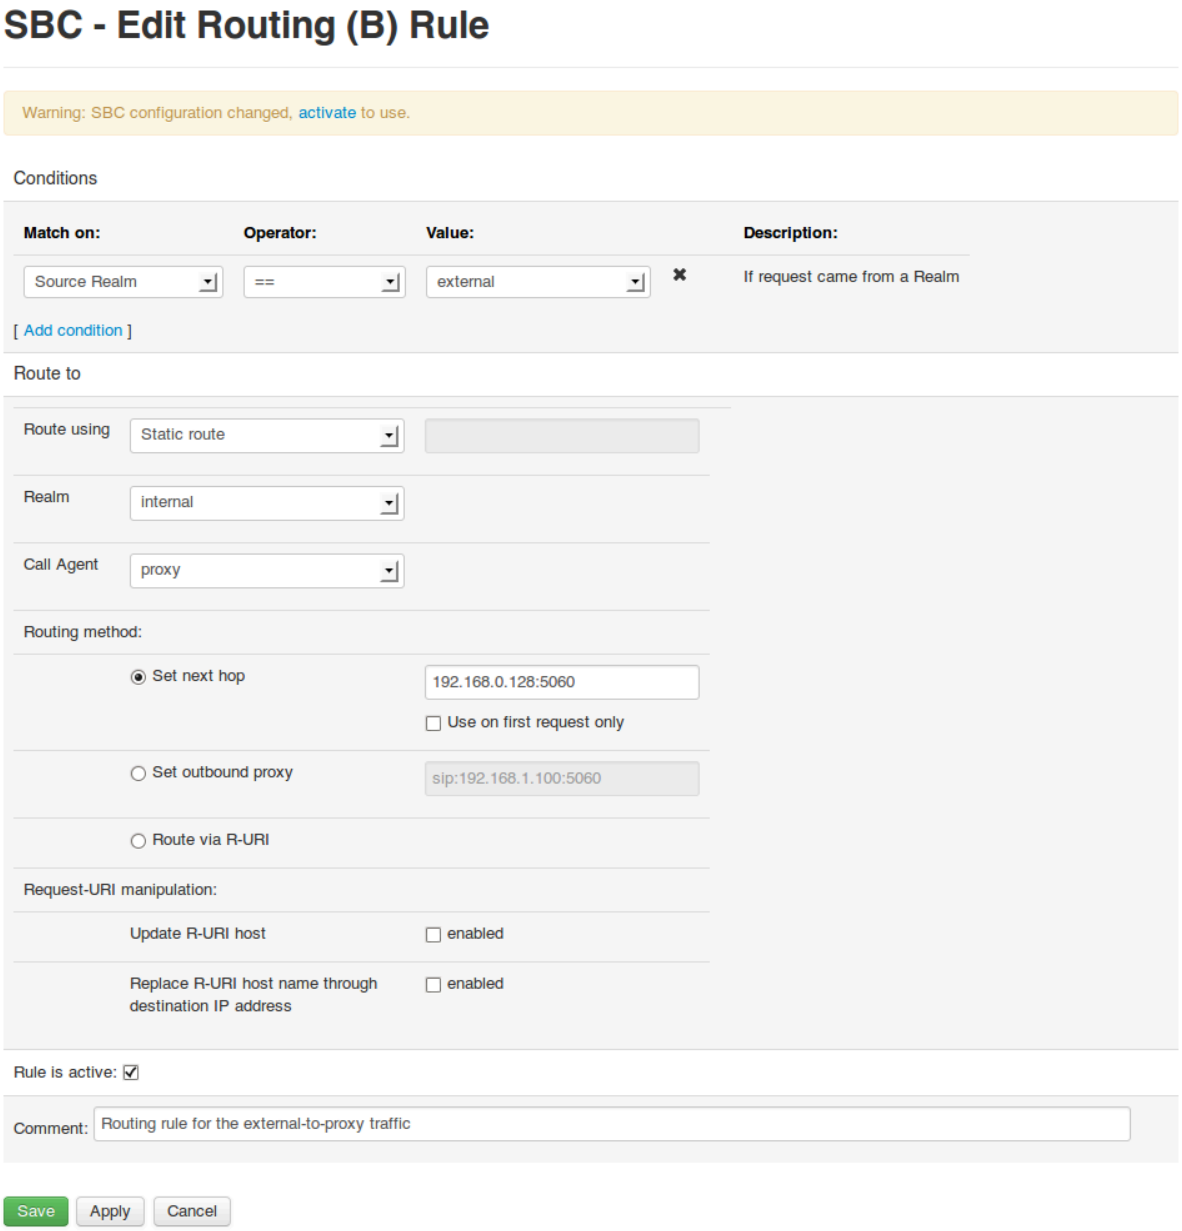

Figure 55: Routing Rule for external to internal traffic

Nevertheless, if needed it can use more sophisticated matching criteria to specify routing decisions: It can divert Message-Waiting-Indication traffic to a different server based on SIP method, different media servers based upon codecs in use, different destinations based on custom-defined header fields, and so on, and so forth.

4.10.7 View A-B-C rules

There is a possibility to view A-B-C rules together for particular SIP message, by clicking on the “ABC” icon for a routing rule on “Routing” screen:

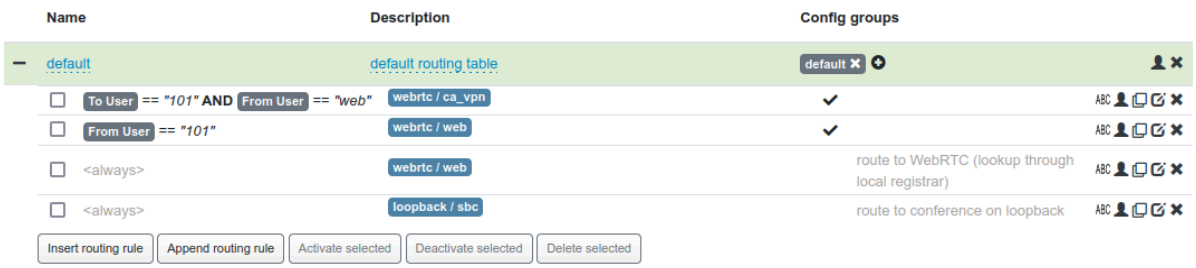

Figure 56: List of routing rules

By clicking that icon, new screen is displayed showing A and C rules and the selected routing rule.

If the routing rule contains “Source Realm” or “Source Call Agent” conditions, then this Realm / Call Agent is pre-selected in the top dropdown for A rules and A rules of this Realm / Call Agent are displayed in the upper part of the screen.

Likewise, if the routing rule use static routing and a Call Agent is selected as the route destination, this Call Agent is pre-selected in the bottom dropdown for C rules and C rules of this Realm / Call Agent are displayed in the bottom part of the screen.

4.10.8 SIP Mediation

SIP Mediation features of the ABC SBC allow administrators to introduce massive changes to the signaling protocol. This is often necessitated by devices with imperfect SIP support, differing practices such as dialing plans between peering providers, or need to implement network-based services such as Private Asserted Identity (**RFC 3325**).

The actual mediation rules are placed in inbound and outbound rules. The inbound rules are used to modify incoming traffic coming from a Realm or a Call Agent to comply to local policies. For example, the inbound rules may transform telephone numbers from a local PBX’s dialing plan to the global E.164 standard. All subsequent actions already work with modified SIP messages. The outbound rules are used to modify outgoing traffic to a form that the receiving Call Agent can or shall process. For example the outbound rules can remove all but low-bandwidth codecs for the target known to be on a low-speed link.

It needs to be understood that mediation is a double-edged sword: massive changes to the signaling protocol can, if not configured properly, cause substantial harm to interoperability. If the ABC SBC encounters, that a SIP message modified by mediation rules breaks standard too far (such as if it generates an empty header-field), it discontinues processing of the message and sends a 500 response back. Still many changes may be syntactically legitimate, remain undetected and result in impaired interoperability.

This section discusses mediation of the signaling protocol, SIP. Mediation of media, that includes codec negotiation and transcoding, is documented in the section Media Handling.

Why is SIP Mediation Needed?

There are multiple root causes why SIP devices have often troubles communicating with each other. There are different standardization groups working on SIP. Different developers often interpret the same specifications differently. Operators deploy different operational and naming practices.

The ABC SBC has the capability to overcome some of these interoperability problems by manipulating the content of SIP messages so that they better fit the expectations of the receiving side. One can distinguish between several frequent interoperability issues: compatibility between various SIP protocol extensions, dealing with deviations from the specification and best current practices caused by non-compliant devices and operating procedures, and incompatibility between different transport protocols used for conveying SIP signaling.

SIP Standard Extensions:

There are various flavors of the SIP protocol. Even the basic SIP IETF standard is extended by tens of accompanying specifications, some of them are deployed, some of them not. Several other standardization bodies have chosen to add even more extensions specific to their use of SIP. In the fixed environment, the TISPAN specifications are used. In the mobile network environment the 3GPP IMS specifications are the most favored. SIP-I is proposed for trunking scenarios in which SIP is used as the signaling protocol used to connect SS7 based networks over an IP core network.

The differences between the SIP specifications from IMS, IETF and TISPAN are mainly restricted to the addition of certain headers, authentication mechanisms and usage of certain SIP extensions such as NOTIFY/SUBSCRIBE or certain XML body formats.

In the context of interoperability of SIP flavors, the ABC SBC can provide the following services:

- Stateless SIP header manipulation: The ABC SBC can be configured to remove certain headers and add others. This way, The ABC SBC can for example delete headers that are useful in an IMS or TISPAN but not in an IETF SIP environment.
- Message blocking: Certain SIP messages might be useful in one network as they provide a certain service. However, if this service is not provided across the interconnection points then exchanging them across the networks does not make sense. SBCs can be configured to reject certain messages such as NOTIFY if presence services are not provided across the network for example.

Deviations from the SIP Standard and Best Practices:

The experience from various interoperability events shows that different vendors interpret the SIP specifications slightly differently. Especially parts that are specified with the strength of "SHOULD" or "MAY" are often implemented as a "MUST" or ignored completely. This makes the communication between two components from different vendors sometimes impossible. Sometimes even if the SIP equipment implements the standard correctly, operators practices for deploying SIP differ to the extent that the protocol needs to be fixed.

The ABC SBC can be configured to overcome some of these issues and to fix certain issues that cause these interoperability problems by offering the following features:

- Existence of certain headers: Some SIP components expect to see certain SIP headers with certain information, for example a *Route* header pointing to them. Others might not bother to add this header. The ABC SBC can be configured to take these special interpretations of the implementers into account before forwarding a request and add or remove problematic headers.
- Location of information: Some SIP components expect to see their address in the Request-URI whereas others want to see it in the *Route* header or both. This might not always be how the location information is included in the SIP request especially if a request was redirected from one component to another.
- Tags and additional information: Again some SIP components might expect to see certain tags and parameters attached to certain headers such as **rport** with a *Via header* whereas other SIP components might not add them.

SIP Transport:

SIP can be transported over UDP, TCP and TLS. The capabilities of different SIP implementations might vary with this regard. That is, some components could support UDP but not TCP and others prefer to use TLS. Therefore, the ABC SBC can be used to convert the transport protocol used by the source to the transport protocol preferred by the destination.

Default SIP transport for outgoing requests is **UDP**. This can be changed via one of the following:

- SIP Routing is done via the `Route via R-URI` method and the R-URI contains `transport` parameter,
- SIP Routing parameter `Force transport`,
- Call Agent configuration parameter `Force transport`.

Note that when forcing the transport via one of the `Force transport` configurations, the `transport` parameter in R-URIs will not be updated unless at least one of the following holds true:

- The routing method is `Route via R-URI`,
- R-URI `transport` parameter is set explicitly via `Set RURI parameter` action.

Request-URI Modifications

The most common manipulation is that of request-URI. Request URI describes who should receive the SIP request. It may include an E.164 telephone number (like `sip:+1-404-1234-567@pbx.com`), a PBX number (`sip:8567@pbx.com`) or be formed as an email-like address (`sip:amadeus@mozart.at`). A typical reason for changing the request URI is normalization of different dialing plans. As an example you may translate a local extension number (`768`) for a PBX with prefix (`+1-404-1234`) into a globally routable E.164-based URI `sip:+1-404-1234-567@national-gateways.com`. You can use several types of modifications to the request-URI, all of them are applied only to the first session's request.

The most important request-URI actions are the following:

- **Strip RURI user** – strips the specified number of leading characters from the user part of request URI. For example "Strip RURI user" applied to the PBX URI `8567@pbx.com` yields the extension `sip:567@pbx.com` without the local "8" prefix. The action is applied as many times as it is called.
- **Prefix RURI user** – inserts a prefix to the user part of request URI. For example, "Prefix RURI user" with `+1-404-1234-` applied to the URI from the previous step yields `sip:+1-404-1234-567@pbx.com`. The result is accumulated if the action is applied several times.
- **Append to RURI user** – appends a suffix to the user part of request URI. The parameter takes suffix value. It may include replacement expressions. The result is accumulated if the action is applied several times.
- **set RURI** – entirely replaces the request URI with a new value.
- **set Contact URI host** – entirely replaces the Contact URI host with a new value. Note that the update isn't run on REGISTER replies.

Also note that the resulting URI not only describes the recipient, but its host part is used to determine the next hop IP address if a route is used with the **Route via R-URI** option.

It is also worthwhile mentioning that URIs often represent additional services a caller gets. For example if a caller prefixes number of an O2 subscriber in Germany with 33, his call will be directly routed to the recipient's voicemail. However administrators would be ill advised to overload request URI with more than routing functionality. An infamous example is using a plain-text password as phone number prefix for authentication. The fraudster Edwin Pena found that out, yielded more than 10 million minutes of VoIP service and in 2009 eventually two years in federal prison.

Several other mediation actions can process sub-parts of request-URI. They include:

- **Set RURI host**
  - Replace host(:port) part of Request-URI with a new value specified in the GUI.
  - Parameters: new host or host:port
- **Set RURI parameter**
  - Add or replace parameter of Request-URI.
  - Parameters: RURI parameter name, RURI parameter value
- **Set RURI user**
  - Replace user part of Request-URI with a new value.
  - Parameters: new user part.
- **Set RURI user parameter**
  - Add or replace parameter of user part of Request-URI.
  - Parameters: parameter name, parameter value.

Changing Identity

Identity of SIP session participants is also described in many other SIP header fields that sometimes need to be changed.

Every SIP request must include URIs of session initiator in the *From* header-field and URI of intended recipient in the *To* header field. The SIP standard has intended to use the *From* and *To* header field only as informational description of how a session was started . URI of the originator in the *From* header field has limited identity value as the plain-text URI is not covered by a message integrity check and can be easily changed by elements in the SIP-path. Even a user client is quite free to put anything in the URI unless there is a client's outbound SIP proxy enforcing specific address for a digest-verified caller.

The URI in *To* header-field may have little relation to the actual recipient of a SIP request as the actual next hop is stored in the request URI.

Notwithstanding how "light-weight" information *From* and *To* header fields convey, some operators deploy policies based on them. They may only accept requests with *From* and *To* URIs that comply to their local convention. There were even cases when *To* URI was used for routing. Therefore it is often useful to modify *To* and *From* header fields. These modification rules apply to the first request of a SIP dialog. *From* and *To* in all subsequent messages of a session are transformed transparently in compliance with the SIP protocol specification.

The most important *To* and *From* changing actions are the following:

- **Set From / Set To** :replaces the whole *From/ To* header field with a new value, for example "Jasmine Blue" sip:jasmine@blue.com. Only "tags" in the *From/ To* header-fields remain unchanged to guarantee unique identification of SIP dialogs.
- **Set From User / Set To User**: replaces the user part of the From/To URI.
- **Set From Host / Set To host**: replaces the host(:port) part of URI with a new value
- **Set From / To display name**
  - Set only the display name of the *From* / *To* header.
  - Parameter: new display name.

Additionally, the SIP protocol is using digest authentication identity (**RFC 2617**) to verify who is initiating a request. If the digest identity of a request originator needs to be changed, the action **UAC auth** is used. It takes the following parameters needed for the authentication procedure: username, password and realm. A request forwarded downstream and challenged to authenticate by a downstream server is then resubmitted by the ABC SBC using these credentials. Note that the input fields support replacement expressions. If i.e. password contains special characters such as `$`, they need to be escaped with a backslash.

Substitution Expressions

SIP message modifications typically "glue" pieces of the original messages and intended changes. For example, a new *To* URI is to be formed using destination's hostname (say "target-gw.com") and telephone number in request URI (say "+1-404-1234-567"). The corresponding **set-To** action needs to access the telephone number in the original request. To address cases like this, the mediation parameters may refer to elements of the original message by so called Replacement expressions. These always begin with a **$** character.

In our example, the user part of the request URI is referred to as "**$rU**" and the action has the form: **Set To** `sip:$rU@targetgw.com`.

Other important replacement expressions are **$fu** for *From* URI, **$tu** for *To* URI, **$si** for source IP address, **$H(headername)** for value of a header field.

If you need to access some sub-parts of the original SIP message without an addressable name, simple substitution expression are not enough. Then regular expressions have to be used to select them. This is called „regexp backreferences. The backreference expressions refer to parts of SIP messages that were matched in rules' conditions. For example, to access the protocol discriminator in a URI, you need to create a rule condition matching it using regular expression, and then refer to the matched expression. You would be forming a rule like this:

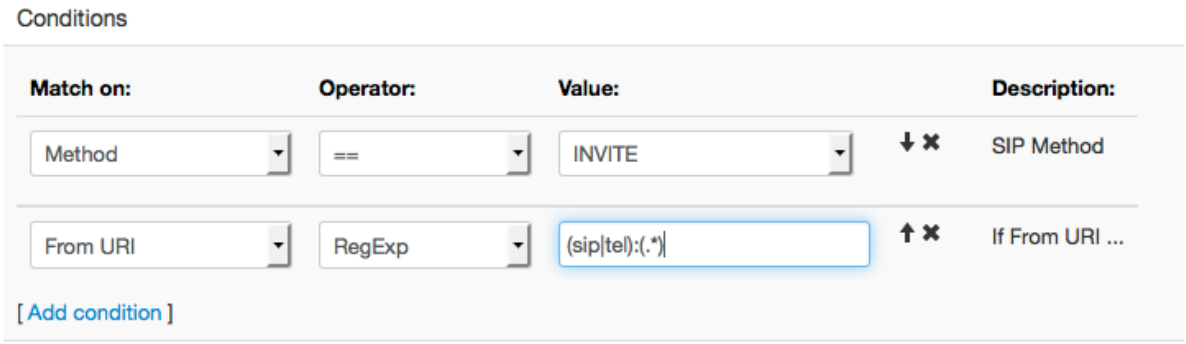

Figure 57: Example of a Condition Being Referred to by a Backreference Expression

the second condition's first sub-part (i.e. matched by the expression in the first parentheses) of the regular expression would identify the protocol discriminator and yield "sip" for SIP URIs. The expression would be formed as this

```
$B(2.1)
```

SIP Header Processing

URI adaptation shown in previous paragraphs is important for harmonization or routing and identity representation between different SIP devices and administrative domains. Yet there are many other header fields conveying important information in need of adaptation. Worse than that, some of them are not even known at the time of writing this documentation. That's because some of them may be proprietary – for example Sipura SIP phones add QoS reports to every BYE message they send. Some header fields may even be specified in recently published standard. Yet even then the ABC SBC can help – it can use general purpose text-processing methods thanks to SIP's text-based nature. Particularly the following actions are available:

- **Remove Header:** Remove-header removes all occurrences of a header-field identified by its case-insensitive name from all requests and responses in a session. Exceptions apply: mandatory header-fields are not removed: Call-ID, From, To, CSeq, Via, Route, Record-Route and Contact. If a header-field with compact name form occurs, both forms must be removed explicitly. Newly added header-fields are not removed by this action.
- **Set Header Blacklist:** is a convenience function removing multiple header fields by a single action. It takes comma-separated list of header-field names as parameter and achieves the same effect as if you used multiple occurrences of the Remove-Header action. Blacklists are applied one by one in the order in which they appeared in the rules and are executed after applying both A and C rules. Blacklists can be a nice short-cut for removing a header-field which has both normal and compact name. For example, you may want to configure deletion of both forms of the Subject header field by using

  ```
  Set-Header-Blacklist("Subject,s")
  ```

- **Set Header Whitelist:** is an even more aggressive convenience function for removing multiple header fields. If used, all but mandatory and whitelisted header fields are removed from all requests and responses belonging to a session. The action is applied after processing of both A and C rules completes.
- **Add Header:** adds a new arbitrary header-field to a dialog-initiating request. This action only applies to the first request of a session. Its greatest power comes from the ability to craft complex header-fields using the substitution expressions.

SIP Header Modification Examples

Let us show the power of these actions on an example. A real-world case is translation of identity between the pre-standard *Remote-Party-ID* header field, still used by some SIP equipment, and the standardized *Asserted Identity*, see **RFC 3325**. Both fulfill the same purpose, yet differ in their syntax which needs to be translated from one form into the other.

The pre-standard header-field looks like this

```
Remote-Party-Id: "Mr. X" <sip:+1-404-1234-000@sipsip.com>;privacy=full
```

The standard form looks like this

```
P-Asserted Identity: "Mr. X" <sip:+1-404-1234-000@sipsip.com>
```

The simplest way for translation is

- finding out if there is an occurrence of *Remote-Party-Id* header-field by a rule with condition

  ```
  if Header(Remote-Party-Id) does not match RE ^$"
  ```

- removing the header field by action

  ```
  Remove-Header(Remove-Party-Id)
  ```

- and eventually forming the newly crafted header field using the URI in the previous header-field by

  ```
  Add-Header(P-asserted-identity: $Hu(remote-party-id)"
  ```

Note that while this example mostly works, it ignores some parameter details for sake of brevity.

It is important to keep in mind that mediation changes have impact on subsequent SIP processing: replacement expressions and header-field tests in condition consider the changed value.

The following example rules change request URI and From URI.

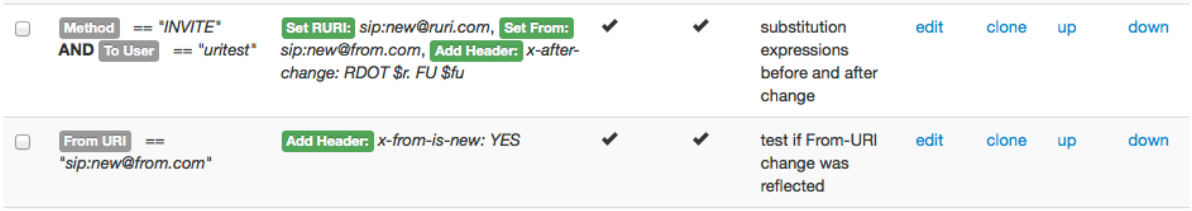

Figure 58: Impact of Mediation Changes on Subsequent Processing

Substitution expressions in the *Add Header* Action print the request URI and From URI in a troubleshooting header-field named *x-after-change. Because they refer to the **current** value, the new URIs appear in the outgoing INVITE regardless what they included originally. Similarly, the header-field value condition that tests if the From URI has assumed the new value prints YES in another troubleshooting header-field name *x-from-is-new*:

```
INVITE sip:new@ruri.com SIP/2.0.
Via: SIP/2.0/UDP 192.168.0.84;branch=z9hG4bKtwvtoaKk;rport.
From: <sip:new@from.com>;tag=170A7805-533943A10009B434-D85F9700.
To: <sip:uritest@abcsbc.com>.
CSeq: 10 INVITE.
Call-ID: 77443978-533943A10009B47B-D85F9700.
User-Agent: Blink Lite 3.1.1 (MacOSX).
x-after-change: RDOT sip:new@ruri.com FU new@from.com.
x-from-is-new: YES.
....
```

Option tags

Option tags are unique identifiers used to designate extensions in SIP. These tags are used in Require and Supported header fields.

To simplify manipulation with these headers ABC SBC offers since version 4.5 following conditions:

- **Supported header**: allows to check whether an extension is (is not) present in Supported header field.
- **Require header**: allows to check whether an extension is (is not) present in Require header field.

and actions:

- **Update Supported header**: allows to add option tags (*Add tags*) or remove them (*Remove tags*) from Supported header field or overwrite them completely (*Set tags*)
- **Update Require header**: allows to add option tags (*Add tags*) or remove them (*Remove tags*) from Require header field or overwrite them completely (*Set tags*)

Early Media, Ring Back Tone and Forking

In SIP, so called "early media" and "forking" are quite complex SIP features which make interoperability sometimes a challenge, especially when occurring together.

Early media appeared in the SIP protocol as a PSTN backwards-compatibility feature. In PSTN the difference between early media like "please wait your call is important to us" and the actual call is simple: the latter is charged for, the former is not. This is by the way the reason, why "early media" is sometimes humorously referred to as "late charging". Early media appear often when the called party is a PSTN gateway. The same protocol vehicle is often also used to implement "ring back tone". The protocol flow is rather simple: The callee sends a provisional response with a reply code equal to *180* or *183* including an SDP answer and starts sending RTP with the ring back tone to the caller. Usually, the caller User Agent only starts rendering the ring back tone to the user when this response is received. The protocol usage examples and details are well described in the **RFC 3960**.

Forking is a feature anchored in the SIP specification **RFC 3261**. It permits SIP proxy servers to forward one incoming requests to multiple different destinations. For example, one can setup this way all his phones to ring in parallel: soft-phone, hard-phone, smart-phone and even a PSTN phone behind a PSTN gateway. Forking can occur in parallel or in series. If serial forking is used, a forking proxy following best current practices sends a 181 inbetween. The "forked" INVITE requests may look almost identical but each of them always must have a unique "branch" identifier in the topmost Via header field.

Various unpredictable situations appear when forking and early media appears at the same time. For example two PSTN gateways send both early media to the caller. To deal with such situations the ABC SBC does only accept the first early media stream and discards the subsequently received ones.

The actions described in this Section help to customize the behavior of the ABC SBC to some special cases.

The action **Drop SDP from 1xx replies** drops SDP payload from all listed 1xx SIP answers. The action takes as parameter a comma-separated list of reply codes. SDP payloads are dropped from all responses with any of these codes. This action is especially useful if specific replies should be handled,

for example a locally generated ring back tone should be preferred to a ring back tone from the far end. Note that the RTP relay is not started if all provisional response are dropped, i.e. a provisional response needs to be processed for the RTP relay to be initialized, also for relaying early media.

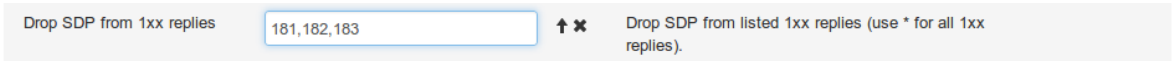

Figure 59: Drop SDP from reply

Another action **Drop early media** drops the RTP packets of early media, that is until the call is established. Note that if early media shall be dropped from signaling entirely, the actions "Drop SDP from 1xx replies" in combination with "Translate reply code" 183->180 must be used.

Figure 60: Drop early media

Support serial-forking proxy: This action allows to reset early media when a downstream SIP proxy server indicates by a 181 response that it has chosen to try some other destination for the call. By default, only the first early media arriving to the SBC is permitted, all other early media is dropped. This strict policy assures that downstream SIP forking cannot create multiple early media streams mutually interfering with each other. With this option, one can make an exception to the rule and permit early media coming later to override the previously established early dialog. It works safely as long as there is no parallel early media and 181 indicates that a later early media stream legitimately replaces the previous stream.

The following SIP flow-chart from Section 2.9 of **RFC 5359** show a situation in which a SIP proxy generates a 181

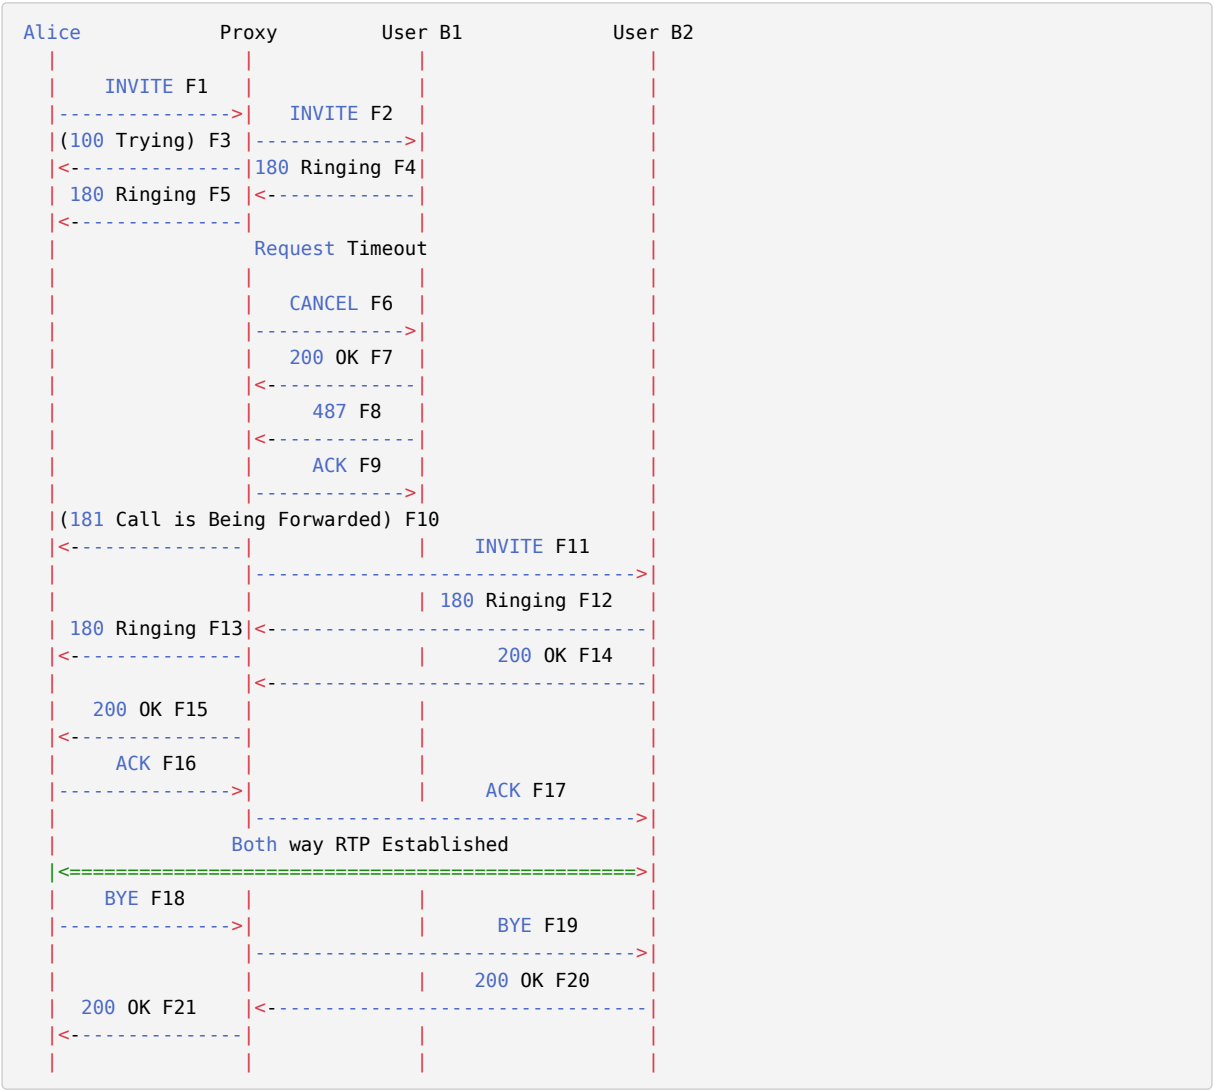

The action **Fork** allows to add a new branch to a processed request and start forking. Multiple occurrences of the action result in multiple branches of the request. The action takes only one parameter, the request URI of the forked request. The parameter can use replacement expressions, however if an invalid SIP URI is formed the call will fail.

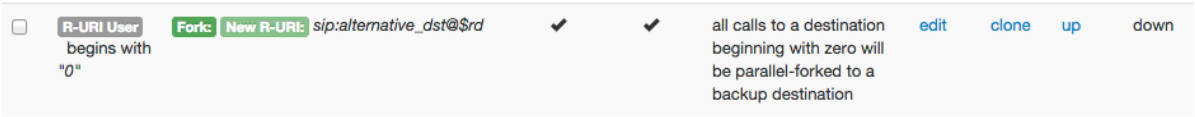

Figure 61: Forking

Call transfers

Using the action **Call transfer handling** it can be configured how in-dialog REFER requests are handled in the ABC SBC. The configuration is per call leg, i.e. if used in inbound (A) rules REFER handling is set for the A leg, if used in outbound (C) rules it is set for the B leg.

Following methods of REFER handling can be used:

- **pass-through** – Pass REFER through the ABC SBC to the remote peer (default).
- **reject** – Reject the REFER request with a `403 Forbidden` reply.
- **handle internally** – In case of an attended call transfer to another call established through the ABC SBC (REFER with `Replaces` in `Refer-To` pointing to a local call) the call legs are connected

locally. Only offer-answer exchanges (re-INVITEs) that synchronize session description on both ends are generated. In case of an unattended call transfer (no `Replaces` in `Refer-To`) the ABC SBC generates a new INVITE to the requested destination. This INVITE can be handled in routing (B) rules and outbound (C) rules similarly to regular calls. For detection of such locally generated calls the condition **Request source** can be used. In case of an attended call transfer to a non-local call (`Replaces` in `Refer-To` refers to a non-existent call leg) the ABC SBC generates a new INVITE with `Replaces` to the requested destination. This INVITE can be handled the same way as an INVITE generated for an unattended call transfer mentioned above.

Limitations:
- only in-dialog REFER requests are handled
- attended call transfer is not possible with transparent call IDs

INVITE with Replaces handling

ABC SBC is able to handle INVITE with Replaces header locally, if the Replaces header points to a call established on the SBC.

The action **Handle INVITE with Replaces header** is used for this purpose - it activates local INVITE with Replaces handling.

Limitations:
- INVITE with Replaces can not be handled when replacing call with transparent call IDs

Mapping Dialog-IDs in INVITEs with Replaces

If an INVITE with Replaces passes the ABC SBC, and the call to be replaced is also traversing the SBC, with transparent call IDs not enabled the Dialog Identifiers in the Replaces header refer to the call leg on the side before the SBC, but do not have a meaning after the SBC.

Using the *Map Replaces header* action, the dialog identifiers are replaced with the corresponding ones on the other side of the SBC so that the Replaces still is valid.

Other mediation actions

The ABC SBC supports various actions related to SIP processing:
- **Enable transparent dialog IDs**
  - Use the same dialog identifiers (Call-ID, From-tag, To-tag) on both sides of a call (e.g., for the incoming and out going messages). If this action is not enabled, the FRAFOS ABC SBC changes dialog identifiers. Unchanged Call-ID may be a security concern because it may contain the caller's IP address. However, transparent identifiers make troubleshooting and correlation of call legs much easier. Also, for call transfers using REFER with *Replaces* as used in call transfer scenarios to work through the SBC, transparent dialog IDs need to be enabled.
  - Transparent dialog IDs should be avoided unless absolutely necessary. It is known to break unattended call transfers with "call transfer handling" action.
- **Forward Via-HFs**
  - This option makes the SBC keep all *Via* headers while forwarding the request. This behavior mimics what a proxy would do, especially in combination with the **Enable transparent dialog IDs** action (the only remaining difference to a proxy is the non-transparent *Contact* header field). Note that forwarding the *Via* headers exposes the IP addresses of entities on the incoming leg side of the request.
- **Translate reply code**
  - Change SIP response code and reason for all SIP responses with a specific code. Note that changing responses between SIP reply classes may seriously break proper operation.
  - Parameters: SIP response code to change, SIP code and reason phrase to use for the reply sent out.
- **Allow unsolicited NOTIFYs**

  - The ABC SBC keeps track of subscriptions and usually only lets NOTIFY messages through if a subscription for it has been created before (through a SUBSCRIBE or a REFER). This action tells the SBC to let pass NOTIFY messages even if no subscription has been created before.
- **Relay DTMF as AVT RTP packets (RFC4733/RFC2833)**
  - relays DTMF tones as RTP avt-tones packets (**RFC 4733**/ **RFC 2833**)
  - Parameters: none
- **Relay DTMF as SIP INFO**
  - relays DTMF tones as proprietary SIP INFO payload
  - Parameters: none
- **Diversion to History-Info**
  - converts SIP diversion header-field (**RFC 5806**) into the History-Info header-field (**RFC 4244**) using the guidelines set in **RFC 6044**.
  - Parameters: none
- **Set Max Forwards**
  - sets the value of Max-Forwards header field in forwarded SIP requests to the configured value. This limits the number of hops a request can be forwarded until it is bounced back. It may make sense to set it to a lower value than **RFC 3261** recommends (70). If this action is not used, the value in incoming request is decremented by one before forwarding.
  - Parameters: number of hops.
- **Set Content Type whitelists** and **Set Content Type blacklists**
  - limits SIP content types to well known payload types (whitelists) or to all but specifically prohibited payload types (blacklists). Most VoIP SIP requests include the type of `application/sdp`.
  - Parameters: comma-separated list of content-types
- **Add dialog contact parameter**
  - Allows to add a parameter to the contact URI generated by the SBC.
  - Parameters: The side of the call (caller/A leg, callee/B leg) can be specified, the parameter name and value.
- **Set Contact-HF parameter whitelists** and **Set Contact-HF parameter blacklists**
  - defines which Contact HF parameters are forwarded through the ABC SBC . By default no parameters are forwarded. With whitelisting, only specified parameters are forwarded. With blacklisting, all but specified parameters are forwarded.
  - Parameters: comma-separated list of Contact parameter names
- **Forward Contact-HF parameters**
  - makes sure all Contact HF parameters are forwarded as received in incoming request. If no action is used, no parameters are forwarded at all.
  - Parameters: none
- **Call transfer handling**
  - The actions defines in which mode incoming REFERs will be processed. They are either rejected, forwarded or handled locally.
  - Parameters: REFER-processing mode

4.10.9 SDP Mediation

SDP mediation allows to manipulate how applications codecs will be selected during session negotiation.

Codec Signaling

In SIP call parties are free to negotiate their capabilities using the offer-answer model, see **RFC 3264**: The caller offers its capabilities such as supported codecs and the caller party matches those against its own. In some cases it may be reasonable to restrict the list of offered codecs. Mostly, this is done when there are bandwidth constraints.

If media anchoring is used, every single media stream enters and leaves the SBC. With the most common but "hungry" codec G.711, it means 172 kbps x 2 in each direction, which corresponds to a maximum of about five thousand calls on a gigabit link (the actual limit is in fact even lower due to packet rate constraints).

G.729 is probably the most widely used codec with lower bandwidth consumption. The bit rate for G.729 yields 62 kbps in each direction (the rate includes UDP, IP and Ethernet overhead).

For mobile clients, bandwidth hungry codecs with large packet size like G.711 can pose additional problems: Due to longer use of the wireless interface, battery life is reduced, and also the packet loss rate is greatly increased with the bigger packet sizes. On the other hand, CPU intensive codecs may also strain the battery on mobile clients if they are not implemented in hardware.

For these reasons, the ABC SBC offers you these functions

- setting codec preferences (**Set codec preference** action). Specifies in descending order which codecs offered in SDP payload should be "picked".
- transcoding (**Enable transcoding** action) – allows to convert sender's media from encoding the received does not support to encoding he does. See more in Section Transcoding.
- codec white/blacklisting (**Set codec whitelist** and **Set codec blacklist** actions) explicitly specifies which codecs are permitted or not.

In order to save bandwidth and improve battery life and call quality, to mobile clients G.711 should not be used by using a **Set codec blacklist** action with "PCMU,PCMA" as blacklisted codecs.

Another example is emergency calls (911), where due to call quality concerns G.711 is the mandatory codec. If, for bandwidth saving reasons, the G.711 codec is usually blacklisted, it should be whitelisted for calls sent to an emergency gateway.

If codec restrictions result in a failure to find a common codec, the ABC SBC offers you to use built-in software based transcoding to increase interoperability.

Please refer to the Media processing section, see Sec. Media Handling for a complete reference of functionality the ABC SBC offers to restrict the set of used codecs, give certain codecs a preference or transcode between codecs.

Media Type Filtering

For an audio call, the media type in the SDP is "audio". For a normal video call with audio, the two media types "audio" and "video" are negotiated, for other types of calls ("image", screen sharing etc), other media types are possible.

Media types may be filtered using the actions

- **Set media blacklist** – remove all blacklisted media types from SDP
- **Set media whitelist** – remove all but whitelisted media types from SDP

The media blacklist/whitelist actions have as parameters a comma-separated list of media types (audio, video, image, …) to be blacklisted. It is applied to all SDP messages exchanged at any time during the call. In the case that after applying the action no media type is left in the SDP message then the request will be rejected with a response message *488*.

Example: Allow only audio payload to pass and prevent video streams to be negotiated; for the User Agents it will appear as if the other side does only support audio.

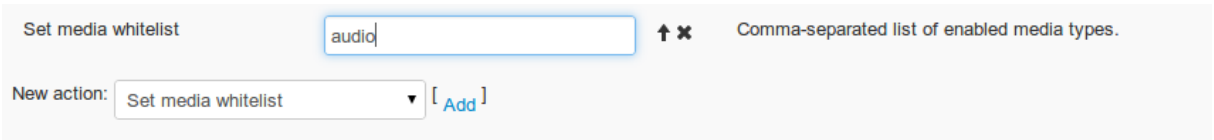

Figure 62: Remove video streams

Example: Let audio and audio/video calls through.

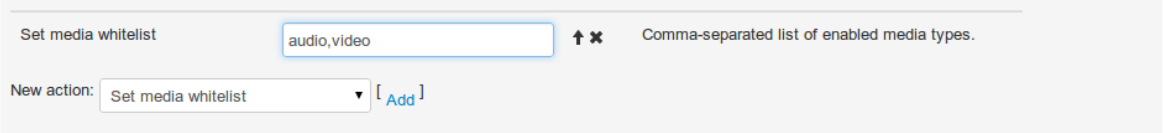

Figure 63: Allow only audio and video

Example: Remove "image" media type.

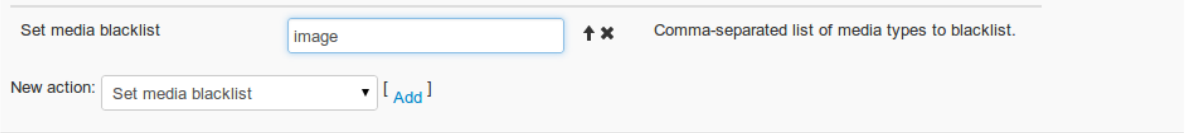

Figure 64: Do not allow exchange of images

CODEC Filtering

The actual audio or video content of a call can be encoded with different codecs, which have each different properties regarding:

- audio or picture quality
- bandwidth consumed
- latency introduced
- processing power required
- resilience regarding packet loss

For example, the G.711 codec has "toll quality" (audio quality roughly equivalent to PSTN, 8khz sampling rate/narrowband) at 64kbit/s (roughly 80 kbit/s including headers in each direction), introduces low latency, requires little processing but is not resilient against quality degradation with packet loss. The G.729 codec has a bit less than "toll quality" at 8kbit/s, 6.4kbit/s or 11.8 kbit/s (depending on the used annexes) with modest latency introduced, some processing power required, and some resilience against packet loss.

In order for two endpoints to successfully establish a call, both endpoints need to support the same codecs. The codecs actually used in a call are negotiated using the SDP protocol and the SDP offer/answer method to the subset of codecs supported by both endpoints, and thus it is usually best to let the endpoints negotiate with the most options possible.

If for some reasons codecs need to be filtered, the following actions can be used:

- **Set CODEC whitelist** – remove all but whitelisted media types from SDP
- **Set CODEC blacklist** – remove all blacklisted media types from SDP

Each of these actions takes a comma-separated list of codecs to white- or blacklist, which must be the names of the codecs as they are used in SDP[1] . Codec names are case-insensitive, a blacklist of "g729,ilbc" is equivalent to "G729,ILBC". In the case that after applying the action no codecs are left in the SDP message then the request will be rejected with a response message *488*.

Figure 65: Setting a whitelist

[1] e.g. PCMU for G.711 u-law, PCMA for G.711 A-law. See http://www.iana.org/assignments/rtp-parameters and the IETF payload type specifications (RFCs) for the used names of the codecs.

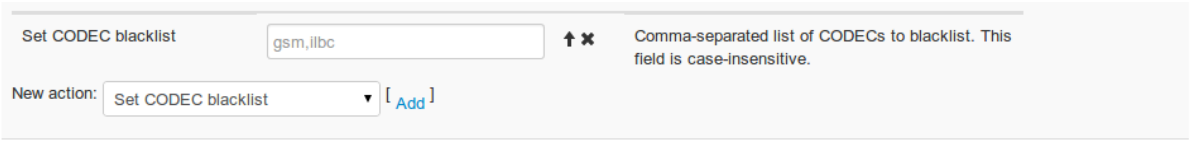

Figure 66: Setting a blacklist

All of the white- and blacklists applied for a call are executed one after another. For example, if one action sets the whitelist "PCMU,PCMA,speex" and another action sets the whitelist "PCMA,G729", it would result in only PCMA (G.711 A-law) be let through for the call.

The codec white and blacklists are applied on both legs, the incoming and the outgoing call legs, and on all SDP messages going in both directions (from caller to callee and from callee to caller) at any time during the call.

CODEC Preference

Codec negotiation in a SIP call usually works this way

- the offerer (the caller in calls with normal SDP offer-answer negotiation; the callee in calls with delayed SDP offer-answer negotiation) lists the supported codecs for a call in preferred order, so the first codec in the SDP is the codec preferred by the offerer
- the other party (the answerer) selects from this list the subset of codecs that it supports, and orders it according to its preference or local policy, it may for example accept the order that the offerer asked for
- both parties encode and send media with the codecs that are in this subset, and usually the first codec in the answer is used, but (even without re-negotiation) any party may switch in-call to any codec in this subset

Because the User Agents usually respect the codec preference of the other side, the operator may influence the codec actually used for a call in the SBC by reordering the codecs as they are listed in the SDPs. Codec preferences may be influenced in order to

- improve audio quality by preferring better performing codecs
- save bandwidth
- save processing power on the User Agents, e.g. especially if mobile/battery powered devices are used

The **Set CODEC preferences** action has as parameters two comma-separated lists of codecs, for the A (caller) leg and the B (callee) leg respectively. An entry may have the specific sample rate appended separated with a slash, e.g. speex/32000, speex/16000, speex/8000

Figure 67: Setting the codec preferences

Any codecs in this list found in the SDP messages exchanged in either direction are prioritized in the order listed, by placing them at the beginning of the codec list in the SDP document. It is a good practice to configure the same order for both legs.

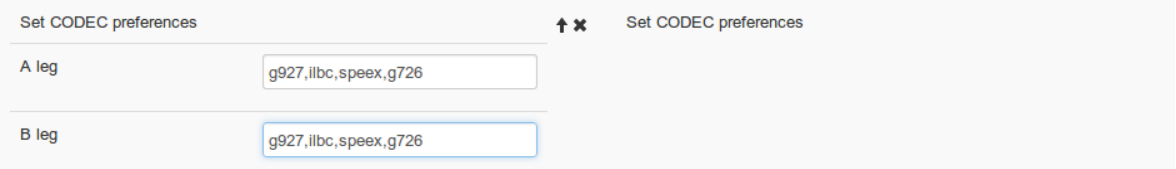

Figure 68: Example: Prefer bandwidth-saving codecs (codecs that compress more): G.729,iLBC,speex,G.726

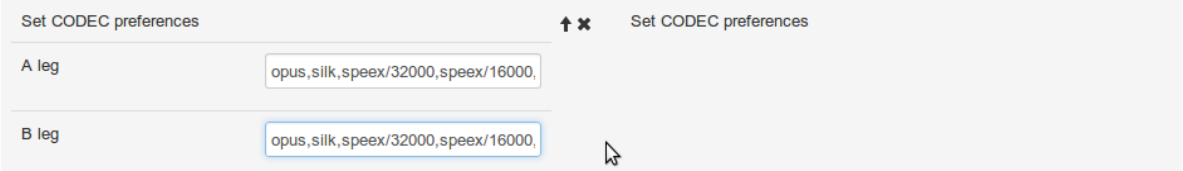

Figure 69: Example: Prefer codecs with high audio quality: OPUS,SILK,speex/32000,speex/16000,G.722,AMR-WB

Additionally codec attributes offered in SDP can be filtered out using actions **Set SDP attribute whitelist** and **Set SDP attribute blacklist**.

SDP Bandwidth attribute limiting

The SDP may contain a (session-level, or media-level) b=<modifier>:<value> attribute, which sets the bandwidth to be used (see RFC4566). Different types of bandwidth signaling are standardized, denoted by different modifiers; the most common being TIAS (RFC3890), AS (application specific, RFC4566) and CT (conference total, RFC4566). The action **Set SDP bandwidth limit** can be used to limit the signaled bandwidth: If there is a bandwidth attribute for the specified type, it will be set to the limit if it is signaled to be more than the limit. If there is no bandwidth attribute for the specified type, one will be added.

Especially when the actual RTP bandwidth available to a call is limited using the action **Limit Bandwidth per call (kbps)**, using this action the SBC can signal the maximum available bandwidth to the endpoints.

If 'Media type' is not set, this action sets the session-level bandwidth attribute. If 'Media type' is set, it sets the media-level bandwidth attribute for that media type. E.g., if 'video' is set as Media type, then all m-lines with type 'video' will have a properly limited bandwidth attribute. 'Media type' can be set to only a single media type value (i.e. 'video' or 'audio'), no list can be given here (i.e. 'video, audio' is wrong); if multiple media types should be limited, multiple actions must be used.

4.10.10 Media Handling

Introduction

In SIP networks, the signaling and media packets may traverse different paths and may be handled by different servers in the path. The ABC SBC can be on both the signaling path and the media path, or only on the signaling path of a call.

In the SIP signaling, the IP addresses between which the actual media is exchanged is negotiated using Session Description Protocol (SDP). The default signaling mode establishes a direct media path between the two call parties as shown in Chart I of Figure RTP Anchoring with and without Symmetric Mode. If the ABC SBC is configured to intervene and insert itself in the media path, it replaces IP addresses in SDP signaling with its own, attracts RTP packets to itself and forwards them to the other call party, as shown in Chart I and II.

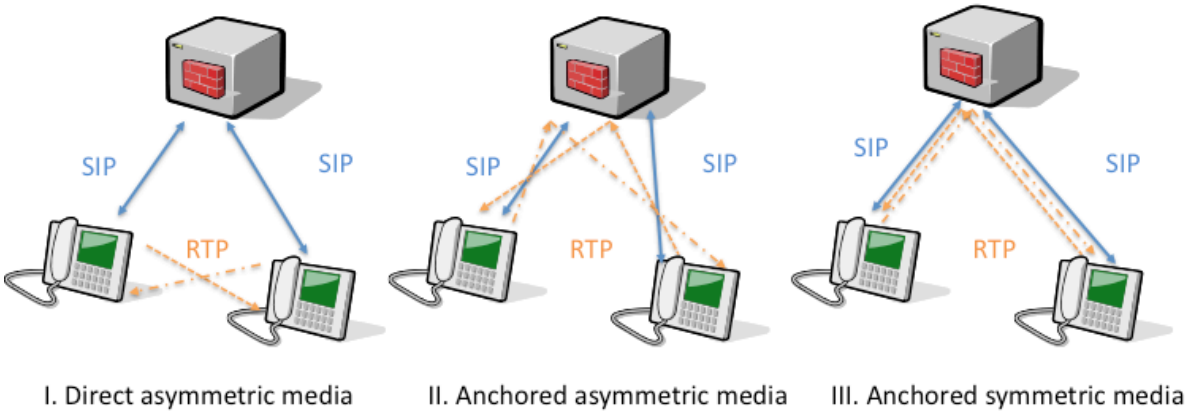

Figure 70: RTP Anchoring with and without Symmetric Mode

Generally it is desirable to have RTP processed and relayed at as few network elements as possible, in order to maintain lowest possible total latency and delay variations (jitter). Performance impact of media relay is discussed in more detail in Section SBC Dimensioning and Performance Tuning. However, the ABC SBC must be inserted in the media path in the following quite common situations:

- **NAT handling** - User Agents that are behind a NAT are not able to send RTP directly to other User Agents behind NAT
- **Connecting unroutable networks** – when the ABC SBC connects networks that cannot directly send packets between themselves, media anchoring must be enabled.
- **Topology hiding** - to improve end point and network security, the ABC SBC prevents entities from learning the addresses of other entities in the network
- **Bandwidth limitations** - the ABC SBC can limit the bandwidth used by one call to the amount that is necessary in order to prevent denial of service attacks, see Sec. Traffic Limiting and Shaping for more details.
- **Traffic monitoring** - the ABC SBC can be used to monitor the amount of media traffic used
- **RTP filtering** - the ABC SBC can filter unknown or not negotiated RTP packets
- **Logging and tracing** - for troubleshooting call audio quality issues, the ABC SBC can be used to get a trace of the traffic including RTP packets
- **Recording** for sake of monitoring, archival or lawful interception - in any of these case the operator must relay RTP packets in order to record the audio.
- **Quality assurance** - especially when protecting a hosted PBX service it is often needed to record incoming calls. To support this feature, an RTP anchoring is needed.
- **WebRTC gateway** - in this mode the ABC SBC must receive the media flows to be able to convert them between plain RTP and secured DTLS-SRTP.

Media Anchoring (RTP Relay)

Media anchoring is activated by applying the **Enable RTP anchoring** action on a call in the A or C rules of either source or destination Call Agent or Realm. This action may be activated several times on a call, however, once activated it can not be deactivated for that call. Executing this action is a prerequisite for many other actions described in this section, they will otherwise not work properly.

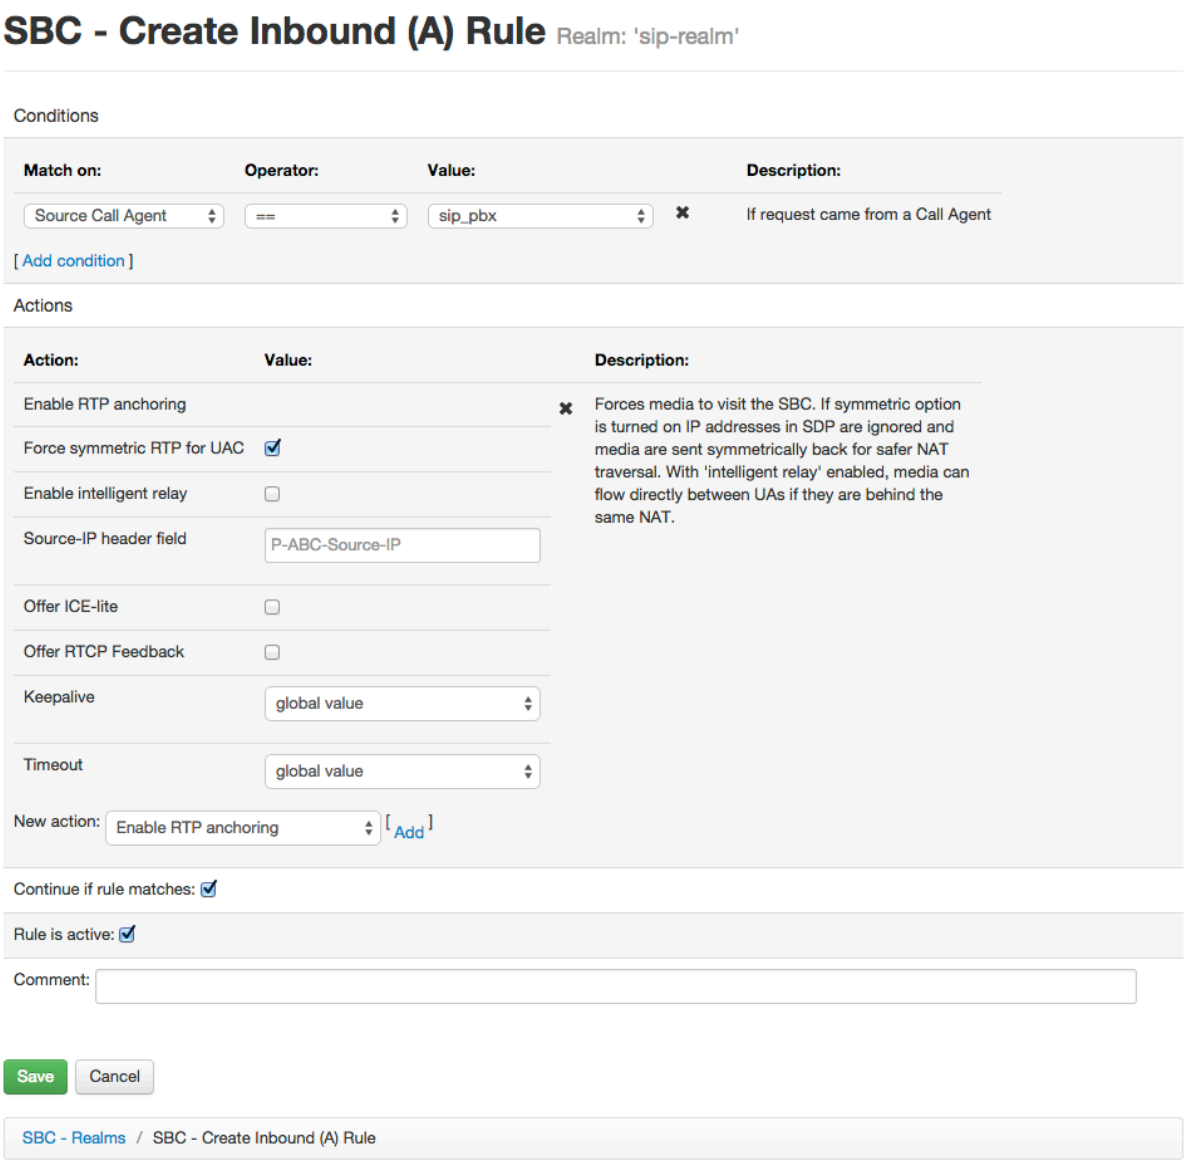

Figure 71: Media anchoring configuration

The following sub-sections described in more detail respective configuration options of media anchoring.

RTP, RTCP and FAX (T.38) Relay

If media anchoring is activated, both RTP and RTCP packets are relayed for a call. Also, Facsimile over RTP and over UDPTL (T.38) is relayed by the ABC SBC. No configuration step is required, the ABC SBC forwards Fax automatically.

Symmetric RTP Mode and NATs

For User Agents behind a NAT - especially if both user agent are behind NATs - relaying media through the ABC SBC may alone not be enough to accomplish NAT traversal. The problem is the ABC SBC cannot easily determine the public IP address to which to relay RTP media for a SIP phone. The IP address advertised by the User Agents in their SDP payload is non routable.

[2]The name "Symmetric RTP" is derived from the property of a UA that it sends RTP from the address/port combination where it expects to receive RTP at.

In a solution here called "Symmetric RTP"[2] (also called Comedia-style[3] RTP handling) the ABC SBC ignores IP address advertised in SDP and learns the IP address and UDP port number of the UA by observing RTP packets coming from the UA. It then starts sending the reverse RTP stream to that address. Symmetric RTP is activated by either one of:

- SIP device, by inserting an "a=direction: active" property in the SDP
- ABC SBC, by enforcing it using the "Media far end NAT traversal" option in the "Enable RTP anchoring" action

We recommend to leave this option turned on for all Call Agents in their both inbound and outbound rules. That not only works safely in most cases, but is also more secure in that it prevents use of bogus IP addresses in SDP payloads. Only when a Call Agent is a) known not to be implemented symmetrically AND b) is directly reachable without NATs in between, it makes sense to turn this option for the Call Agent off.

The RTP flows are depicted in Figure RTP Anchoring with and without Symmetric Mode. The chart I. shows RTP flows when no media anchoring is engaged. The RTP packets take then the "shortest path" without any SIP intermediary. This flow fails in presence of NATs because telephone's private IP address advertised in SDP is not reachable for the peer device. The chart II shows RTP flows when media anchoring is enabled. The RTP flows from and to a SIP phone are not symmetric, i.e., they are sent from and to different UDP ports as advertised in SDP payload. Like in the chart I, NAT traversal will fail. The chart III shows symmetric RTP that is the safest option for NAT traversal. IP addresses in SDP payload are ignored and the ABC SBC relays media to a telephone to address from which phone's RTP packets come.

Note well: it is important to realize that enabling **Media far end NAT traversal for UAC** will open a security weakness subjecting the call to a so called **RTP Bleed** attack. It can be mitigated partially by using the **Lock on addresses learned from RTP** option. Forcing usage of **Secured RTP** will effectively mitigate this attack as the SRTP packets will be authenticated prior to the address learning step.

Intelligent Relay (Media Path Optimization)

If the ABC SBC handles a call which is originating from the same network that it terminates to, it may be useful to skip media anchoring for that call, in order to save bandwidth and to reduce the total latency introduced. The ABC SBC detects that the caller and the callee are behind the same NAT and is so, bypasses media relay. The test is done by comparing source IP address of incoming INVITE to the intended destination of the request. This algorithm works only if both User Agents are behind the same NAT and there are no intermediary elements between the NAT and the ABC SBC. If this condition doesn't hold, the optimization will fail. That may for example happen if two user agents from behind different NATs speak to the ABC SBC through an intermediary proxy server and appear to the ABC SBC as if they were behind the same IP address. Further, this algorithm does not detect UAs behind a NAT which controls multiple public IPs. Also, the signaling IP address of the callee used for the comparison is projected and does not support domain names.

[3]The name "Comedia" came from an Internet Draft proposing use of "Connection Oriented Media". The Internet draft draft-ietf-mmusic-sdp-comedia became eventually RFC 4145.

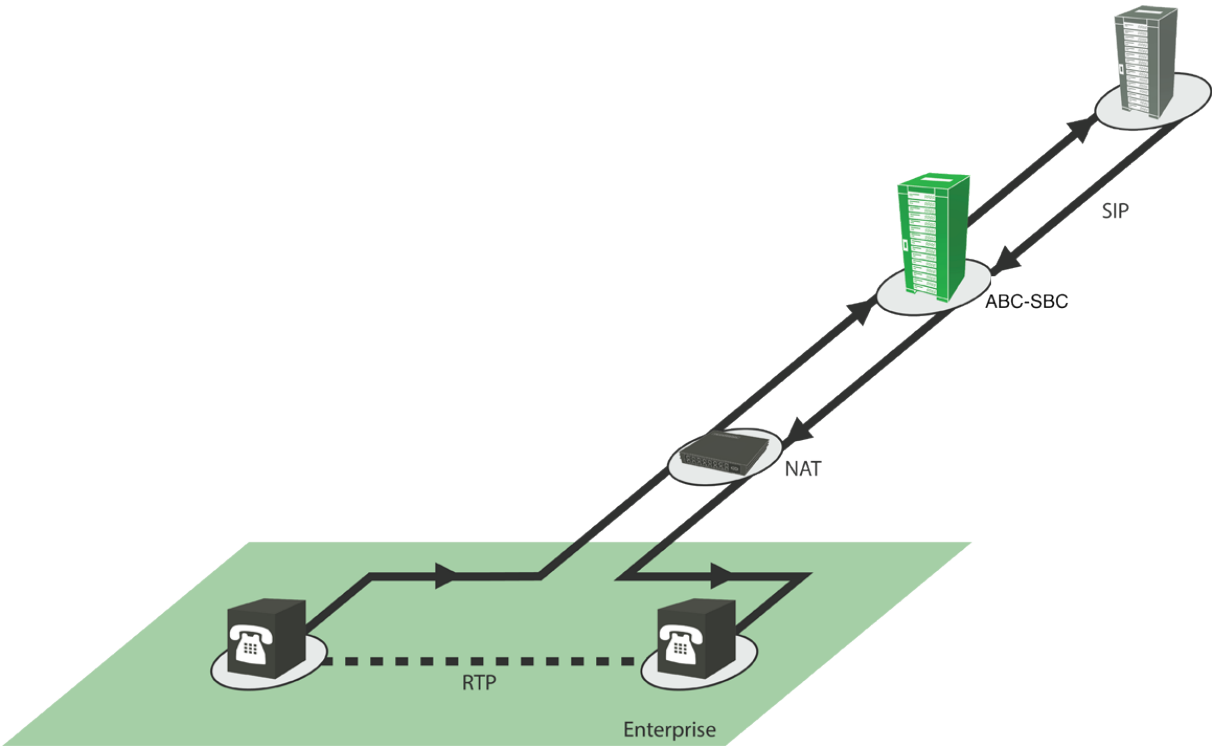

Figure 72: Intelligent relay

A separate header field, the "Source-IP header field", is used to transport the information about the caller's network through additional proxies in the signaling path. The header name may be configured as a property of the "Enable RTP anchoring" action, so that it can be customized and subsequently filtered out if necessary. This way, the ABC SBC can perform the "same-NAT" test even in scenario shown in Figure Intelligent relay and which a call passes the ABC SBC twice. Once on the way in, when source IP address is known but not the final destination , and then on the way out where the destination is already known but the source IP address would be unknown without the additional header-field.

Advanced Anchoring Options

Further media-anchoring options are useful for interaction with advanced clients that use newer protocols: ICE and STUN for NAT traversal, and also additional RTCP feedback for measuring QoS. Such clients are yet rare, however a new interoperability profile for WebRTC clients does actually include all of these. See also Section SIP-WebRTC Gateway. Other group of options are those related to keeping calls alive: they make sure that the ABC SBC and its communication peers will properly detect active calls as such, and timely detect calls that ended abruptly.

The following advances options are available:

- *Offer ICE-Lite* – adds ICE-lite (server-side ICE) capability to SDP. This is a must for WebRTC clients that expect their peers to communicate using ICE. WebRTC Call Agents must thus have this option enabled in both A and C rules. It can be also useful for SIP-based User Agents if they support ICE – however generic ABC SBC NAT techniques do not require use of ICE for facilitating NAT traversal.
- *Offer RTCP Feedback* – adds additional RTCP capabilities for sake of finer QoS monitoring than available in traditional RTP implementations. This is mostly useful for WebRTC implementations which include this extension in their interoperability profile.
- *Keepalives* – allows to send keep-alive RTP traffic. This is useful if one side of a call detects and discontinues inactive calls whereas the other side suppresses RTP due to Voice Inactivity Detection or On Hold scenarios. With this option turned on, the calls will not be discontinued.
- *Timeout* – allows call termination when no RTP traffic appears. Useful to eliminate "hanging calls" due to abruptly disconnected SIP devices.

RTP and SRTP Interworking

The ABC SBC can also transform media between "plain-text" RTP and encrypted SRTP. This is particularly useful in SIP/WebRTC interworking scenarios detailed in Section SIP-WebRTC Gateway.

The action *Force RTP/SRTP* performs protocol admission in A-rules and protocol conversion in C-rules. When placed in A-rules, it only permits calls corresponding to the requested protocol, the calls will be rejected otherwise. When the action is placed in C-rules, it converts media to the chosen protocol. If the chosen protocol is SRTP, the keying protocol must be also chosen: DTLS or SDES. When the destination is a WebRTC client, the keying protocol must be DTLS since spring 2014.

SRTP with RTP fallback ("SRTP fallback to nonsecure RTP") is a method of optional SRTP (opportunistic encryption) where the offerer sends an SRTP offer, but the answerer can fall back to RTP in case SRTP is not supported. This means that, contrary to SDP offer-answer requirements of RFC3264, the transport of the answer can be different to the offer: it can be RTP/AVP(F) when the offer is RTP/SAVP(F). This Cisco-specific method also adds a Supported tag "X-cisco-srtp-fallback". Whether SRTP or RTP is used can be renegotiated at every Offer-Answer exchange (e.g. re-Invite). This fallback method does not try to re-negotiate non-secure RTP if a 488 is received.

Figure 73: Enforce SRTP

SRTP End to End encryption

The action *End to End encryption* provides the capabilities to stay in the media path while not interfering in the SRTP negotiation. The SRTP key is negotiated by both peers without any intervention of the SBC, which is not able to encrypt/decrypt the media if this action is enabled. The RTP or SRTP packets are then just relayed as-is.

Transcoding

To enable broader interoperability, the ABC SBC can transcode between different codecs, that is it will decode incoming RTP packets and encode RTP packets into a different codec. The ABC SBC currently supports transcoding for audio only, there is currently no support for transcoding video streams.

For transcoding to be available in the ABC SBC, the operator needs to get and install the proper license for the media processing package, see Sec. Container Installation for details. Also, for some codecs, patent licenses need to be acquired separately. For some codecs, special software packages need to be installed, please contact FRAFOS support if in doubt.

The FRAFOS ABC SBC supports software based transcoding. Transcoding adds non-negligible processing power requirements to the SBC hardware, see Sec. SBC Dimensioning and Performance Tuning for details.

Depending on the codecs used, transcoding also reduces voice fidelity, especially if transcoding is applied a multiple times on the path of the call.

The action **Activate transcoding** takes as parameter a comma-separated list of codecs which are added to the SDP offer, if not present in the original offer. The codecs listed here must be supported by the ABC SBC . If the other party accepts one of these, the media stream is transcoded. Both (or, all) codecs which the ABC SBC should transcode between need to be added to the list of transcoding codecs.

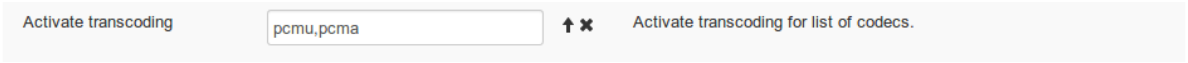

Figure 74: Activate transcoding

For example, if "PCMU,PCMA" is configured as transcoding codecs, and the caller offers only PCMU and the called party PCMA, the ABC SBC transcodes from PCMU at one side to PCMA at the other side and the other way around. If both sides happen to support the same set of codecs then transcoding will not be needed and will not be used.

Audio Recording

Recording may be useful for a variety of purposes: most often for archival, monitoring and legal interception. If a call is selected for recording, the ABC SBC collects audio and stores it in a WAV file. Each direction is stored in one channel, the file is stored with sampling rate 8kHz, two bytes per sample (PCM), two channels. To allow recording, media anchoring must be turned on. Recording works only if supported codecs are used.

Recording is activated by the action *Activate audio recording* that takes a comma-separated destinations as parameter.

The destination may be a filename, a HTTP server to which the WAV file is uploaded using the PUT method or a SIP URI if the call recording is to be outsourced to a SIPREC call recording server. If a SIP URI is used, only one is supported and should not be entered as a list.

The call recording action supports two more parameters allowing for start and stop announcements. These announcements are played when the recording starts or stops. Please note that the stop announcements can only be played if SIPREC is used, and the call recording server (SRS) does stop the recording before the call has been ended.

Replacement expressions can be used to provide easier identification of the system. USE CAUTION when devising the filename: filename conflicts will result in different sessions overwriting each other's WAV file. If no filename is included, the ABC SBC uses its own ephemeral filename. Filenames are made relative to the directory /data/recordings to make sure that the recording doesn't interfere with the filesystem.

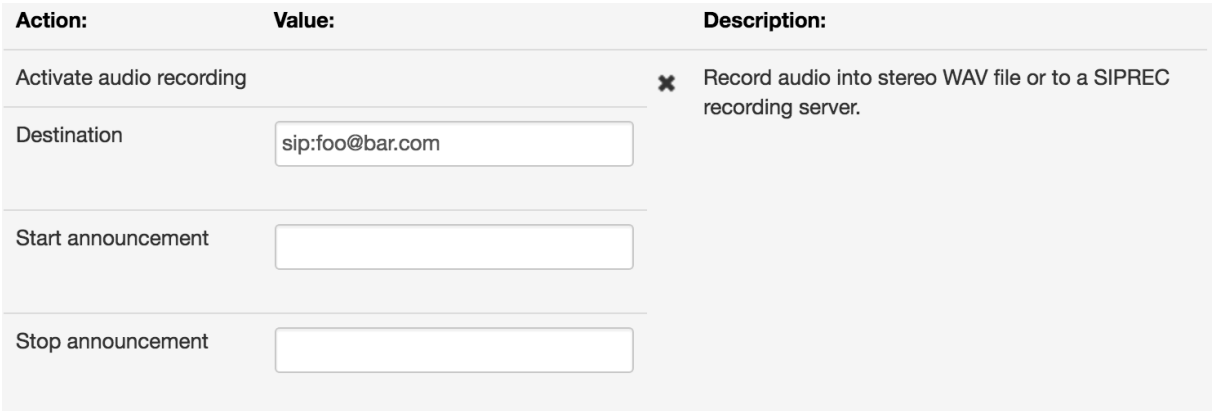

Figure 75: Activate Audio Recording

When recording and generating the WAV files completes, an event is produced.

Note that to avoid premature deletion of important archival data, the system does not delete any audio files and keeps them stored. This may potentially exhaust the disk space. Consult professional services if you need help on managing audio archives.

SIPREC specific options

When a SIPREC is used for audio recording, a set of specific options can be configured.

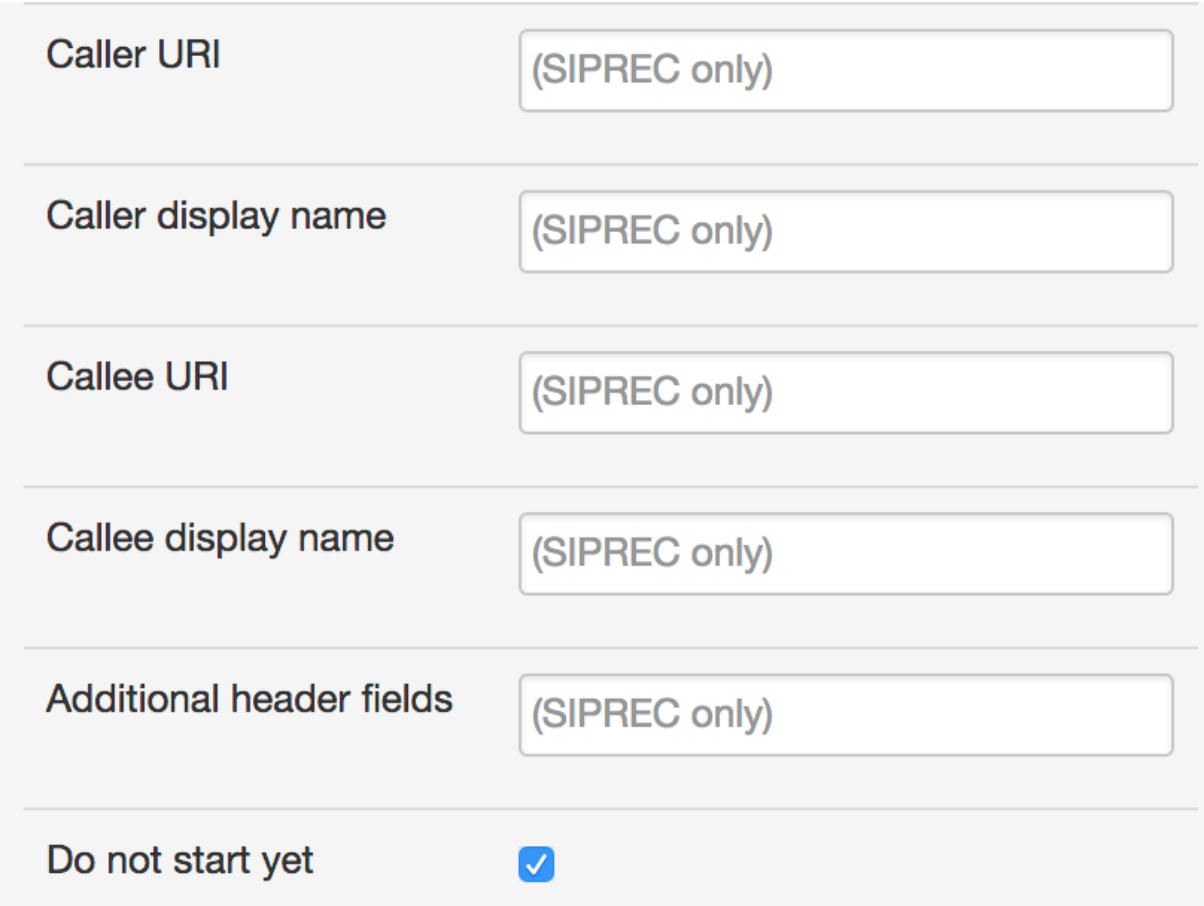

Figure 76: SIPREC options

The fields related to caller and callee determine the values transmitted to the recording server within the SIPREC metadata. The Caller & Callee URIs are used to set the participants' `nameID aor` tags, while Caller/Callee display names are used to set the `name` tags.

The field "Additional header fields" can be used to add header(s) to the SIP INVITE message sent to the SIPREC server.

The last option "Do not start yet" allows to delay the start announcement until the SIPREC server signals it has started the recording. This allows for notifying the user properly. Please note that the media streams are transmitted during the complete call to the recording server, even if this feature is used.

Since ABC SBC 4.5 it is possible to configure SIP timers towards SIPREC server. With appropriate values this may help to speed up error detection. See SEMS Parameters for configurable options.

Playing Audio Announcements

Often it is practical to inform caller about an error by an audio announcement rather than a numerical SIP code. The ABC SBC can play audio files on several different occasions for each of which there is an appropriate action:

- "Refuse call with audio prompt" plays an audio file immediately on receipt of an INVITE

- “Play Prompt on Final Response” plays an audio file on an unsuccessful call attempts
- “Generate Ring-Back Tone” plays an audio file instead of the default ringing tone
- “Activate Music On Hold” plays an audio file when a party chooses to put a call on hold

The action “Refuse call with audio prompt” plays a prerecorded audio WAV file immediately on receipt of a SIP INVITE. It is typically used to decline a call using a pre-recorded message. This action plays a prerecorded WAV audio file and terminates the call. It takes the following parameters:

- filename of the announcement relative to the global configuration option “Prompts/Base Directory”. The filename must refer to an existing file which has been uploaded to the prompts directory by administrator.
- a checkbox specifying whether the announcement shall be played as early media or establish a regular call
- a checkbox specifying whether the announcement shall be played once or in a loop
- SIP response code, phrase and header-fields to be used for terminating early media announcement (unused if a regular call is established)
- also instead of playing a pre-recorded WAV file, a beep tone can be generated. To turn it on, activate the “generate ringtone” checkbox and describe the tone length, on-off periods and frequencies.

Figure 77: Example Action for Playing an Announcement

The same effect can be achieved for SIP calls that failed downstream. The action **Play Prompt on Final Response** plays an announcement for call attempts that failed downstream due to one of the listed failure codes. The announcement can be played as a regular call or as early media, in which case a specified SIP response code will conclude the announcement.

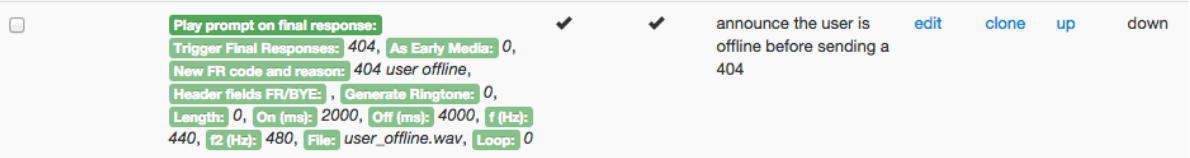

Figure 78: Example Action for Playing an Announcement on Receipt of a 404 Response

Also it may be useful to play specific tones or audio when called party’s telephone is ringing. To enable this functionality, use the action “Generate Ring-Back Tone” and either define a tone or reference to an audio file to be played instead of the default ringing tone. The screenshot in Figure Example Action for Playing an Audio File during Ringing is showing such a configuration that plays a predefined audio file during the ringing phase after the receipt of the UAS’s 180 response. Alternatively one could have played a predefined dual-frequency tone.

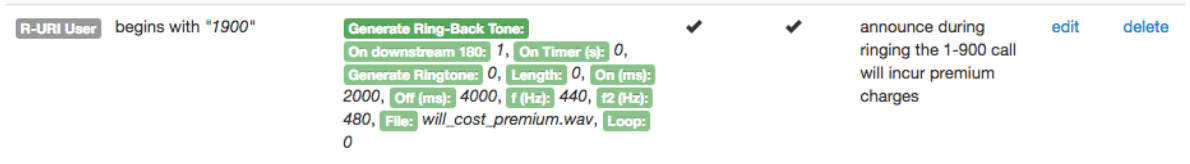

Figure 79: Example Action for Playing an Audio File during Ringing

Similarly it is possible to play an audio file whenever a party chooses to put a call on hold. The action takes several parameters that allow it to define how the on-hold status is signaled to the other call party and if the audio file is played once or in a loop.

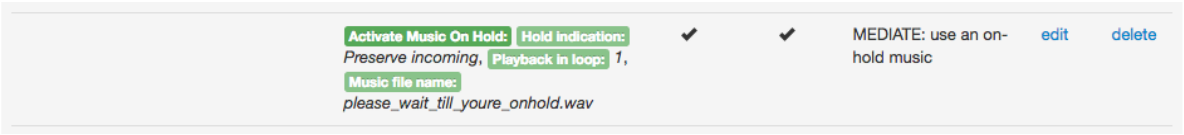

Figure 80: Example Action to Activate an Audio File when a Call is Put on Hold

Onboard Conferencing

To accommodate smaller-scale dial-in conferences without need for an external conference bridge, the ABC SBC can mix audio calls. To enable a conference, place a "join meet-me conference" action in A rules. The action's parameters allow to specify which conference an incoming call shall join: either by two-stage DTMF dialing, or by a "hard-wired" conference ID, or by conference ID determined using a replacement expression.

If the room is entered via keypad (DTMF), then some more parameters control how that is done: A minimal length of the room can be set, and also some unacceptable room numbers (e.g., too simple, can be guessed). Once the room is entered via the keypad, a prefix can be prepended to it: This way, separate 'namespaces' of conference IDs can be used, e.g. if the same SBC is connected to two different networks which should never share a conference room, but in both of them the room ID should be entered via the keypad.

The room entered via keypad can also be split at a specific position into room ID and participant ID by using the "Split room and participant ID" setting. This way, a web interface can send out invitations with individual PINs and later identify the callers by their participant ID, while the different participants still hear each other in the room.

The following example configuration shows a conference bridge configured to serve calls from native SIP devices, SIP-based PSTN gateways and WebRTC browsers. SIP devices and WebRTC browsers are handled the same way: userpart of request URI ($rU) identifies the conference a caller is joining. Calls from the PSTN call-agents however are prompted to type in the conference id when dialing in.

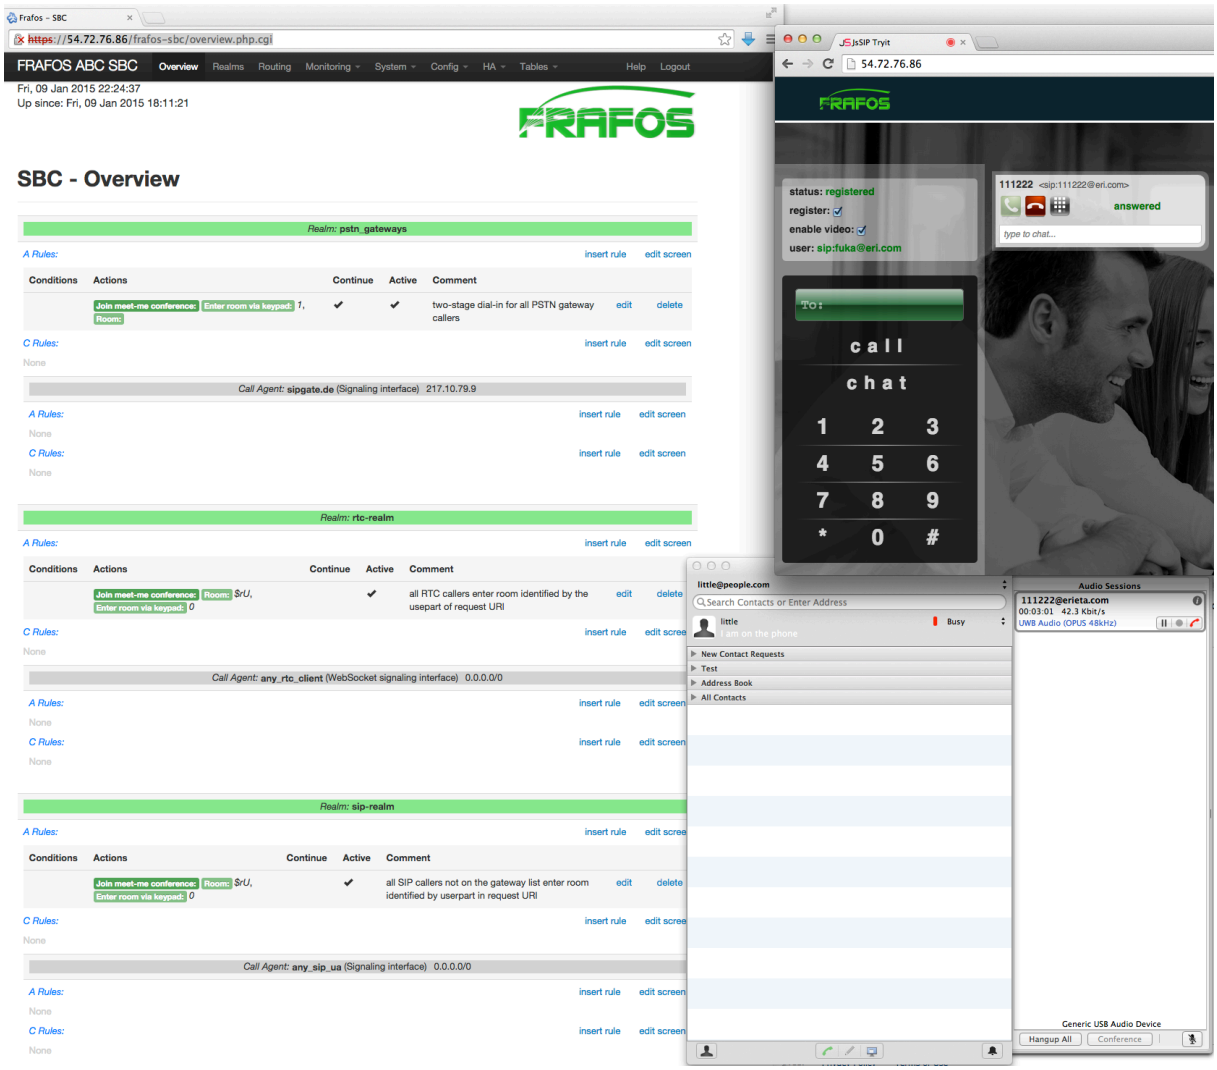

Figure 81: Example Screenshot: Conference Bridge configuration

Note that the default WAV announcement files (like conference ID prompt, or "you're welcome" message) are stored on the system in the directory /usr/lib/sems/audio/webconference and can only be changed through Command Line Interface. For more information, refer to Default Audio Files. Also note that locally processed onboard conferencing calls do not appear in the list of live calls (see Section Live Calls) in which only relayed calls are represented.

The status of the conference bridge can be inspected using CLI as shown in the following example

```
[root@ec2-54-154-137-127 ~]# sems-stats -c "DI webconference listRooms verysecret"
sending 'DI webconference listRooms verysecret\n' to 127.0.0.1:5040
received:
[200, ['1234'], 'Server: Sip Express Media Server (3.1.1-79-c86706a-417e607-87219e5
(x86_64/linux)) calls: 1 active/0 total/0 connect/0 min']
[root@ec2-54-154-137-127 ~]# sems-stats -c "DI webconference roomInfo 1234 4447"
sending 'DI webconference roomInfo 1234 4447\n' to 127.0.0.1:5040
received:
[0, 'OK', [['020D64C4-54E33359000B632C-E54DD700', '"homer" <sip:homer@simpson.org>',
3, 'direct access: entered', 0, '']],
'Server: Sip Express Media Server (3.1.1-79-c86706a-417e607-87219e5
(x86_64/linux)) calls: 1 active/0 total/0 connect/0 min']
[root@ec2-54-154-137-127 ~]#
```

Conferencing room pin protected

Conference room may also be protected in a form of a security PIN. That security PIN is asked to each participant before joining a room.

To use this option, please enable the “Room is PIN protected” option of the “Join meet-me conference” action. PIN management is subject to 3 possible options :

- First user to join is prompted to set the security PIN. To do so, please enable the “Room is PIN protected” option and leave the “PIN” field empty.
- PIN is set via action rule. To do so, please enable the “Room is PIN protected” option and set the “PIN” field to the desired security PIN.
- PIN is set via another action. User may use the “Meet-me conference set PIN” action to set and persist a security PIN into a specific provisioned table.

In the following scenario, user is required to first set the pin of a room before accessing it. We request the user dial the `9011234` so the security pin of the room `1234` may be set. He’s then able to dial the `9001234`, where he’ll be prompted for the security PIN of the room `1234` before being able to join.

Please start by creating a provisioned table of type “pins”. Use the “Meet-me conference set PIN” action to set the PIN of a room and persist that value into the provisioned table. You may then fetch the PIN value from the table and use it as call variable (see following rules screenshot)

See the following rule configuration as example :

A Rules: insert rule append rule edit screen

Conditions	Actions	Continue	Active	Comment	
Method == "REGISTER"	Save REGISTER contact in registrar	✓	✓		edit delete
R-URI User RegExp "^900(.+)"	Read call variables from table: security_pins: "$B(1.1)", Set Call Variable: room = "$B(1.1)", Log message: Log level: Error, Message: room '$V(gui.room)' security pin: '$V(gui.pin)' - set from $V(gui.set_from)($V(gui.set_at))	✓	✓	read pins table	edit delete
Call Variable Existence Does not exist "pin" AND R-URI User begins with "900"	Refuse call with audio prompt: As Early Media: 0, Header fields FR/BYE: , Generate Ringtone: 1, Length: 0, On (ms): 500, Off (ms): 500, f (Hz): 480, f2 (Hz): 620, Log message: Log level: Error, Message: byebye	✓	✓	pin not set	edit delete
Call Variable Existence Exists "pin"	Join meet-me conference: Enter room via keypad: 0, Room: $V(gui.room), Minimal room length: , Unacceptable rooms: , Room prefix: , Split room number and participant ID: 0, Position to split room: , Room is PIN protected: 1, PIN: $V(gui.pin)	✓	✓	pin set, room exist, join room	edit delete
Method == "INVITE" AND R-URI User RegExp "^901(.+)"	Read call variables from table: security_pins: "$B(2.1)", Set Call Variable: room = "$B(2.1)", Log message: Log level: Error, Message: pin is $V(gui.pin) and room is $V(gui.room)	✓	✓		edit delete
Method == "INVITE" AND R-URI User == "901$V(gui.room)"	Meet-me conference set PIN: Room: $V(gui.room), PIN: $V(gui.pin), Source IP: $si, Path to digit wav directory: /usr/lib/sems/audio/webconference/, Provisioned Table API user: sbcuser, Provisioned Table API user password: verysecret, PINs provisioned table: security_pins	✓	✓		edit delete

Please note that for this example to work, we’ve created a new CCM’ user “sbcuser” and granted him full actions on the following permissions:

- “Tables: definitions”
- “Tables: values”

Record and play username

Conference room offer the possibility to register participants’ name so it is played, under various circumstances to the other callers.

In the following example, participants have their username preserved as file on the file system for 90 days. The files is named after the “From User” rule action replacement.

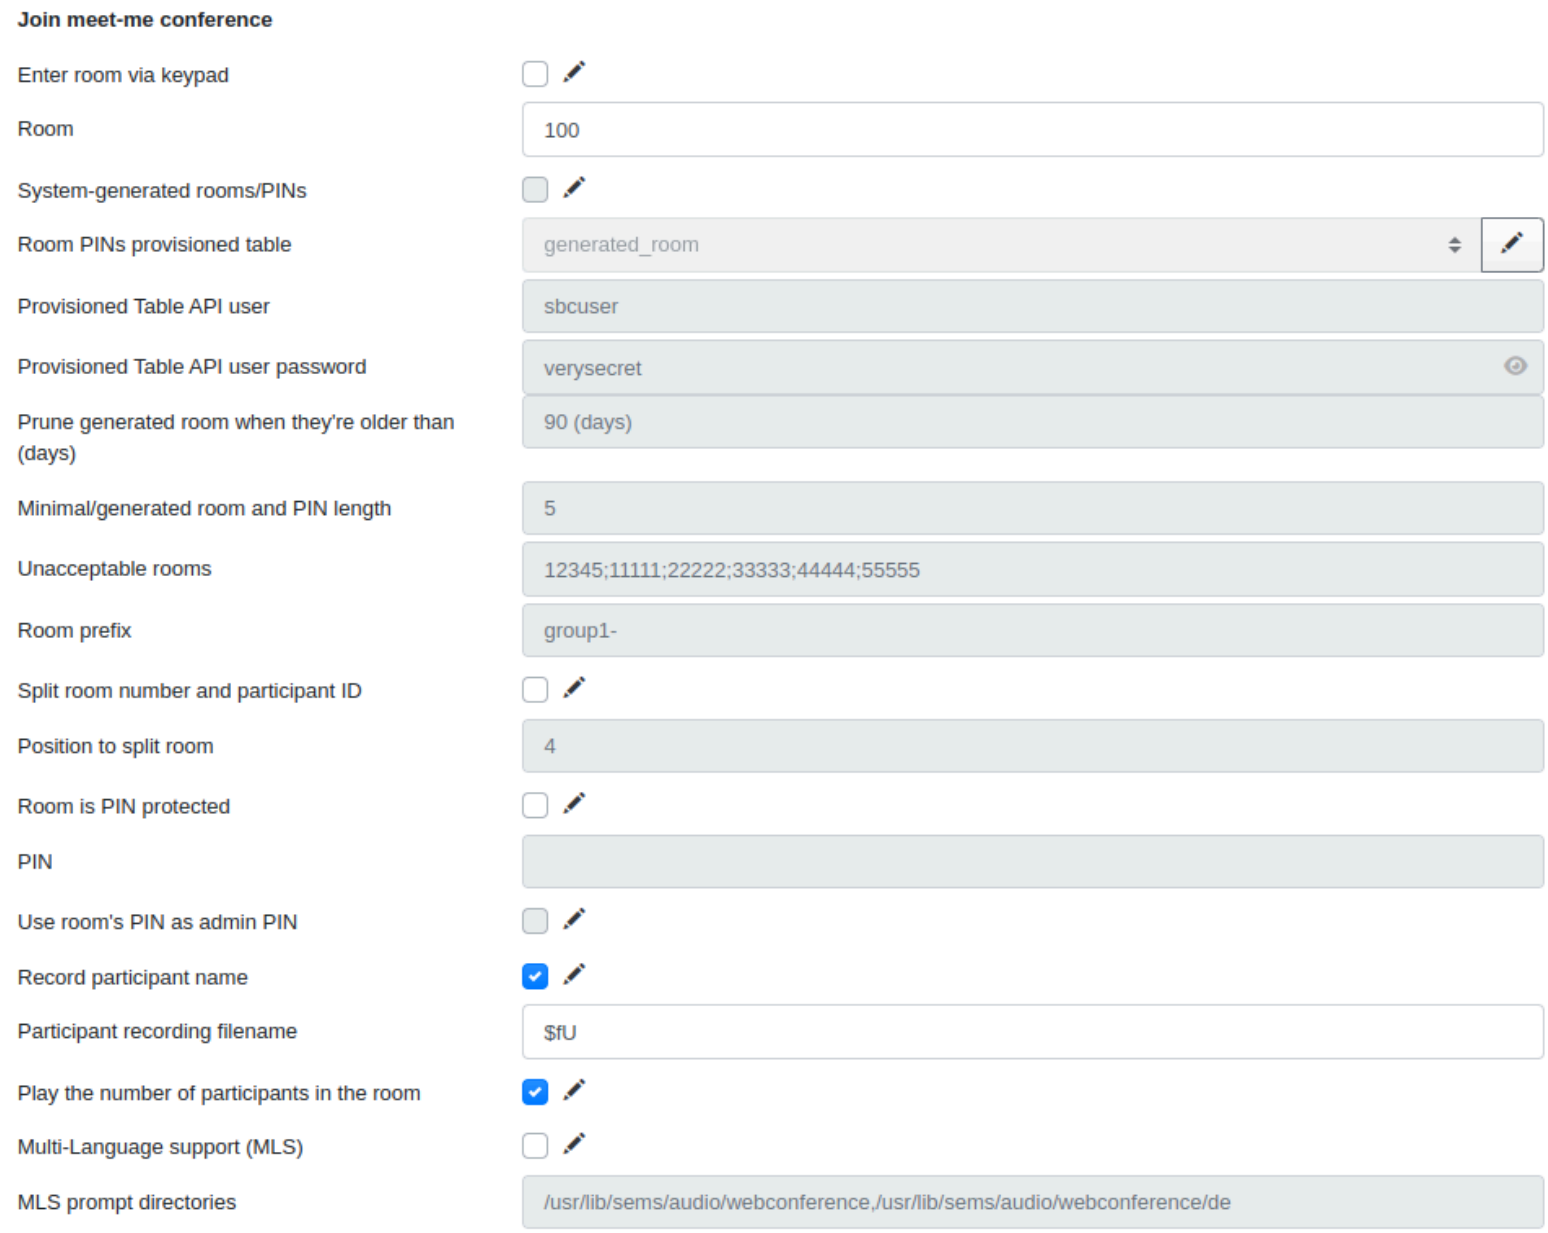

Associated gconfig:

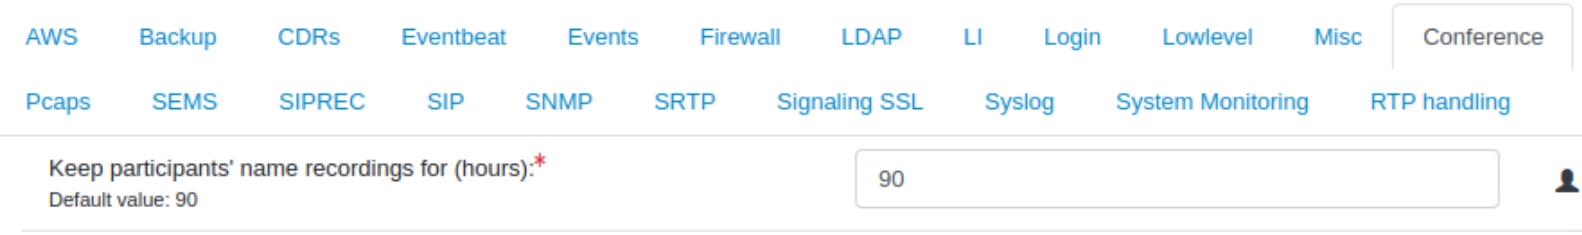

The following impacts rooms using such options:

Username recording

A joining participant will be asked for it name before joining the room. When re-joining, a participant will be prompted between either keeping an existing recording (see later description) or registering a new one. Depending on the configuration, the recorded username will be kept as `.wav` file on the file-system,

using a dynamic file naming (see example). To allow recording to be re-used, we recommend using some caller' unique information: ip, number … File will be kept until the room is closed (default behavior) or, if specified otherwise, preserved for a given amount of time (Conference Parameters) (general to all rooms).

New announces

Some new announces have been introduced, taking advantage of those new recorded usernames. As so, current room users are now notified of the following events:

- a new participant joined the room
- a participant leaved the room
- pound detected (#), the list of current participants will be played back to the emitter

Please note that, depending on the configuration (Conference Parameters), pressing * while in call play back the number of participants to the emitter. Both can be used together.

Multi lingual conferencing announcements

Conference room now support multi lingual prompts! If enabled, user will have the possibility to change the prompts' regional for his ongoing call.

The feature is disabled out of the box. To use it, please enable the "Multi- Langague support (MLS)". Optionally, one may configure custom regional prompts via the "MLS prompt directories". Defaults value set English as primary regional, user have the ability to switch to German by pressing 2.

In the following example, the default regional is set to German, French can be selected by pressing 2, English by pressing 3.

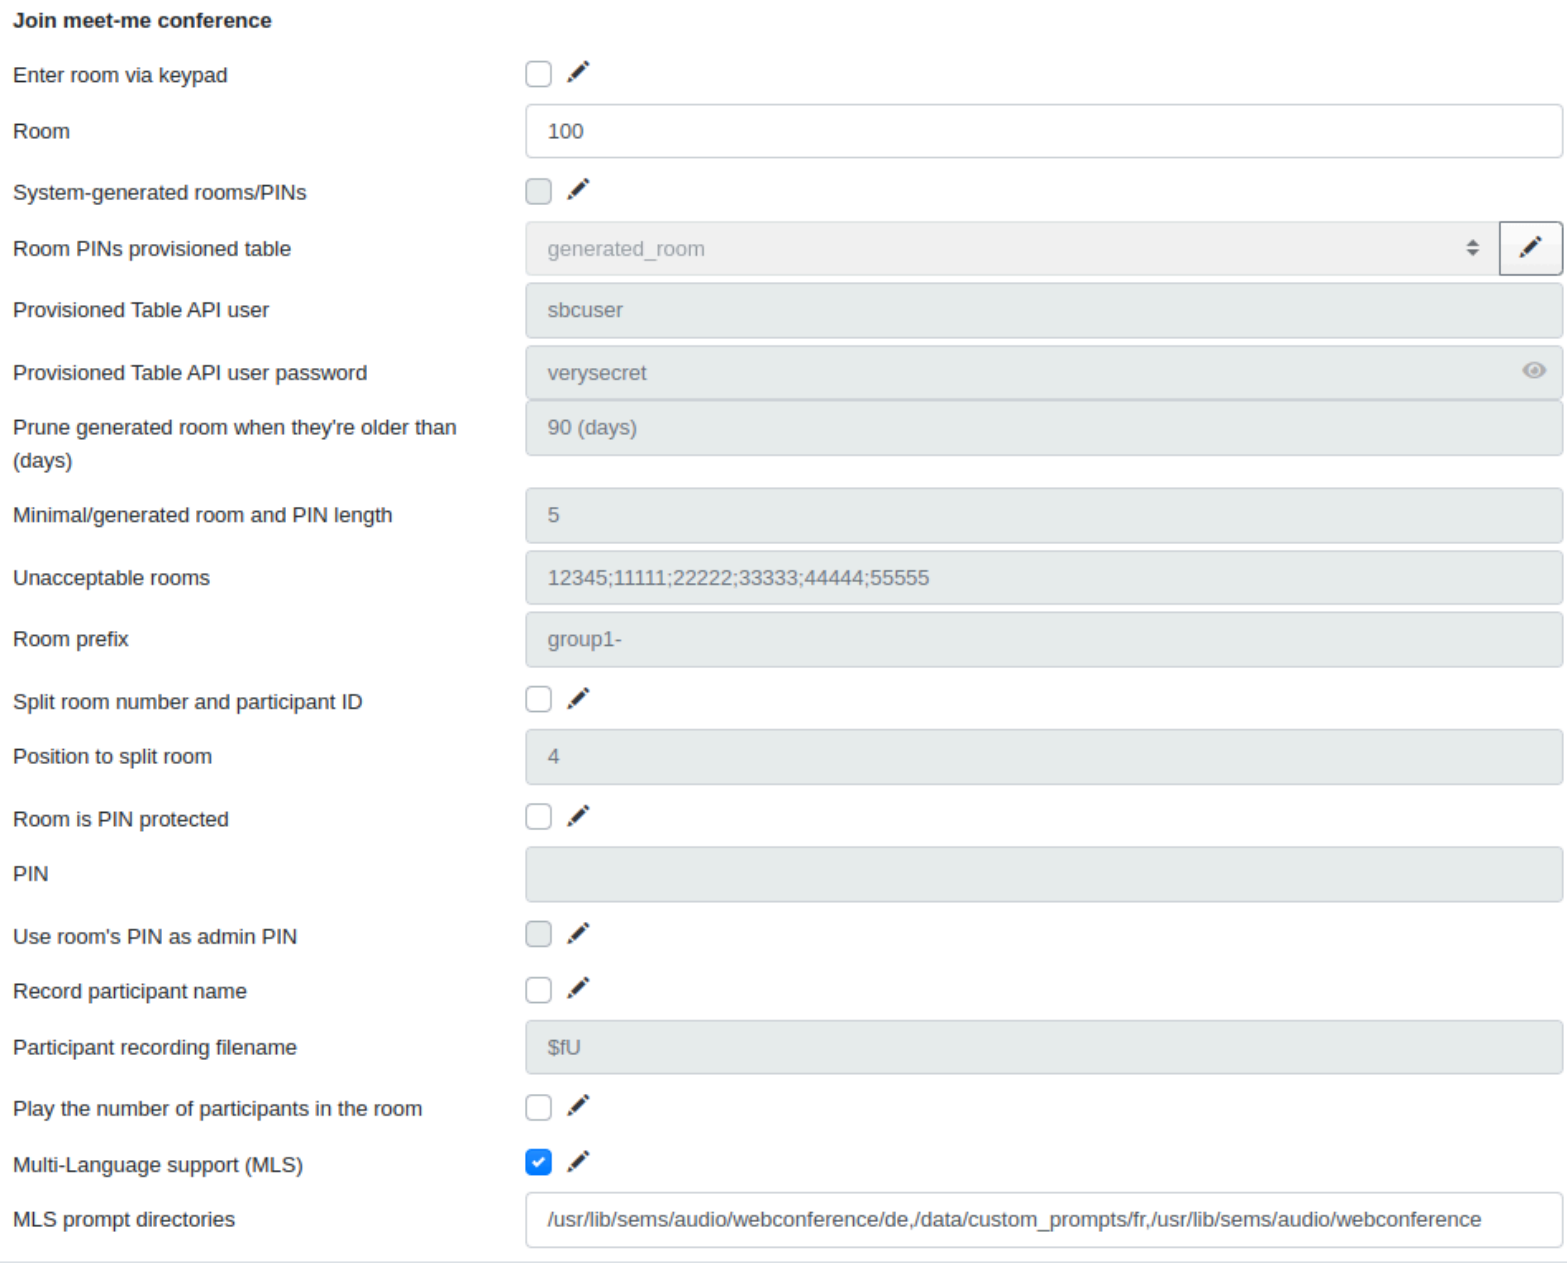

Note that for this example to work, custom French prompts were manually deployed to the `/data/custom_prompts/fr` directory of the ABC SBC container.

Please refer to Default Audio Files for a list of expected prompts.

4.10.11 NAT Traversal

SIP devices behind Network Address Translators (NATs) cannot reach other SIP devices reliably. The root reason is SIP protocol advertises SIP device's IP addresses in several places in the protocol: *Contact* and *Via* header fields SDP *c* line. These addresses are non routable once they cross a NAT device and break signaling. The ABC SBC is designed to assist the to facilitate NAT traversal for SIP devices by several techniques: it centers itself in the middle of communication path, sends signaling and media reversely to where it came from even in violation of the SIP standard, replaces private IP non routable addresses with its own and keeps all bindings alive.

As depicted in Fig. SBC and NAT traversal, when an INVITE request traverses a NAT then only the IP addresses in the IP header will be changed. Any IP addresses included in the message itself, e.g., *Contact*, SDP *c* line, will still reference the private IP address of the caller. As the callee would use the information in the *Contact* header for replying back to the caller and send media packets to the address in the SDP the call establishment will fail.

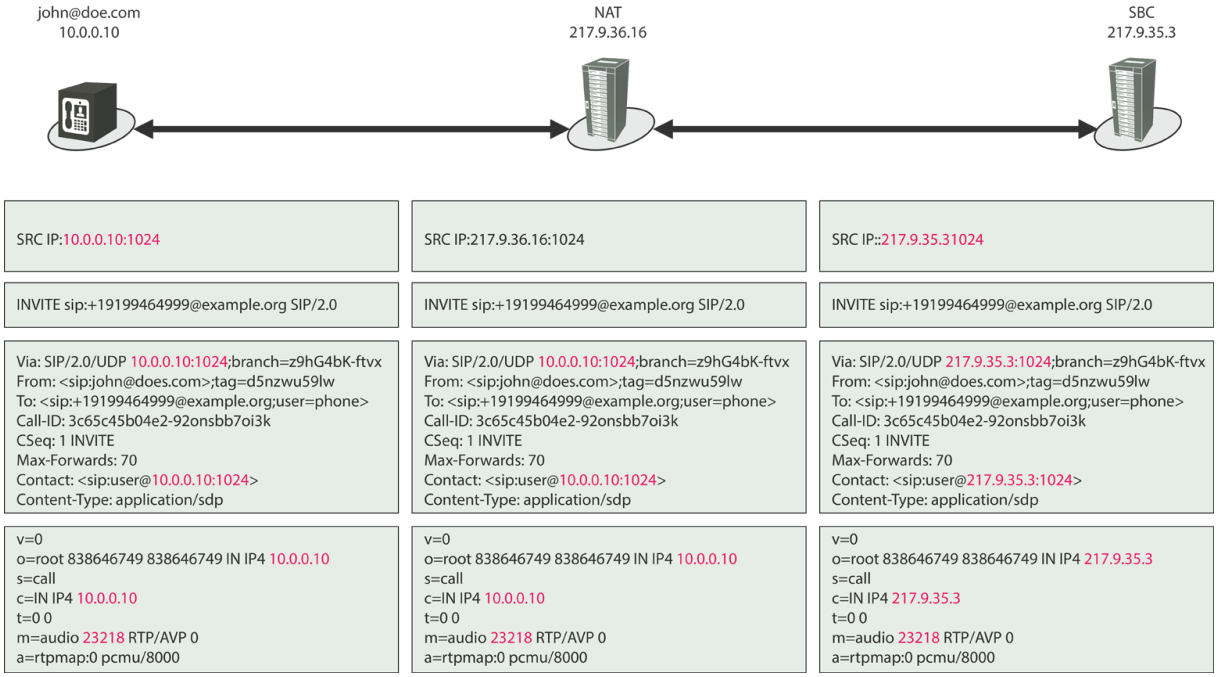

Figure 82: SBC and NAT traversal

In the context of NATs, there are basically three different possibilities with respect to the network topology that will influence the possible measures that can be taken by the ABC SBC to deal with those NATs:

- **Far end NAT:** This is the most common case in public SIP service scenarios. The SBC is located on the public Interned and the end-devices access the SBC from behind NATs. The SBC must facilitate the NAT traversal for the end-devices: it must accomplish RTP traversal, SIP traversal and registration off-load.
- **SBC is placed on the NAT:**. This is the most common case in enterprise deployments in which the SBC acts as firewall between a private network with SIP telephones and PBXes, and the outside network. It has at least one signaling and media interface inside and one outside the NAT. SIP signaling is handled natively without any additional configuration. However, it is necessary to enable RTP anchoring (relay) so that the media payload can flow from between the otherwise non routable networks.
- **Near end NAT:** here, the SBC is placed right behind the NAT and a port forwarding is configured from the NAT to the SBC. This is considered by the ABC SBC as a special case of the previous configuration. In this case, it is perfectly sufficient to configure the signaling and media interfaces with the public IP address on the outside of the NAT. It is also necessary that the configured port range on the media interface corresponds to the forwarded ports for media transport, as no port translation is supported at this place.

In order to enable users behind a NAT to be reachable the ABC SBC needs to perform the following tasks.

- Detect if a SIP user is behind a NAT. This allows to eliminate expensive NAT traversal facilitation for users who do not need it. The test is performed by the condition **NAT** in inbound rules.
- Fix outgoing calls: The ABC SBC must fix SIP INVITEs from users behind NATs so that subsequent SIP messages coming back will cross the NATs successfully. Particularly, the SBC fixes Contact header-

field and stores NAT information in dialog context. This functionality must be enabled in inbound rules using the **Enable Dialog NAT Handling** action.

- Fix incoming calls: The ABC SBC must deliver incoming INVITEs for a user behind a NAT through the NAT devices. This is only possible if the NAT devices keep the UDP or TCP binding open over which the user registered. Otherwise the user becomes unreachable once the binding expires. Therefore the ABC SBC pays great attention to the SIP registration process: It can force more frequent registrations to keep the bindings alive and it also keeps source address from which the registration came. Subsequent requests for the registered client are forwarded to the address (as opposed to the private IP address advertised in Contact header-field). To enable this functionality the actions **REGISTER throttling** and **Enable REGISTER caching** must be applied to REGISTER messages, and the action **Retarget R-URI from cache**, with the **enable NAT handling** option turned on must be applied to calls towards the client. See section Registration Caching and Handling for more details.
- Media anchoring: The ABC SBC redirects RTP stream to itself and sends symmetrically one's party RTP media to where the RTP media from the other party came from. This symmetric mode of operation overrides SIP signaling but works more reliable, because it is better compatible with how most NAT devices work. The downside of this approach is the extra bandwidth consumption on the ABC SBC and increased RTP latency. Media anchoring is enabled by the action **Enable RTP Anchoring**. "Force symmetric option" can be turned on and off for UACs in inbound and UAS in outbound rules. Media handling and RTP anchoring ABC SBC is described in more detail in Sec. Media Handling.

In summary the following conditions and actions are used to configure NAT traversal:

- **NAT** condition - check if the first Via address is or is not behind NAT. This checks if the first Via address is the same as the IP address that the SIP message was received from.
- **Enable dialog NAT handling** action - force all subsequent in-dialog messages to be sent to IP/port from which the dialog-initiating request came.
- **Enable RTP anchoring** action - force RTP from a SIP user to be sent through the ABC SBC. The **Media far end NAT traversal** option forces media from the other side to be sent reversely to where the user's media came from. Turn it off only if a Call Agent is known to reject symmetric media.
- **Enable REGISTER caching** must be applied to REGISTER messages for retargetting to function. See section Registration Caching and Handling for more details.
- **Retarget R-URI from cache (alias)** action. This action makes sure that INVITEs coming to a registered client will reach it by sending them to the transport address from which a REGISTER came. The options **Enable NAT handling** and **Enable sticky transport** should be turned on.

The example in the following subsection shows how to place the respective actions in the ABC SBC rulebase.

> **Info**
>
> In an actual deployment, the specific topology needs to be considered carefully. For example, if SIP passes the SBC twice, RTP could without precaution be also anchored twice resulting in unnecessary performance degradation. That's because the SBC recognizes inbound and outbound call legs as two separate calls.

NAT Traversal Configuration Example

This example shows how to put all the NAT-facilitating actions in a consistent rule-base. This example is based on the "far-end" NAT traversal topology as shown in the Figure NAT Traversal Example: Network Topology. In this topology, the ABC SBC is multihomed: it connects to the public Internet with one interface and to a private network with the other interface. Both networks are not mutually routable, the ABC SBC connects them on application layer. SIP telephones are on the public side behind NATs,

service provider infrastructure including a SIP registrar, proxy, media server and PSTN gateway are located in the private network.

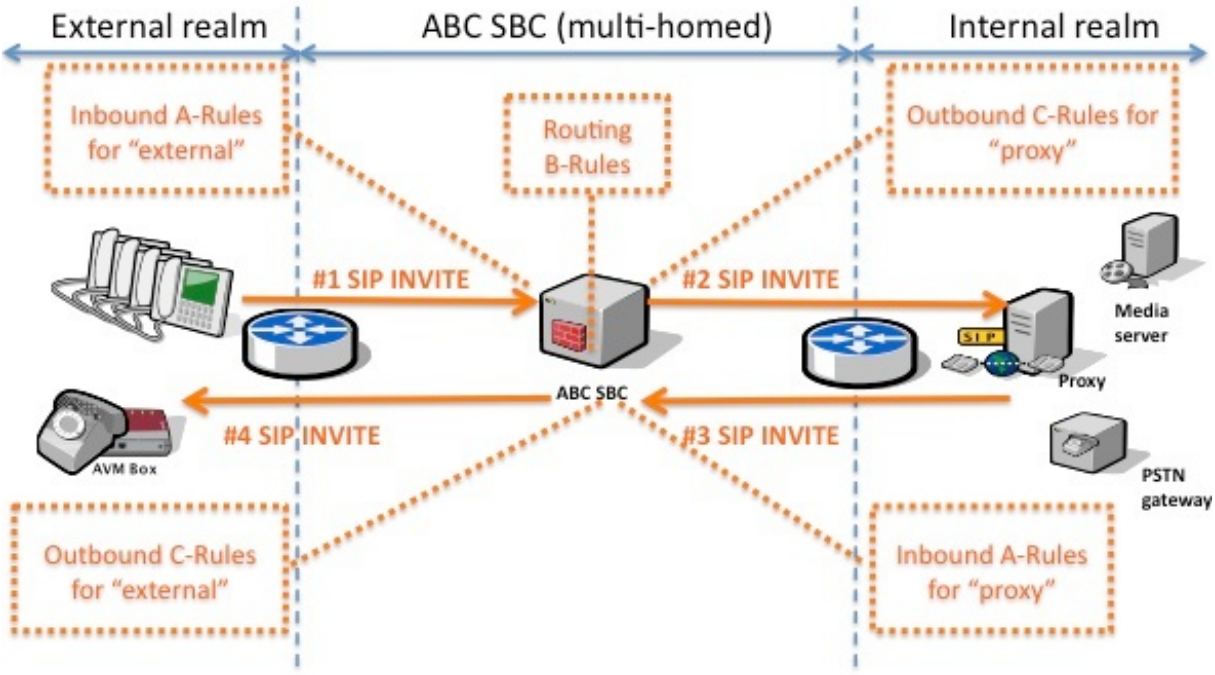

Figure 83: NAT Traversal Example: Network Topology

The first preparatory step to be is handling telephone's outgoing SIP messages coming from the external realm. That is done in the external realm's A rules: Dialog-initiating requests must be fixed in a way that reverse messages will follow the same path. REGISTERs must be stored along with the transport address from which they came. Frequent re-registration must be enforced to keep NAT-bindings alive. Eventually RTP anchoring must be enabled. The configuration fragment is shown in the Figure A-rules for traffic coming from outside.

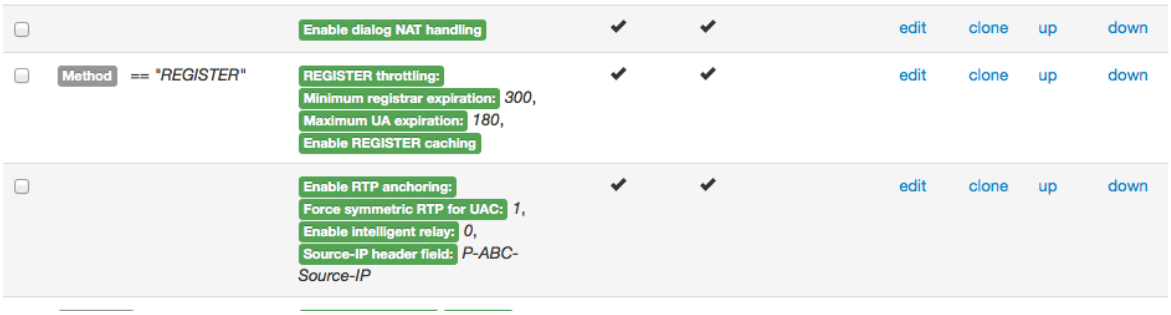

Figure 84: A-rules for traffic coming from outside

When requests pass the SIP proxy and come again from the inside network to the SBC, the addresses in them must be reverted to the form used initially by SIP User Agent. The configuration fragment is shown in Figure A-rules for traffic coming from inside.

Figure 85: A-rules for traffic coming from inside

Eventually RTP anchoring is turned in in C-rules for calls going to the public network, as shown in Figure C-rules for traffic leaving for outside.

Figure 86: C-rules for traffic leaving for outside

4.10.12 Registration Caching and Handling

ABC SBC's registration cache mediates the registration flow between SIP User Agents and SIP registrars. It keeps track of SIP User Agent contacts and shields SIP registrar from overload. It also facilitates NAT traversal. In case a user agent is located behind a NAT it will use a private IP address as its contact address in the *Contact* header field in REGISTER messages. This non routable address would be useless for anyone trying to contact the user agent from the public Internet.

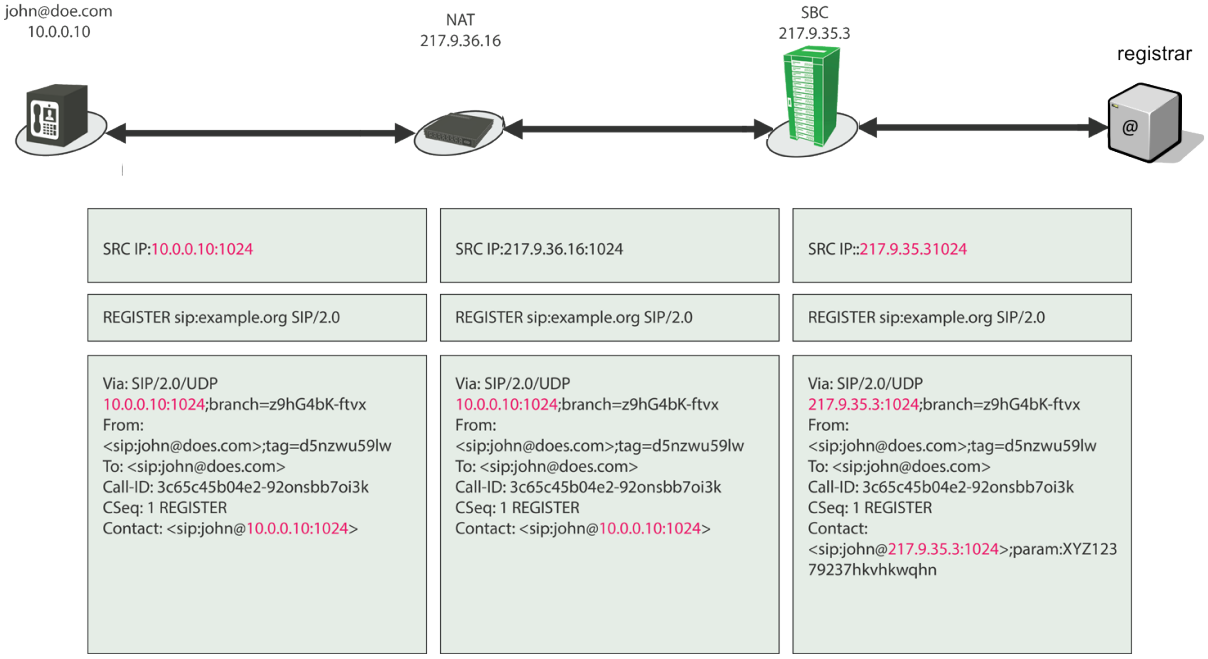

Figure 87: SBC and NAT traversal: Registration handling

In order for a SIP User Agent to be reachable the ABC SBC will manipulate its registration information. The SBC remembers the User Agent's transport address and replaces the information in the *Contact* header with its own IP address before forwarding downstream, see Fig. SBC and NAT traversal: Registration handling. This is the information that is then registered at the registrar. Calls destined to the user will then be directed to the SBC. The SBC then forwards the calls using previously stored information about the User Agent' transport address and initial unmodified contacts.

The registration cache implements the following functions:

- **Contact fixing:** SIP contacts of User Agents behind NATs include private IP addresses which are not routable from the public Internet. Therefore the ABC SBC rewrites the IP address in the Contacts with its own and holds the original Contact in its cache. If the ABC SBC connects to multiple networks using multiple IP addresses, the IP address is used which is associated with the interface over which the REGISTER request is forwarded. When later incoming requests towards the User Agent reach the ABC SBC, the ABC SBC restores the original address.
- **Keeping NAT-bindings alive**. If periodic request-response traffic was not crossing the NAT behind which the User Agent is located, the NAT address binding would expire and the client would become unreachable. Therefore the ABC SBC steers User-Agents to re-register often.
- **Registration off-loading**. Various circumstances can cause substantial registration load on the server: most often it is self-inflicted by the keep-alive functionality, but it may be also on the occasion

of a registration storm caused by a router outage, broken client or Denial of Service attack. The ABC SBC fends off such overload by using high-performance in-memory registration cache that serves upstream registrations at high-rate, handles them locally, and propagates them down-stream at a substantially reduced rate. That's the case if the registrations were to create new bindings, deleting existing ones or if they were to expire downstream. The propagated registration changes become effective on the ABC SBC only if confirmed by the downstream server. If a registration expires without being refreshed the ABC SBC issues a reg-expired event.

The following Subsection, Registration Handling Configuration Options, documents the specific actions that implement the cache functionality. Note that the procedures described here refer to individual URI registration as envisioned in the **RFC 3261**. Provisioning of bulk registration for PBXs as specified in the **RFC 6140** is described in the Section Table Example: Bulk Registration.

Registration Handling Configuration Options

The ABC SBC can handle SIP registrations in two ways, either caching them locally and forwarding them to a downstream registrar, or acting itself as a SIP registrar.

If the ABC SBC fronts a registrar, the action **Enable REGISTER caching** is applied on incoming REGISTERs from a User Agent to cache and translate its Contacts. On the reverse path towards the User Agent, the action **Retarget R-URI from cache(alias)** restores the original Contacts.

- **Enable REGISTER caching** - cache contacts from REGISTER messages before forwarding, create an alias and replace the *Contact* with *alias@SBC_IP:SBC_PORT;contact_parameters*. This method should only be applied to REGISTER messages to be forwarded to a registrar. This action has effects only on REGISTER requests, see Fig. Enable Register caching.

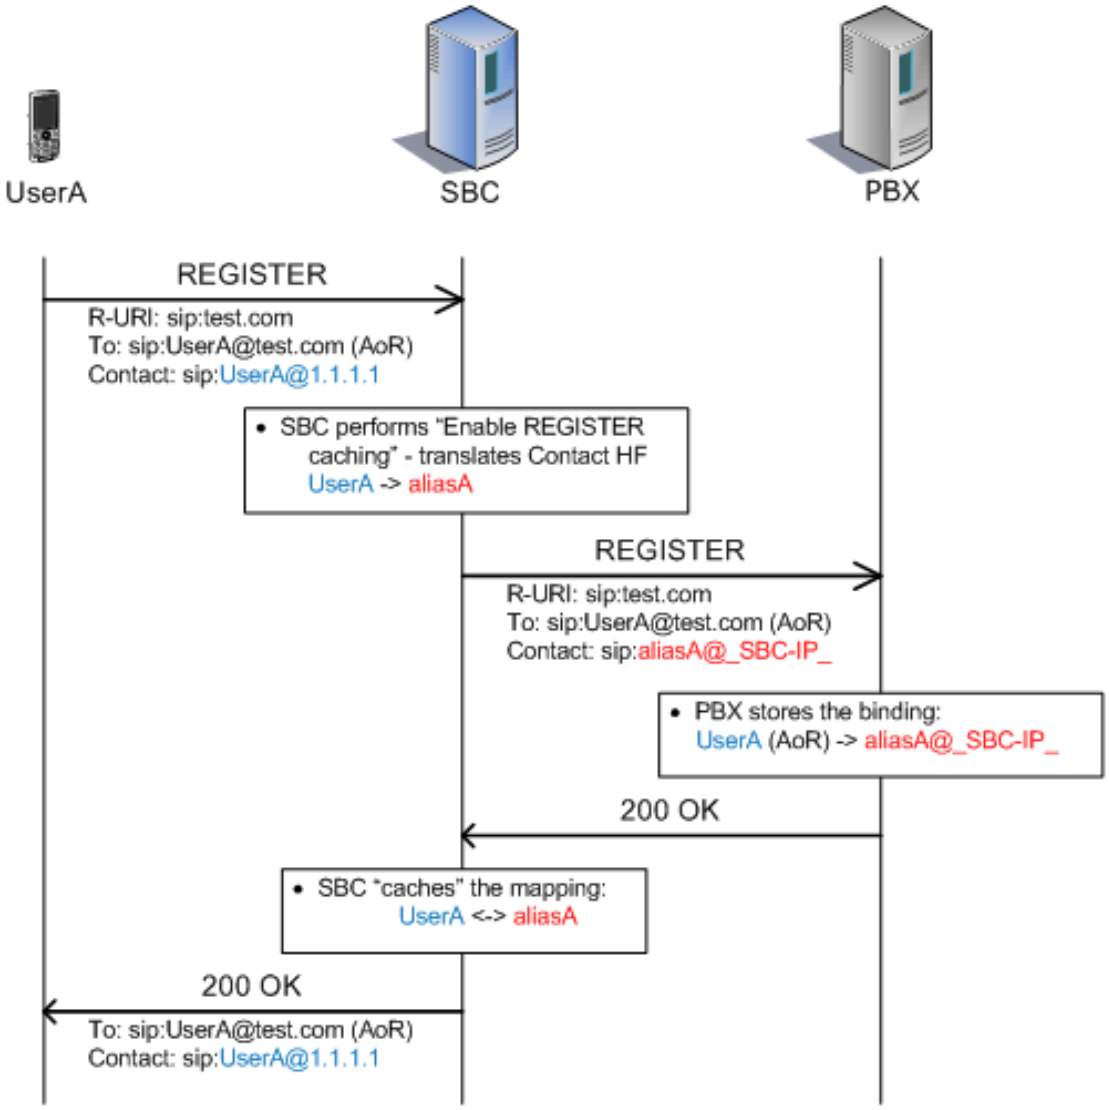

Figure 88: Enable Register caching

- **Retarget R-URI from cache (alias)** – Look up cached contact under alias and rewrite the request URI with it. Apply this action to messages sent to clients whose registration were cached previously using the "Enable REGISTER caching" action. This scenario is depicted in Fig. Retarget R-URI from cache. When an INVITE arrives to a user that has previously registered its contact information (*UserB*) the ABC SBC will forward the INVITE to the PBX which acts in this case as the SIP proxy and Registrar. The PBX will look for the registration information of *UserB* which in this case are *aliasB@_SBC-IP_* and use this information for routing the request. When the INVITE with the Request URI set to *aliasB@_SBC-IP_* arrives at the SBC, the ABC SBC will check its registration cache and retrieve the actual contact information of the user, namely *UserB@2.2.2.2* and use this information as the Request URI and forward the message to this address.
  - Parameters:
    - **Enable NAT handling**:source IP and port of the REGISTER request
    - **Enable sticky transport**: use the same interface and transport over which the REGISTER was received.

  Note: if no matching entry is found in the cache, the ABC SBC returns a 404 SIP response.

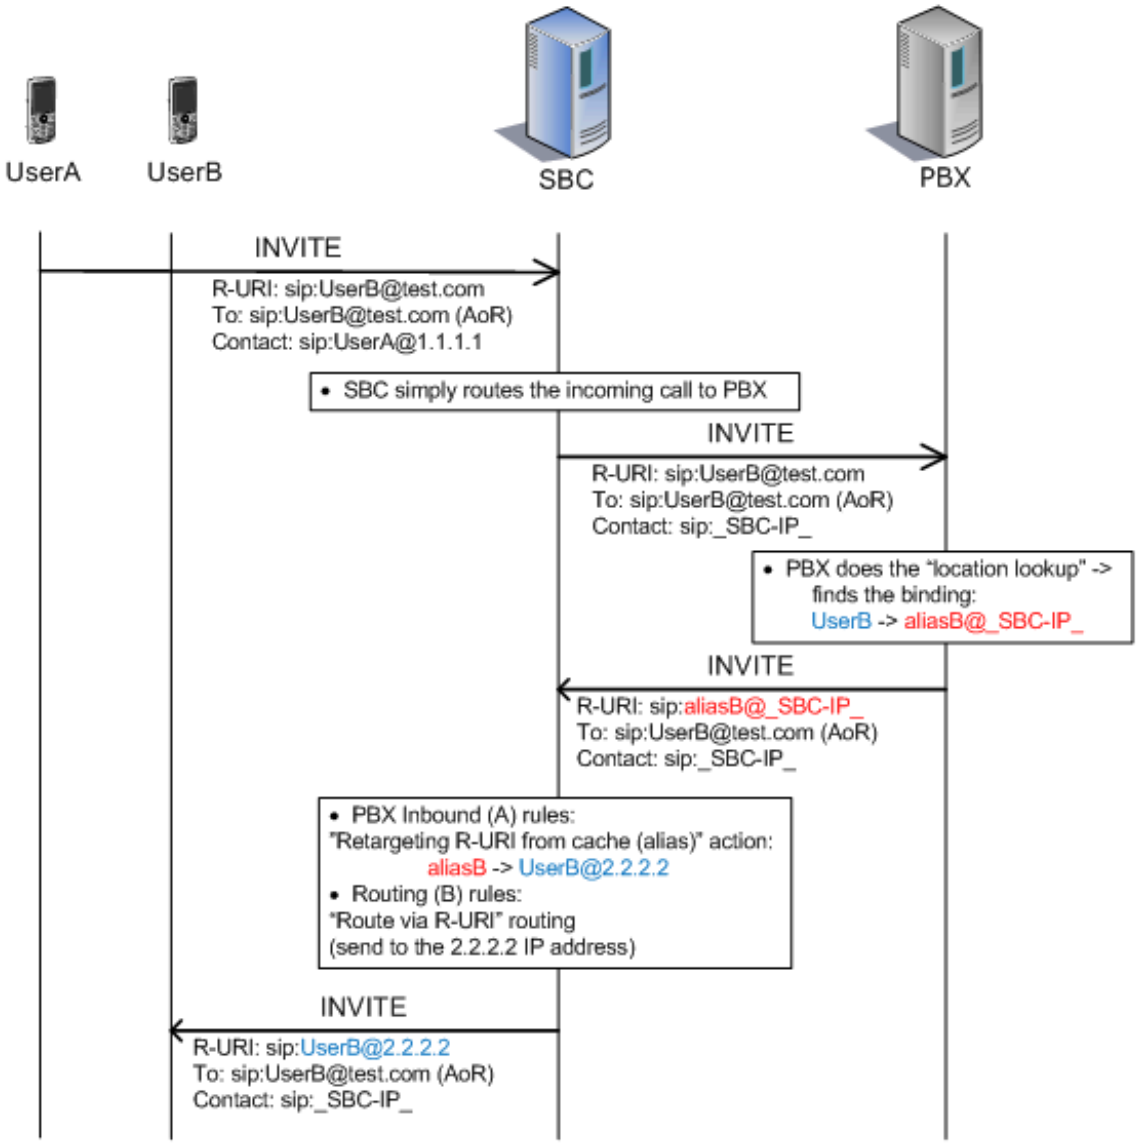

Figure 89: Retarget R-URI from cache

If the ABC SBC is used as registrar, the two following actions are used instead: **Save REGISTER contract** for REGISTERs and **Restore contact** for incoming requests towards the User Agent.

- **Save REGISTER contact in registrar** – Act as local registrar by saving the contact and replying with a 200 response.
- **Restore contact from registrar** – Use contact stored within internal registrar.

Registrar off-load

Registration throttling can be used in both built-in registrar and registrar-cache modes. The action **REGISTER throttling** enforces high re-registration rate towards User Agent Client and a reduced rate towards registrar. The high upstream rate serves the purpose of preserving connectivity by keeping address binding along the SIP path alive. The reduced rate on the downstream side makes sure that the connectivity traffic doesn't overload the downstream registrar. The **REGISTER throttling** action must precede any other REGISTER-processing action, otherwise its load-reducing function will take no effect:

- **REGISTER throttling** – force SIP user-agents to shorten re-registration period while propagating the REGISTERs upstream to registrar at longer intervals. This is useful to keep NAT bindings open without imposing the refreshing load on registrar.

- Parameters:
  - **Minimum registrar expiration**: expiration time used in direction to registrar.
  - **Maximum UA expiration**: maximum expiration time in direction to User Agent Client.

Note that these two parameters have vast impact on the volume of SIP traffic: they steer the registration rates towards upstream SIP client and downstream SIP registrar.

The **Maximum UA expiration** "knob" steers the SIP traffic rate between upstream SIP client and the ABC SBC. It suggests to the UA at which time interval it shall re-register. The lower value is enforced, the higher the registration rate will be. The SIP standard suggests one registration per hour which is not good enough to keep NAT bindings alive. Forcing the re-registration interval down to 180 seconds will cover a satisfactory share of population behind NATs. Further reducing the re-registration window will cause substantial increase in bandwidth consumption. See Section SBC Dimensioning and Performance Tuning for additional details. Note that non-compliant SIP clients may fail to honor the recommendation provided by the ABC SBC and register less often. The ABC SBC will still keep their registration bindings alive and allow incoming traffic to them. However IP and transport layer connectivity may not stay alive without the intense traffic. Contacts of such disobedient clients will show red expired status in the Registration cache window as shown in Figure Client-side Expired Registration Contact.

The **Minimum registrar expiration** "knob" steers how much SIP traffic passes through to the downstream registrar. Basically this parameter suggests to the downstream registrar how long a registration shall remain valid. The configured value is recommendatory only: the downstream registrar may accept it or change it in its final responses to REGISTER requests. The value in the response from the downstream registrar is used as the actual time-to-live for the cached registration binding. Registration renewals are not passed from the User Agents to the registrar until this time-to-live is about to expire. The longer this window is, the less traffic will be forwarded to the registrar. On the other hand, a client's failure to renew its contact will remain undetected by the downstream registrar and result in "hanging contacts" if the client is actually unavailable.

In the normal case when reduction of the upstream rate is desirable, the ratio of registrar-to-UA must be greater than 2.0 – otherwise the server registration window will not be long enough to capture more than one client registration. Typically the ratio is higher, 10.0 at least. The throttling action MUST be placed before any other register processing actions to take effect.

It is also important to understand the time-to-live of the cached registered contacts in detail. Briefly, the bindings stay active for the time requested by SIP registrar in response to forwarded REGISTER requests. They will be deleted if any of the following conditions occurs:

- the UAC fails to re-register its contact before the time-to-live of the contact set by the downstream registrar expires. That also means that a failure of a User Agent Client to renew its contact within the "client window" are tolerated by the ABC SBC as long as the time-to-live period is not over. A "reg-expired" event is generated. Example of a registration cache view when only the "UAC-side" timeout expires is shown in Figure Client-side Expired Registration Contact.
- the UAC explicitly de-registers using procedures described in RFC3261. In this case a "reg-del" event is generated.

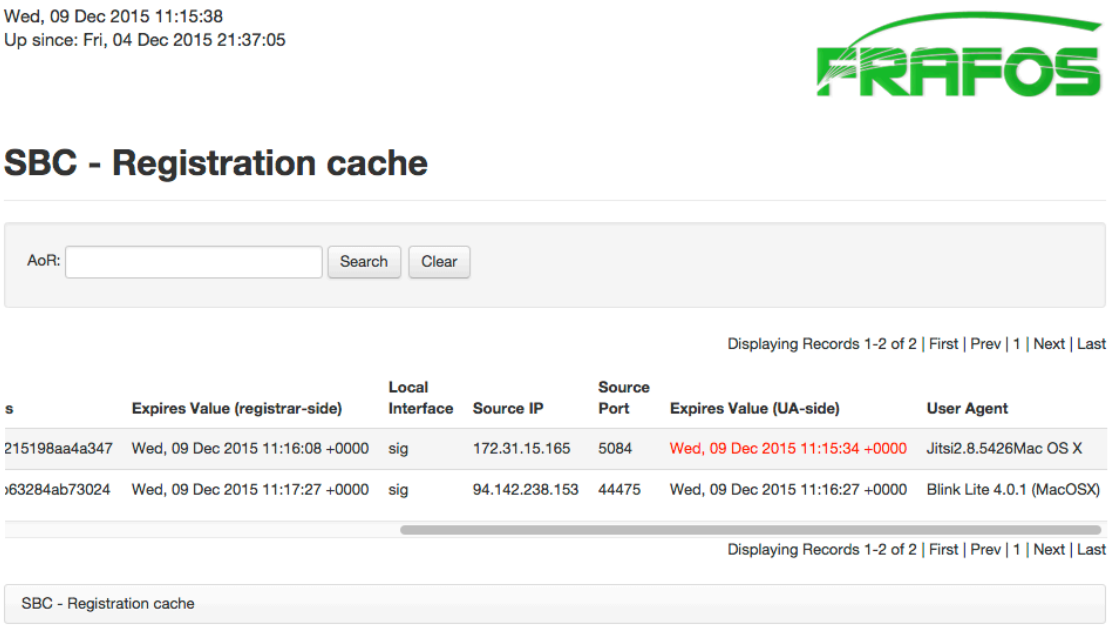

Figure 90: Client-side Expired Registration Contact

Registration Caching and Handling by Example

To enable registrar caching, you must let all REGISTER requests be processed by using the actions **REGISTER throttling ** and ** Enable REGISTER caching**. The throttling action takes additional parameters: **Minimum registrar expiration** and **Maximum UA expiration**. The former specifies the "throttle" which reduces the registration traffic propagated towards a registrar. The value is normally in order of tens of minutes, we chose a whole hour in our configuration example. The other parameter, **Maximum UA expiration**, determines the traffic pace towards the SIP User Agents. The value is normally in order of minutes to keep REGISTER messages flowing and holding IP connectivity through firewalls and NATs upright. We chose an extremely aggressive value in our example, half a minute.

Figure Registrar handling shows a screenshot with such a 3600/30 throttling ratio configuration. The rules are part of a "public" realm serving REGISTERs coming from the public Internet.

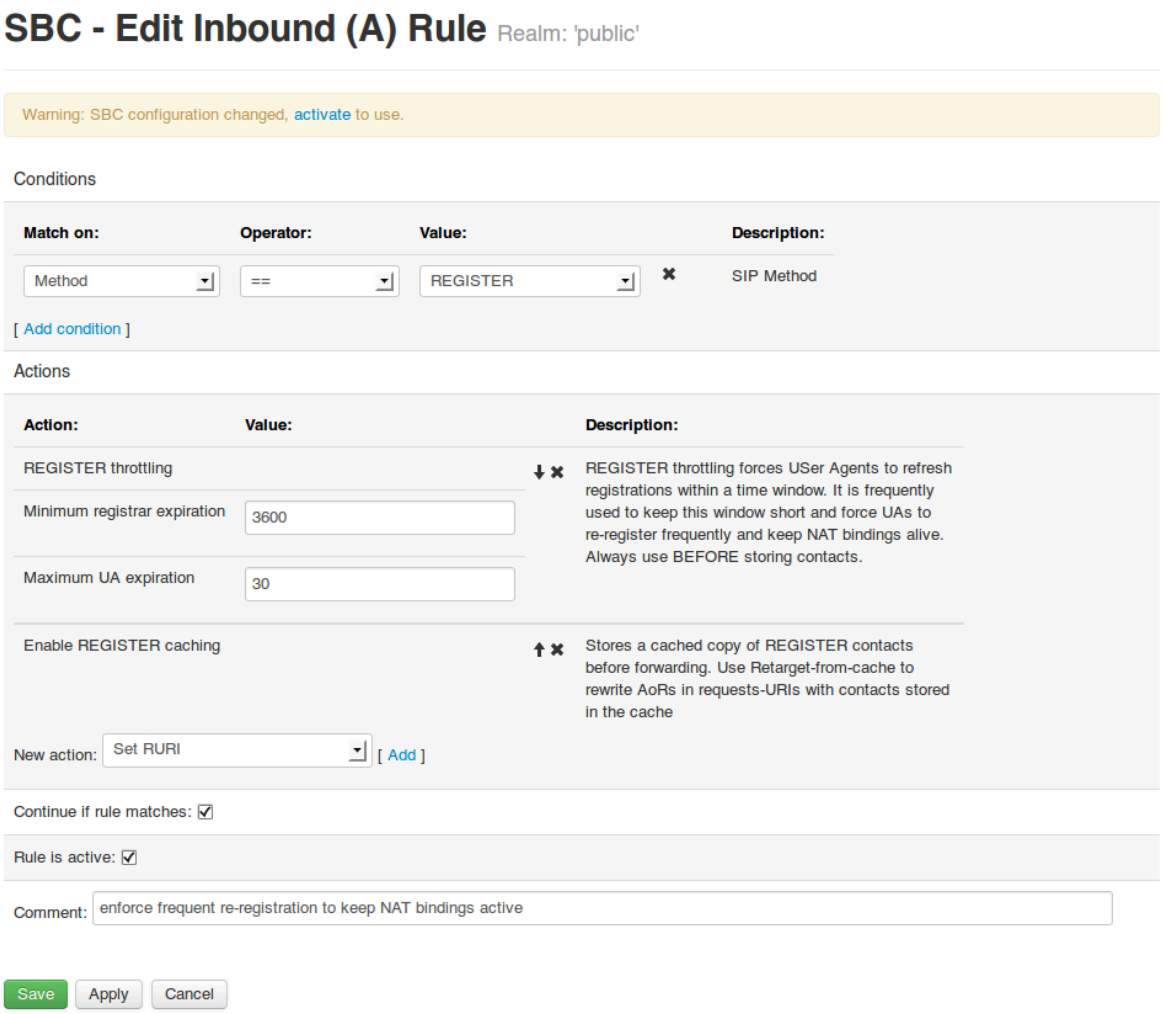

Figure 91: Registrar handling

You also need to configure how incoming messages for the registered users will be processed. Particularly the URIs coming back in incoming requests must be recovered to the original form in the initial REGISTERs received by the ABC SBC. To do so, enable the action **Retarget R-URI from cache**, with the **enable NAT handling** option turned on for all traffic routed to the public realm. The configuration is shown in Fig. Restoring cached contacts.

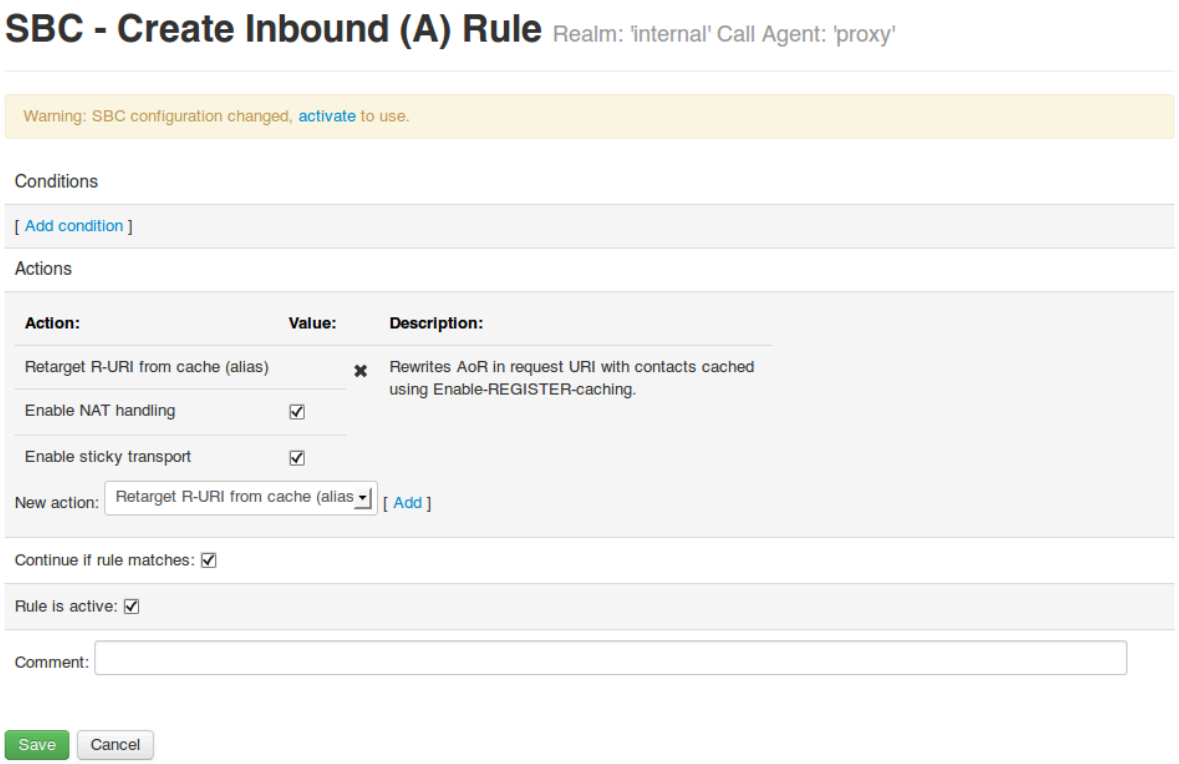

Figure 92: Restoring cached contacts

With this configuration in place, the actual SIP call flows may appear like in the following diagram Call Flow Registration Throttling.

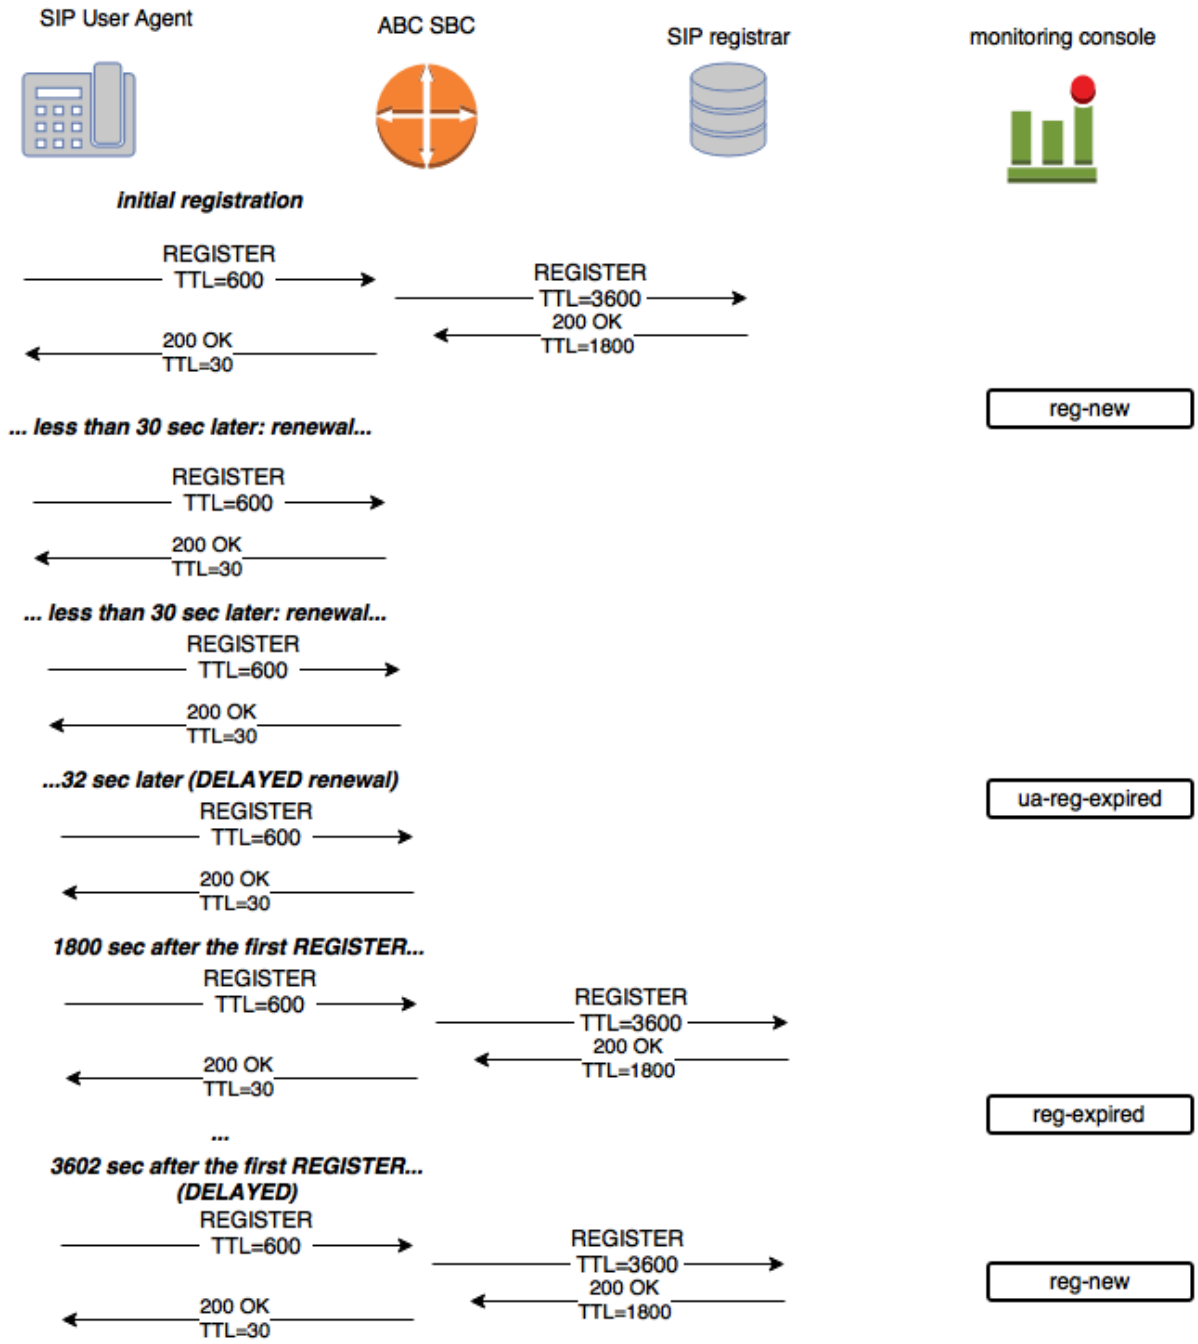

Figure 93: Call Flow Registration Throttling

The example sequence shows the primary function of the registration throttle: increasing the traffic towards SIP User Agent (left-hand side) and reducing it towards the registrar (right-hand side). It also demonstrates how SIP equipment may differ from the expected traffic pattern and how the ABC SBC will deal with it.

The sequence begins with an initial REGISTER. The SIP telephone proposes a “time-to-live” in the message, 600 seconds in this example. (The SIP message element is really called “expires” but we found the “TTL” name more explanatory.) The ABC SBC chooses to send the traffic to downstream registrar less often, and overrides this value to a longer period of 3600 seconds. The registrar downstream finds it too high though and agrees to keep contacts for only 1800 seconds in its 200 SIP response. Now the ABC SBC knows how often it must refresh the registrations downstream: every half an hour.

In the direction towards the client, the ABC SBC compels the client to re-register more often by advertising it would only keep the registered contacts for no longer than 30 seconds.

As a result, the client keeps registering every half a minute, and the ABC SBC passes on the registration to the downstream registrar every half an hour. If the client is late with a renewal request, the address binding remains in ABC SBC registrar cache. If the client is however re-registering too late with respect to the registrar TTL, the cached registration will expire, same way like of there was no cache. The late re-registration will then create a newly registered contact.

Registration Agent

The ABC SBC can register itself with a third-party service using a SIP address of record. That allows it to receive incoming requests for this address subsequently. The ABC SBC does so by sending REGISTER requests periodically and authenticating them if challenged to do so. This may be for example useful, when an ABC SBC installation is configured to use the built-in conference bridge and is also supposed to serve calls from PSTN coming via a SIP-2-PSTN service.

Please note that if transport is needed to be specified, the **Next Hop** “host:port/transport” can be used.

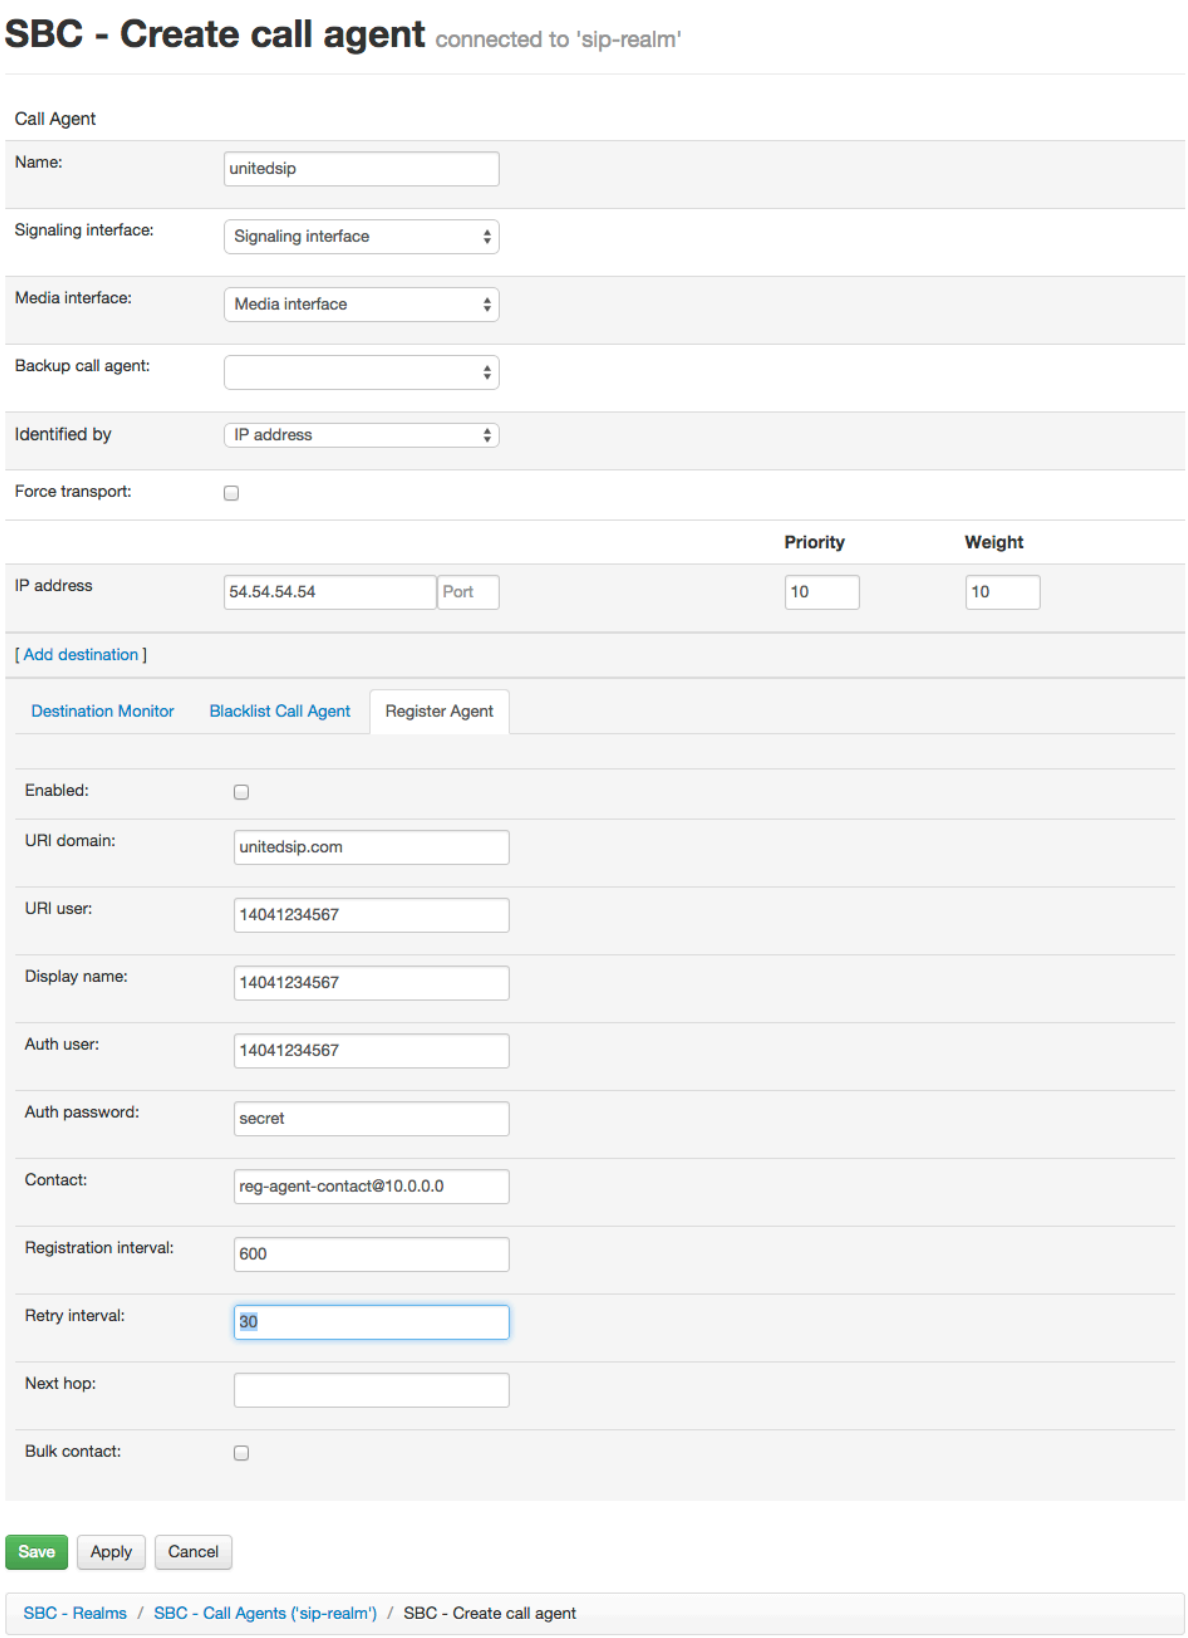

Figure 94: Screenshot of Registration Agent Configuration

The status of registration agent can be inspected under **"Monitoring ▸ Registration Agents"**. For more information see Registration Agents in Live ABC SBC Information section.

4.10.13 Call Data Records (CDRs)

The ABC SBC generates Call Data Records (CDRs) for every call processed by the SBC.

For syslog & conference CDRs (experimental), please refer to New restify CDR process.

CDRs Location

CDRs are generated into the directory:

```
/data/cdr/
```

They are generated on Source Realm basis, so every CDR is filtered to a specific file with the name „*cdr-source_realm_name.log*“. All CDRs also go into one combined file called „*cdr.log*“.

CDR output files are rotated once a day at midnight and exported to the archive directory:

```
/data/cdr/export
```

The exported files are renamed to include the date and time – e.g. ” *cdr.log-201207011200*“. The files are stored for 93 days by default and then are deleted from the disk.

The number of daily rotated files to keep and also the directory for exported files can be changed using Config / Global config, using settings under “CDRs” tab. The lowest possible number of days to keep the exported CDR files is one day.

CDR Format

CDRs are stored in CSV format and contain following items in given order:

- Source Realm
- Source Call Agent
- Destination Realm
- Destination Call Agent
- From user part
- From host part
- From display name
- To user part
- To host part
- To display name
- Local tag (ID for call)
- Timestamp when the call was initiated (format - 2012-05-04 02:22:01)
- Timestamp when the call was connected (format as above)
- End Timestamp of the call (format as above)
- Duration from start to end (sec.ms)
- Duration from start to connect/end (for established/failed call; sec.ms)
- Duration from connect to end (for established call; sec.ms)
- SIP R-URI
- SIP From URI
- SIP To URI

CDR example:

```
pstnprovider.com,gw1,mobile.com,uas,"alice","example.com","","bob","192.168.1.4","",
"6D47CCAA-4FF10747000824C5-80299700","2012-07-02 04:28:23","2012-07-02 04:28:28",
"2012-07-02 04:28:33","10.139","4.895","5.244","bob@192.168.1.4:6000",
"alice@example.com", "bob@192.168.1.4"
```

Access to CDRs

CDRs can be accessed from the host where ABC SBC container is running, under the container filesystem /data/cdr sub-directory.

To get only exported files (i.e. files that are not updated any more and are ready for post-processing), use the /data/cdr/export sub-directory.

Customized CDR Records

The content of CDR records can be changed using *CDR format* global config variable. Following items can be used in the option:

- $srcrlm.name - Source Realm
- $srcca.name - Source Call Agent
- $dstrlm.name - Destination Realm
- $dstca.name - Destination Call Agent
- caller_id_user - From user part
- caller_id_host - From host part
- caller_id_name - From display name
- callee_id_user - To user part
- callee_id_host - To host part
- callee_id_name - To display name
- $ltag - Local tag (internal call identifier)
- $start_tm - Timestamp when the call was initiated (format - 2012-05-04 02:22:01)
- $connect_tm - Timestamp when the call was connected
- $end_tm - End Timestamp of the call
- $duration - Duration from start to end (sec.ms)
- $setup_duration - Duration from start to connect/end (for established/failed call; sec.ms)
- $bill_duration - Duration from connect to end (for established call; sec.ms)
- sip_req_uri - SIP R-URI
- sip_from_uri - SIP From URI
- sip_to_uri - SIP To URI
- disposition - Result of call establishment. Possible values: answered, failed, canceled.
- invite_code - Result code of final reply to initial INVITE (not set for canceled calls).
- invite_reason - Reason phrase of final reply to initial INVITE (not set for canceled calls).
- hangup_cause - Reason for the call termination, set for established calls only. Possible values: BYE, reply, no ACK, RTP timeout, session timeout, error, other.
- hangup_initiator - Reason for the call termination. It is set for answered calls only where hangup was caused by request (BYE) or in case of call was terminated because of local error. Possible values: caller, callee, local
- ucid - Unique Call Identifier. Can be used to map CDRs for transferred calls together. The value is common for all CDRs generated for calls that are results of unattended call transfer from one original call done by the ABC SBC.
- call variable - $gui.<call variable name> - Allows to write user specified call variable into CDR. For example $gui.experimental_variable will write the value of experimental_variable into CDR. These outputs are located under "/data/cdr/cdr.log". The call variable can be set using "**Set Call Variable**" action, see Binding Rules together with Call Variables for more details on using call variables.

Please make sure "List of call variables added into events:" section is filled by requested call variable to see in CDR. Under "Call Events" table, ADVANCED button should be clicked to be able to see call variable on Monitor as requested.

Important: CDRs are written using syslog and split into per-realm files according to first item that is expected to be the source Realm. If you change the CDR format the way the first item is not a "realm"

name, the files will be named according to the values in this column and won't represent per-realm data any more.

4.10.14 Advanced Use Cases with Provisioned Data

The ABC SBC can be integrated with external or internal sources of data and logic. This allows to complement its rigid rules-based logic with richer and more complex applications provisioned by the administrator.

The following methods are available:

- **Generic RESTful queries** to an external server using the **Read call variables over REST** action. Using this interface allows to drive the ABC SBC behavior using in-house developed business logic located in external web programming environments. see Section RESTful Interface for more details.
- **Provisioned tables**. Solving some problems with tabular nature, like Least-Cost-Routing, Blacklisting, Dial Plan Normalization or SIP Connect Bulk Registration is much easier if tables are provisioned separately from the rules. For this reason the ABC SBC supports on-board provider-provisioned databases. See Section Provisioned Tables for more details.
- ENUM queries using the **Enum query** action , as described in the section ENUM Queries. This method allows for a number-to-URI translation using the DNS-based ENUM queries.

RESTful Interface

The RESTful interface embedded in the ABC SBC allows high programmability of the SIP Session Border Controller.

The interface addresses an important dilemma for operators: how to introduce new scenarios, while preserving the existing ones intact. Hardwired product logic compels the operators to request code changes from vendors. Change requests result in unavoidably tedious process, weeks or months of negotiation, changes of changes and delays regardless how small and reasonable the changes are. Therefore ABC SBC comes with the possibility to implement business logic outside the product in an operator-controlled environment.

This capability follows a general trend in which the business logic is concentrated in a single place that defines behavior of relatively "dumb" network elements. The business logic defines security policies (who can call whom), marketing campaigns (at what price), and network behavior (how to route the calls). Placing this logic in a web server relieves operators from inadequate vendor dependencies and allows PHP, Perl, and virtually any web programmer to implement new SIP scenarios in a well understood programming environment.

The operation of a RESTful application is simple and consists of three steps characteristic for any computer program. The steps are depicted in Figure RESTFul Call Flow: The ABC SBC receives an incoming INVITE on input in the first step A, processes it in step B, and generates a correctly processed INVITE on the output in step C. The processing in step B is split in three phases:

- B.1: RestFul query is formulated that contains all pieces of SIP information needed to execute the web-based logic. The information is passed in form of URI parameters to the RESTful server.
- B.2: The RESTful server performs the application logic.
- B.3: Eventually the RESTful server sends an answer back to the ABC SBC. The answer contains an array of variables that represent an advice to the ABC SBC how to handle the message in the final step C.

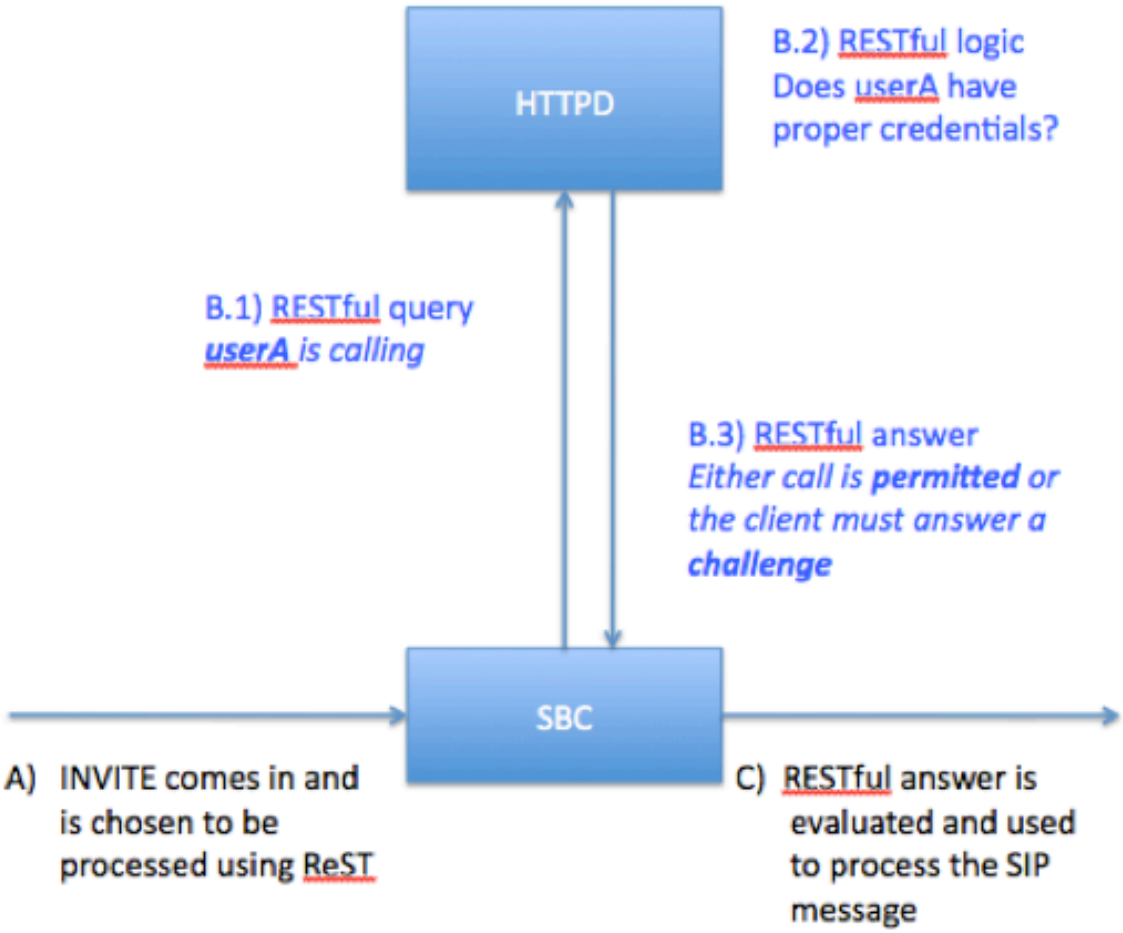

Figure 95: RESTFul Call Flow

RESTful Interface using Digest Authentication Example

In this example we show how to outsource digest authentication to the external RESTful server. This relieves the ABC SBC of implementing a user:password database. It is even designed in a way that leaves the ABC SBC unaware of the cryptographic authentication protocol: all it does is it shuffles header-fields back and forth.

Background: digest authentication in SIP works by conveying shared password from client to server in a hashed form. If both client and server hash the password and obtain the same result, identity of client is proven without sending the password in clear-text.

The whole process follows the steps outlined above. It begins when a SIP request comes in. (only fragment shown):

```
INVITE sip:music@abcsbc.com SIP/2.0.
 Via: SIP/2.0/tcp 192.168.178.22:54251
 From: "foo" <sip:foo@abcsbc.com>;tag=0omQsGfHsCtP8-5k2.t4uJI3ekc66bGZ.
 To: <sip:music@abcsbc.com>.
 CSeq: 14693 INVITE.
 Proxy-Authorization: Digest username="foo", realm="abcsbc.com",
   nonce="UP6CiVD+gk0Uyu4WHAv+48ypPC2vjH+6", uri="sip:music@abcsbc.com",
   response="560ad1cc8777efa6a6cc1857795ec155".
```

The INVITE request signals a call from user with address sip:foo@abcsbc.com to address sip:music@abcsbc.com The ABC SBC checks the request against its rules and initiates the RESTful logic. (see Figure Rule for Evoking a RESTful query).

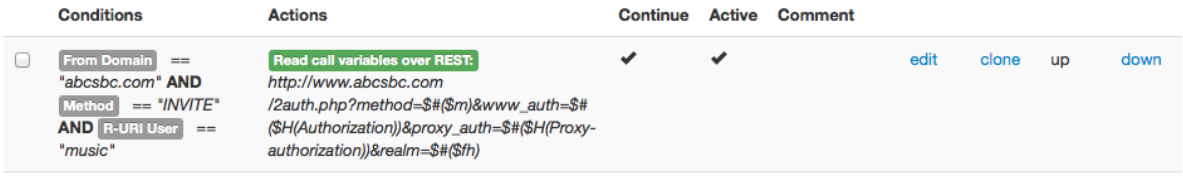

Figure 96: Rule for Evoking a RESTful query

The rule in ABC SBC's configuration matches by From domain, method and request URI, and therefore processing is passed to the action "Read Call variables over REST". ABC SBC is configured to pass several header fields as URI parameters to the RESTful application. Particularly, the user information relevant to authentication are passed: Authorization and Proxy-authorization header fields ($H(Proxy-authorization)), request method ($m) and realm. The domain in From URI ($fh) is used as realm – this way you can build up a multi-domain hosted service, which will work same for any domain without change.

Now the SBC has received a call, chosen to process it using a web server, the REST can begin by sending an HTTP query. Let's see how the query looks on wire. It simply conveys the values chosen in SBC configuration as URI parameters. Symbols contained in the values are substituted using the escape code %:

```
GET /2auth.php?method=INVITE&www_auth=&proxy_auth=Digest+username%3d%22foo%22%2c
+realm%3d%22abcsbc.com%22%2c+nonce%3d%22685f3174-6aaf-4337-a9f4-4cf4d1f150ab%22%2c
+uri%3d%22sip%3amusic%40abcsbc.com%22%2c
+response%3d%225b3cc376e815c949bbc084c747a3a55f%22&realm=abcsbc.com HTTP/1.1.
User-Agent: REST-in-peace/0.1.
Host: www.abcsbc.com
```

When the web server receives this query, it starts an application. In our example we have chosen to build it using PHP, it could be done same well using Perl, Java, python or any other popular web programming language. In its own way the interpreter passed the URI parameters to application's variable and processing begins.

The following PHP code shows key steps during computation. Not all are shown. For a given user, it calculates her hashed password stored in database and checks it against one coming in the request. If they are equal, it answers with a 200 answer suggestion, otherwise it advises the SIP server to re-authenticate the user:

```
// simulation of a database query … ask username and password
$users = array('foo' => '12', 'guest' => 'guest');
// prepare challenge for the case credentials are invalid
// or missing
$challenge='Digest realm="'.$realm.'",nonce="'.new_nonce().'"'

// no credentials supplied? Request some!
if (empty($proxy_auth)) {
  print_answer("407", "Authenticate",
         "Proxy-Authenticate: ".$challenge,$cmt);
}
// parse the credentials
$data=parse_hf($proxy_auth);
// calculate expected answer
$expected_response=calculate_answer($data);
if ($data['response'] != $expected_response) {
  print_answer("407", "authenticate", "Proxy-authenticate",
                $challenge);
  return;
};
// otherwise proceed with OK
print_answer("200", "ok");
```

Now we have an answer: either a positive 200, or a negative 407 with authentication challenge to be passed to the SIP client through the ABC SBC. If you observe the wire you will see the following HTTP answer:

```
HTTP/1.1 200 OK.
Date: Tue, 22 Jan 2013 11:58:47 GMT.
Server: Apache/2.2.3 (Debian) PHP/4.4.4-8+etch6 mod_ssl/2.2.3 OpenSSL/0.9.8c.
X-Powered-By: PHP/4.4.4-8+etch6.
Content-Length: 214.
Content-Type: text/html.
.
code=407.
phrase=authenticate.
headers=Proxy-Authenticate: Digest realm="abcsbc.com",
   nonce="685f3174-6aaf-4337-a9f4-4cf4d1f150ab"\r\n.
```

What you see here in clear-text is, that the programmer has stored the processing results into several variables that are passed back to the ABC SBC rule-base.

We are in the final stage now – the web application has returned processing results back to ABC SBC, the SBC will evaluate the parameters in its rules and use them in further SIP processing. Particularly, the rule in Figure Rule for processing result of RESTful query says, if processing ended up with variable code not being equal to 200, the call will be refused. The negative answer will include parameters determined in the HTTP answer.

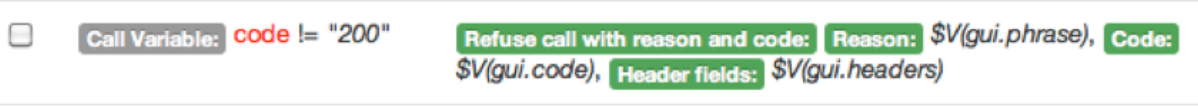

Figure 97: Rule for processing result of RESTful query

Here is the final outcome of our effort: SIP answer calculated in the RESTful server and challenging SIP client to submit proper credentials:

```
SIP/2.0 407 Authenticate.
Via: SIP/2.0/tcp 192.168.178.22:54251;received=83.208.91.146
From: <sip:foo@abcsbc.com>;tag=0omQsGfHsCtP8-5k2.t4uJI3ekc66bGZ.
To: <sip:music@abcsbc.com>;tag=50123210ee9d5f0f8df05cf1f196cfeb-c5a6.
CSeq: 14692 INVITE.
Proxy-Authenticate: Digest realm="abcsbc.com", nonce="UP6CiVD+gk0Uyu4WHAv+48ypPC2vjH+6"
```

Provisioned Tables

The ABC SBC rules can refer to an internal database maintained separately from the rules logic. This greatly simplifies use-cases which would have to be implemented using a large numbers of almost identical rules otherwise. The typical use-cases include tests if a URI is on a blacklist or list of monitored users, static SIP registrations, Least Cost Routing tables, definition of dialing plan normalization and more.

The tables are physically located on the ABC SBC machine for highest performance, can be provisioned using the web interface and can include any number of administrator-chosen attributes in addition to the lookup key. There is also a possibility to provision the data remotely via RPC or REST API.

There are two types of tables:

- **data** - general purpose data tables can be queried to fetch specific data associated with a key. The structure of such tables can be freely defined by the administrator, thus allowing great flexibility. The data tables are used from A-rules and C-rules.
- **routing** - specialized routing tables have a list of mandatory attributes that define routing behavior and are always present. Additional attributes may be added. The routing tables can only be used from B-rules.

Using the tables is as simple as creating a table with the desired structure, filling it with data, looking up a result in the table from A,B or C-rules by a selected value, and processing the found data entry. The data entry is returned to the script processing as variables bearing the names of the table columns. The whole process is described in the following subsections in detail.

Please be aware that restful provisioned tables queries have some limitation compared to the one available via GUI. The rest matching cannot emulate a case insensitive match as the data are stored in a redis database, while it's a SQL for the later one.

Configuring Tables

The process of setting up a new provisioned table consists of the following steps:

- **Analysis of the problem to be solved**. You need to specify what data you are going to lookup in a table by what key and how you are going to use the resulting table record.
- **Definition of the table structure**. This is started from the Web menu under *Provisioned Tables* ‣ *configure* ‣ *Insert new*. There you must identify:
  - ‣ key lookup operator which is one of *equal*, *range*, and *prefix*. The operator defines the method by which a key is looked up in a table. It is not possible to lookup in the same table using some other method. If *range* is chosen, the resulting table will include two key columns for begin and end of a range. If *prefix* is used and overlapping choices are found, the longest match is selected.
  - ‣ table keys and their types. For *range* and *prefix* lookup operator just one key could be defined. For *equal* operator multiple keys could be defined. The key type is one of *uri*, *number*, *call-agent*, *string*. The type is used for syntactical checking when the actual data is entered later. Even more importantly it is used to determine how the lookup operator is used. Particularly, prefix lookup is sensible for string types as it discriminates between "0" and "00". For numerical types, these two values would be the same.
  - ‣ and type of table *data* or *route*, as explained above.
  - ‣ There is also Group-by, which can be *none* or *string*. The option *string* allows to add an informational tag to each entry so that tables can be viewed by groups.

  - Optionally, any number of additional table named columns can be added, whose type must be chosen from *uri*, *number*, *call-agent*, *string*. When the "save" button is pressed, the table structure is created and becomes instantly ready for filling with data.
- **Filling tables with data**. This is started from the Web menu under *Provisioned Tables* ‣ *(table name)*. On the web page that opens, the link *Insert new rules* opens a dialog for inserting a new data entry. When editing the new entry is complete, pressing the "Save" button will store it. A Version 4 UUID (**RFC 4122**) is automatically added to every entry for sake of internal data maintenance.
- **Completing data entry**. To make the ABC SBC understand that the newly created records can be used, the button *Activate Changes* must be pressed. This allows editing the tables without interfering with the table as currently being used by the ABC SBC. *Activate Changes* activates the current version of the table for use by the ABC SBC, and creates a new table version which is used for further provisioning.
- **Introducing the table lookup in the rules**. This requires adding the action *Read Call Variables*. The action takes the table name as the first parameter and lookup value as the second. The lookup value is typically formed using substitution expressions (see Section Using Replacements in Rules for a reference.)

Provisioned Table Example: Static Registration

We will start with a relatively simple example: static registration. Similarly to how SIP devices use SIP REGISTER messages to create a temporary binding between SIP Address of Record and actual Contact URI, administrators can provision a similar association manually and permanently. That means that a table is provisioned and used the following way:

- The table is keyed by an Address of Record URI and includes the next-hop URI as attribute.
- When a SIP INVITE comes in, its request URI is checked against the table and if the next-hop is found, replaced with it.

Note that this structure can be used to implement static call-forwarding.

Screenshots of the resulting table structure, table content and rules using the table are shown in the Figures Structure of Static Registrations, Static Registration Records, and Static Registration Rules respectively.

You can make several observations about the table lookup rule:

- The rule conditions make sure the lookup is only executed for authenticated INVITEs, which helps to eliminate unnecessary database queries - there is no point in making the table query for other than INVITE requests. If authentication is requested, the lookup for the first unauthenticated request would be also useless.
- The table entry is looked up by the replacement expression "$r." which stands for the current request URI. The result is returned in variable next_hop, as defined in the table column name.
- If no result is found, the table lookup placed in the rule's condition will return FALSE and no action will not be performed.
- If a result is found, we apply two actions: change request-URI and add a header-field for troubleshooting purposes. Its presence in the outgoing INVITE shows a lookup was performed, the request-URI which was used as key and the returned value.

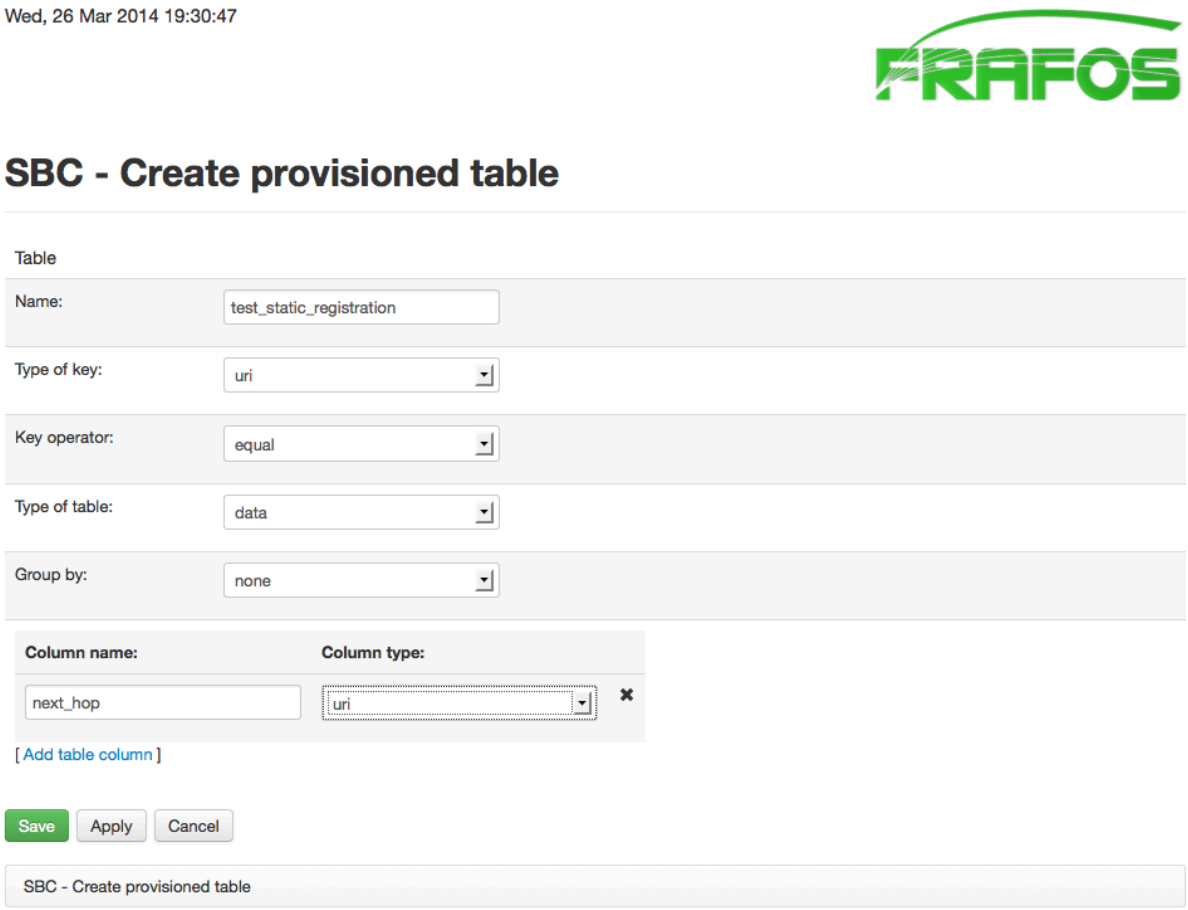

Figure 98: Structure of Static Registrations

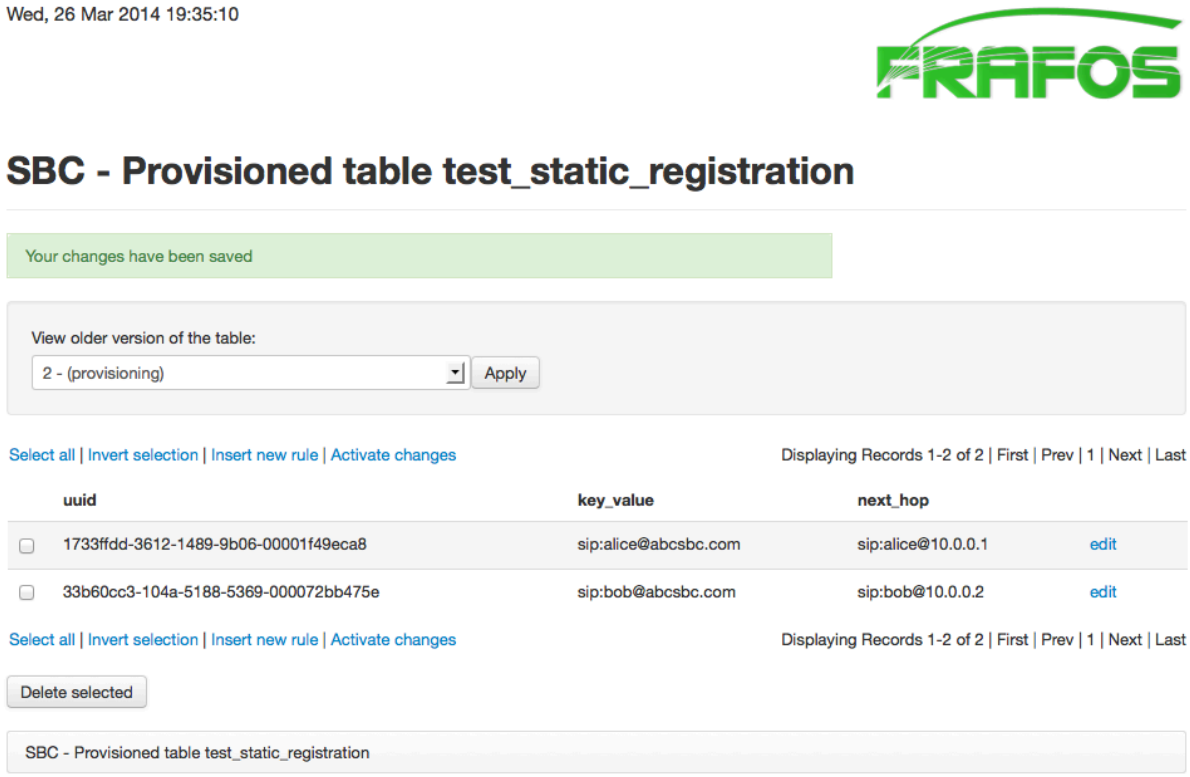

Figure 99: Static Registration Records

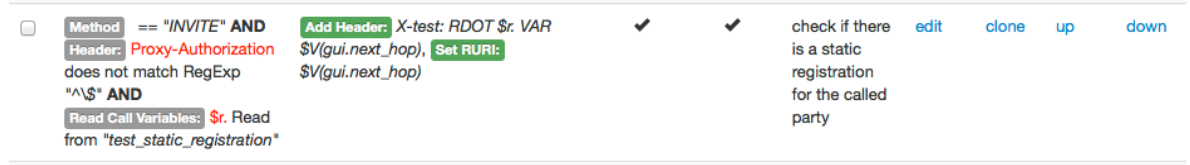

Figure 100: Static Registration Rules

Provisioned Table Example: URI Blacklist

Simple tables can have a great use. In this example we test presence of a SIP element value on a list. That means that the table only includes keys, with which no additional values are associated. We then look up the elements in the keys. That can be used to implement scenarios involving all kind of discrimination like:

- Call recording: is a call coming from a user whose calls shall be recorded?
- Domain discrimination: is a call being routed to a listed domain for which some header fields must be removed or appended?
- URI Blacklisting: is the caller blacklisted?

The following screenshots show the configuration of the URI-blacklisting example: Figure URI Blacklist Structure, Figure URI Blacklist Content, and Figure Blacklisting Rule. The rule is simple: If the SIP URI in the From header-field matches a URI in the blacklist table, the request is declined using a 403 response. Note that the lookup key is concatenated using “sip:” and “$fu”, because the replacement expression $fu does not include a protocol discriminator.

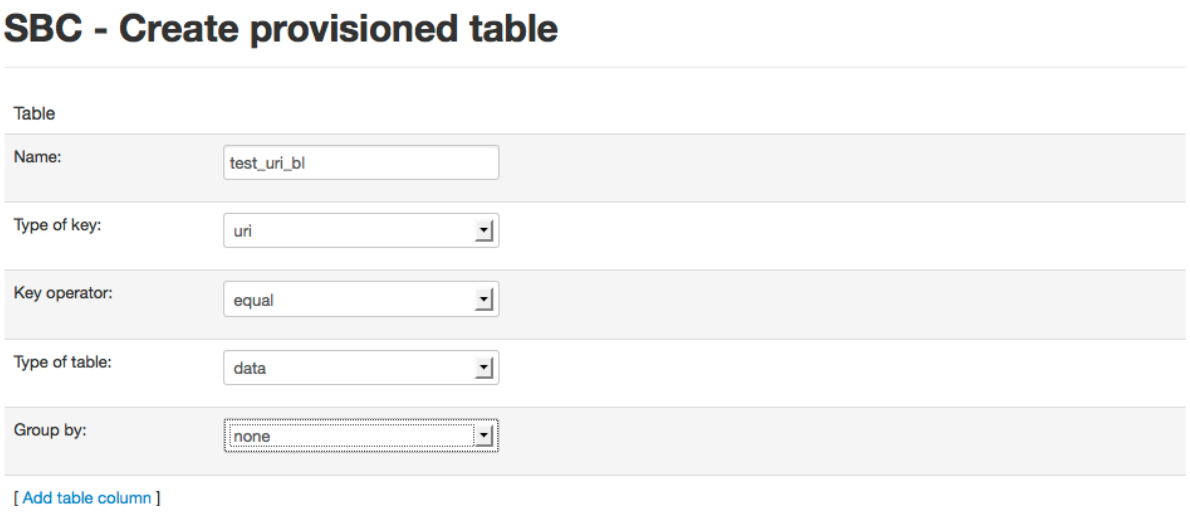

Figure 101: URI Blacklist Structure

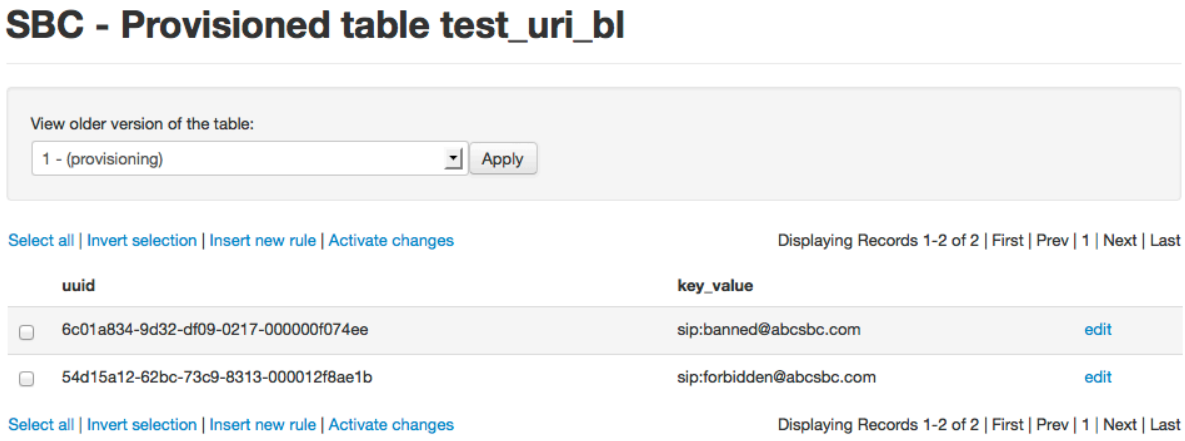

Figure 102: URI Blacklist Content

Figure 103: Blacklisting Rule

Table Example: Dialing Plan Normalization and Least-Cost-Routing

This is a two-in-one example showing two tables that are usually cascaded behind each other: normalization of a PBX dialing plan and least-cost routing.

Telephone numbers as used within a PBX can have different forms, following local national conventions and enterprise policies. For example, a typical user of a PBX in Munich dials with three leading zeros followed by international and area code to reach an international destination, two zeros followed by area code to reach destinations within Germany, one zero to reach destinations within Munich metropolitan area, and phone numbers without leading zeros to reach other PBX users. Using this dialing convention is convenient, the number length only grows with distance. However, these numbers loose significance if one tried to use them globally say to reach an international PSTN gateway. Therefore it is useful to normalize them in the E.164 format by stripping leading digits and introducing an appropriate prefix.

The following table shows examples of telephone numbers and how they are normalized for calls from a Munich PBX:

local number number	E.164 equivalent	digits to be stripped	prefix to be introduced
000140433345678 (US destination)	+1-404-333-45678	3	+
003034567000 (German number)	+49-30-3456-7000	2	+49
078781234 (Munich number)	+49-89-7878-1234	1	+4989

The following screenshots show the configuration of dialing plan normalization: Figure Dialplan Structure, Figure Dialplan Content, and Figure Dialplan Rules.

Note that the key is defined as string to make sure that prefix "00" in request URI does not match all of "0", "00" and "000" as it would if the data type would be numerical.

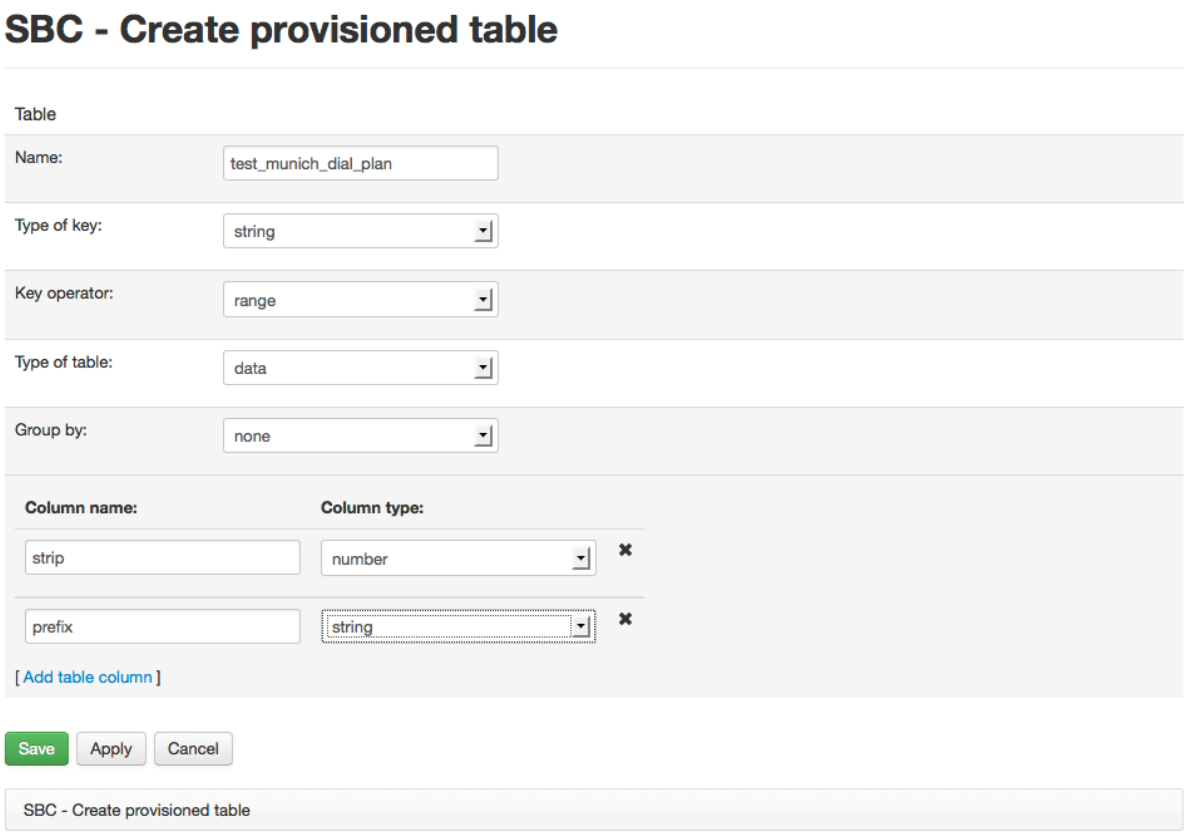

Figure 104: Dialplan Structure

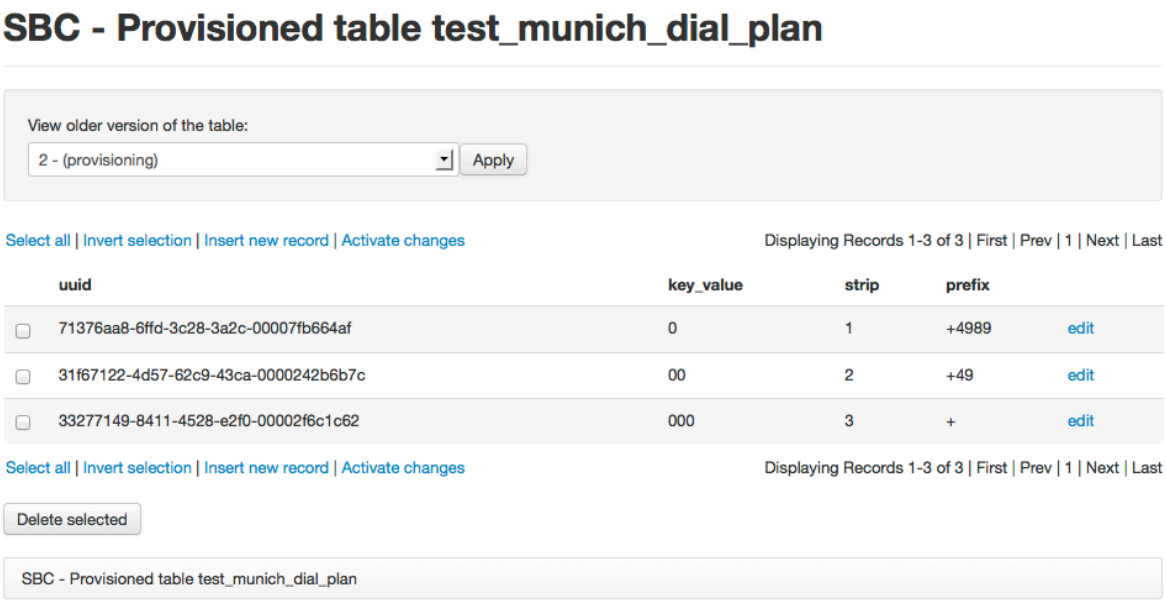

Figure 105: Dialplan Content

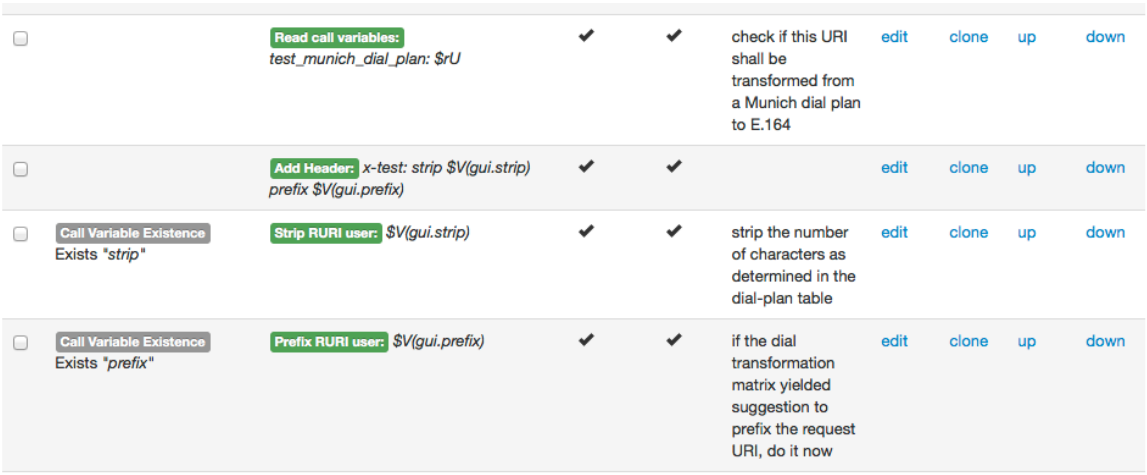

Figure 106: Dialplan Rules

Once the numbers are normalized in the E.164 form, it is also easy to check the destination against a least-cost routing table to find the most economic PSTN gateway. The table may have the following content: prefix that is used to march phone numbers, and DNS name of a gateway chosen to serve the matched destination. Longest match applies which means that the shortest-match is taking lowest precedence and is used as "default route".

prefix	destination	comment
+1	us-gateways.com	US destinations
+43	austrian-united.com	German destinations
+	cheap-pstn.net	Everything else

The provisioning process is shown in the following three Figures: Creating an LCR Table, Creating LCR Table Entries, and Calling the Routing table from routing rules.

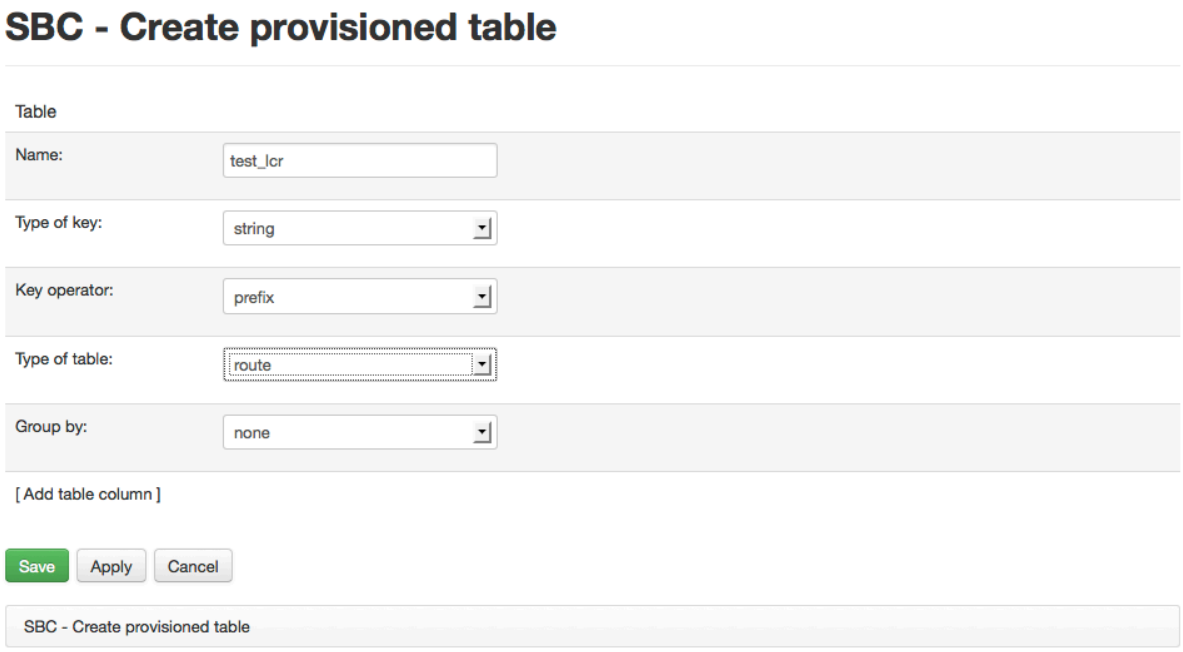

Figure 107: Creating an LCR Table

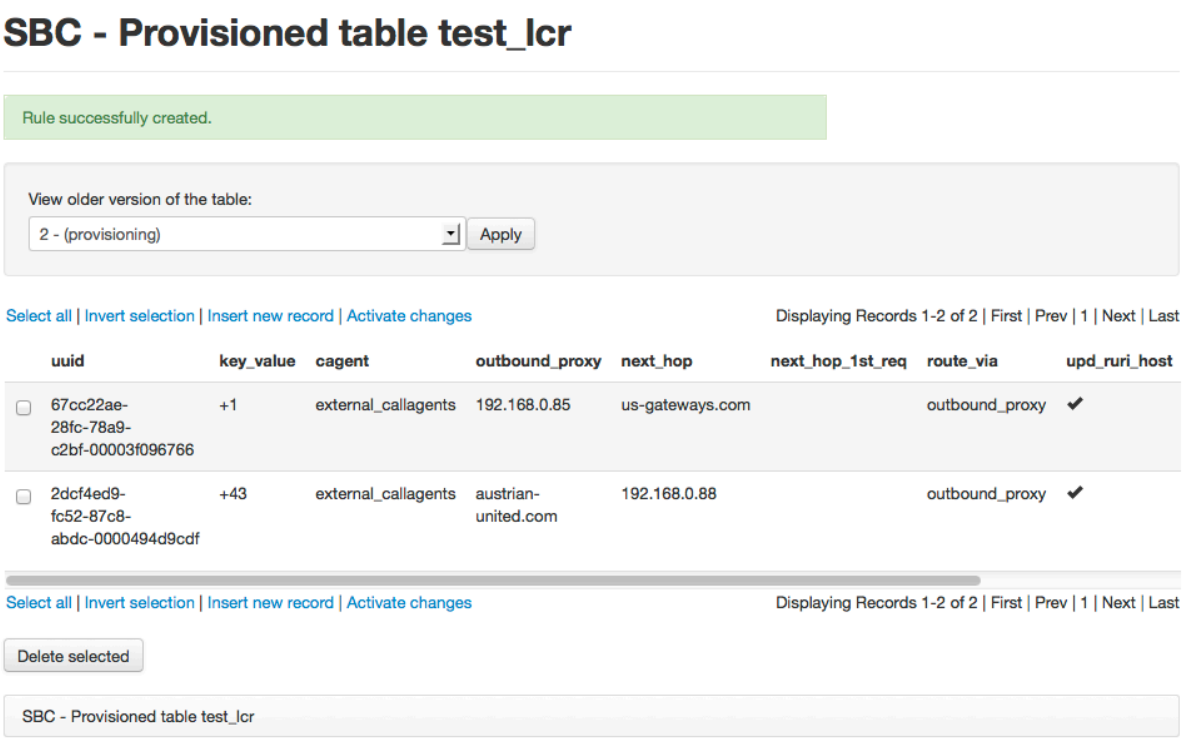

Figure 108: Creating LCR Table Entries

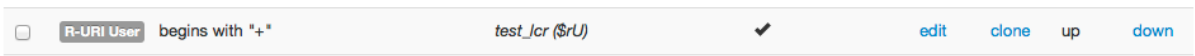

Figure 109: Calling the Routing table from routing rules

Table Example: Bulk Registration

Thanks to an extensions of the SIP standard, it is now possible for a PBX to have one digest identity under which it can serve a whole range of telephone numbers. The extension emerged out of SIP Forum's

effort to create a profile for PBX interoperability, that became known as “SIP Connect” and standardized as **RFC 6140**.

The ABC SBC supports these scenarios and it even makes the deployment scenario much simpler than contemplated in the RFC. It allows an arbitrary SIP client, PBX, softphone or any other SIP device to authenticate under a single URI and receive calls for a whole range of telephone numbers. It works “as is” without requiring any of the “bnc”, “gin” or GRUU extensions.

The following example call flow assumes a network topology in which the ABC SBC guards an internal network, in which a combined proxy/registrar is located. The administrator has provisioned the telephone number range 7200-7400 to be server be PBX reachable under the URI sip:pbx@abcsbc.com. Note that a similar scenario could also be implemented using the ABC SBC's built-in registrar.

The call flow starts with a SIP registration using digest authentication (1)-(4). When an INVITE comes in (5), the telephone number in the Request URI is translated to that of the PBX (6). This allows the proxy/registrar behind the SBC to perform user-location lookup and forward to the PBX through the ABC SBC (7). The SBC then, as usual, retrieves the original URI, whose username is fixed eventually to be the target telephone number (8):

To orchestrate this call-flow, the following configuration steps must be taken:

- A number range must be defined and assigned to the URI the PBX owns. This is done using the table-provisioning feature. The screenshots showing this process are in the Figure Definition of the Number Range Association and Figure Assignment of a Number Range to a URI.
- In a rule, incoming INVITEs (5) must be tested against the available ranges. If such a range is found, the request URI must be translated to that owned by the PBX. At the same time the telephone number must be preserved in a request-URI parameter and/or proprietary header-field (x-pbx-user here), whichever the registrar behind the SBC can better deal with. (6) The configuration is shown in Figure Assignment of a Number Range to a URI.
- Before the INVITE is eventually sent to the PBX, it must include the destination telephone number in the request URI. This is done in a rule that retrieves the phone number from the request URI parameter or header-field, in which it was stored in the previous step. The configuration is shown in Figure Retrieving the Telephone number back in request URI.

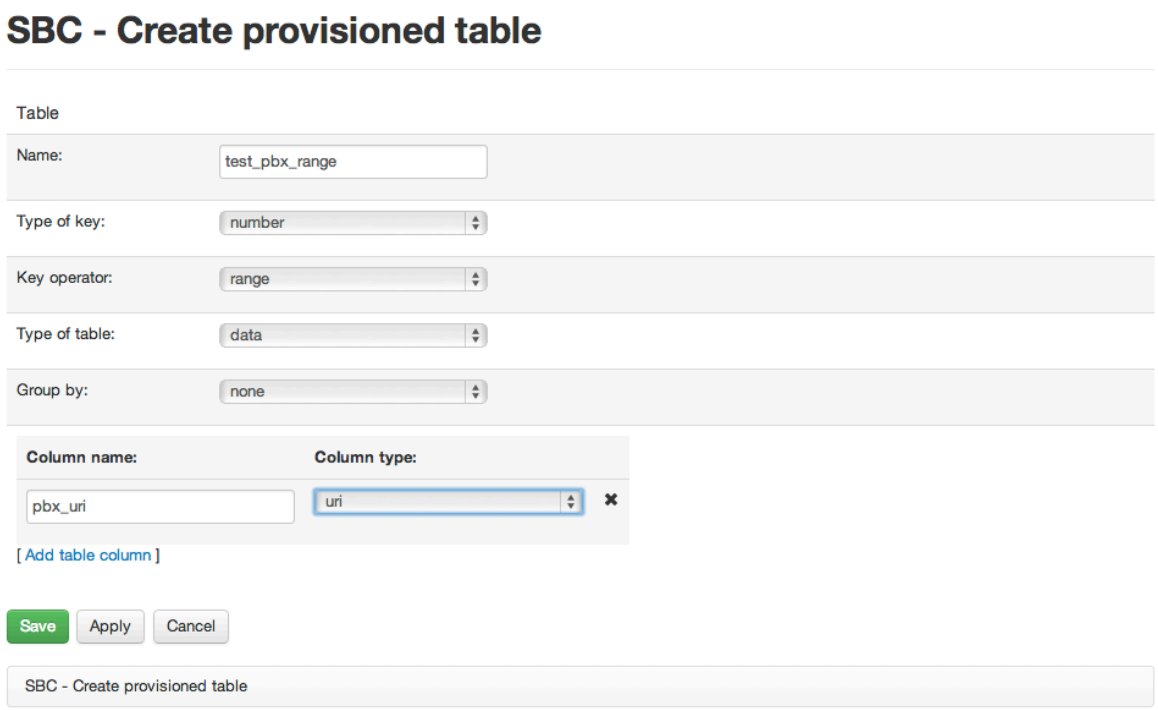

Figure 110: Definition of the Number Range Association

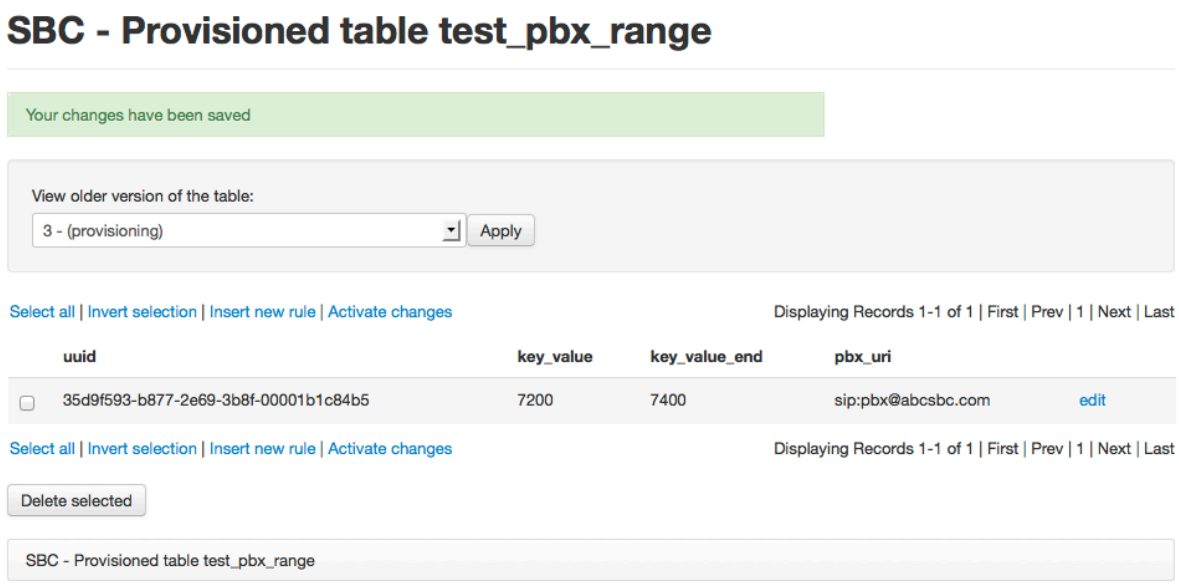

Figure 111: Assignment of a Number Range to a URI

Figure 112: Placing PBX's Address in request URI and Storing the Original Telephone Number

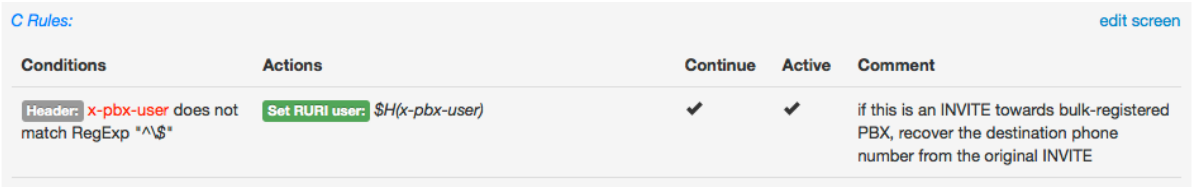

Figure 113: Retrieving the Telephone number back in request URI

Provisioning Tables Using RPC or REST API

In the case that the ABC SBC administrator already has a table available, it will be easier to transfer it automatically to the ABC SBC as opposed to typing it in the web-interface. This can be accomplished using the ABC SBC's XML-RPC data provisioning interface.

Check following sections for more information: Reference XML-RPC functions and its Provisioned Tables sub-section.

The description of REST API can be found in API reference.

ENUM Queries

ENUM (**RFC 3761**) is a DNS-based phone number database that translates telephone numbers into URIs. For example, the telephone number +1-405-456-1234 can be translated to sip:mrs.somone@abcsbc.com. This is often used to find SIP address for a VoIP user when she receives a call from the PSTN under her telephone number.

An ENUM query can be run using the **Enum query** action. This action queries the default DNS resolvers configured in the SBC host with an Enum query, and sets the request URI to the result.

Figure 114: Using Enum queries

By using a different domain suffix than the default one (e164.arpa), private enum servers can be queried. This is in fact the way ENUM is widely used – as of today, no public ENUM service with global coverage has emerged.

The result of the Enum query can be tested using the **Last Action Result** condition. If it returns true, the ENUM query returned a URI, false is returned otherwise. In case of success, the ENUM-returned URI has rewritten the request-URI and may be rewritten using the **($rU)** replacement expression.

4.10.15 SIP-WebRTC Gateway

WebRTC is a relatively new protocol suite added to the VoIP technology that makes a telephone out of every capable web browser. As a result, users can click-to-dial a company representative, easily access video-telephony from within other web applications and receive calls from any web-browser, be it on their PC, smartphone or Internet cafe.

All of that while enjoying confidentiality widely available to consumers as never before in telephony's history. Both analog and digital telephony were inherently insecure, mobile telephony secured at least the wireless hop, yet rather weekly. SIP's security protocols, PGP, S/MIME and Identity (**RFC 4474**) desperately failed to be adopted. With WebRTC, we have proven web-based cryptographic protocols that just work!

The key missing piece for connecting Web clients to the SIP telephony is a SIP-WebRTC gateway – see the left-most element in the Figure Integration of RTC, SIP and PSTN Networks using the RTC Gateway. The gateway connects the populations of web users, SIP telephony users and traditional telephony users behind PSTN gateways. The gateway also provides a practical and yet fairly secure communication model: on the "internal" SIP-based side of the gateway, traditional IT practices for securing controlled networks can be used, while on the public Internet facing side proven cryptographic protocols are used. That is where the ABC SBC comes in: its border control instruments in combination with built-in RTC gateway allow to form a viable security model.

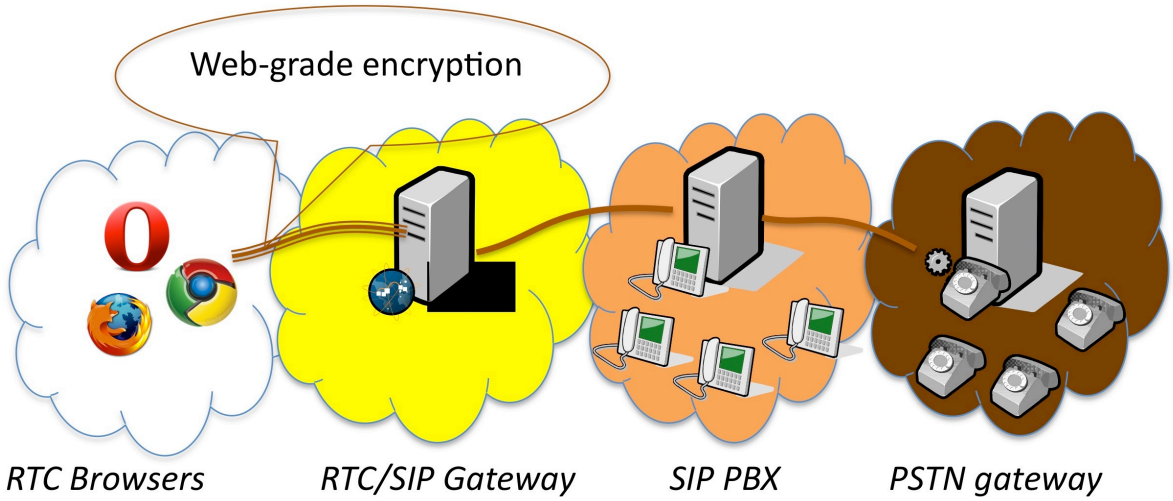

Figure 115: Integration of RTC, SIP and PSTN Networks using the RTC Gateway

The gateway anchors signaling and media and performs translation between different standards for WebRTC and traditional VoIP, particularly security, codecs and signaling protocols as shown in Figure WebRTC Gateway Protocol Stack.

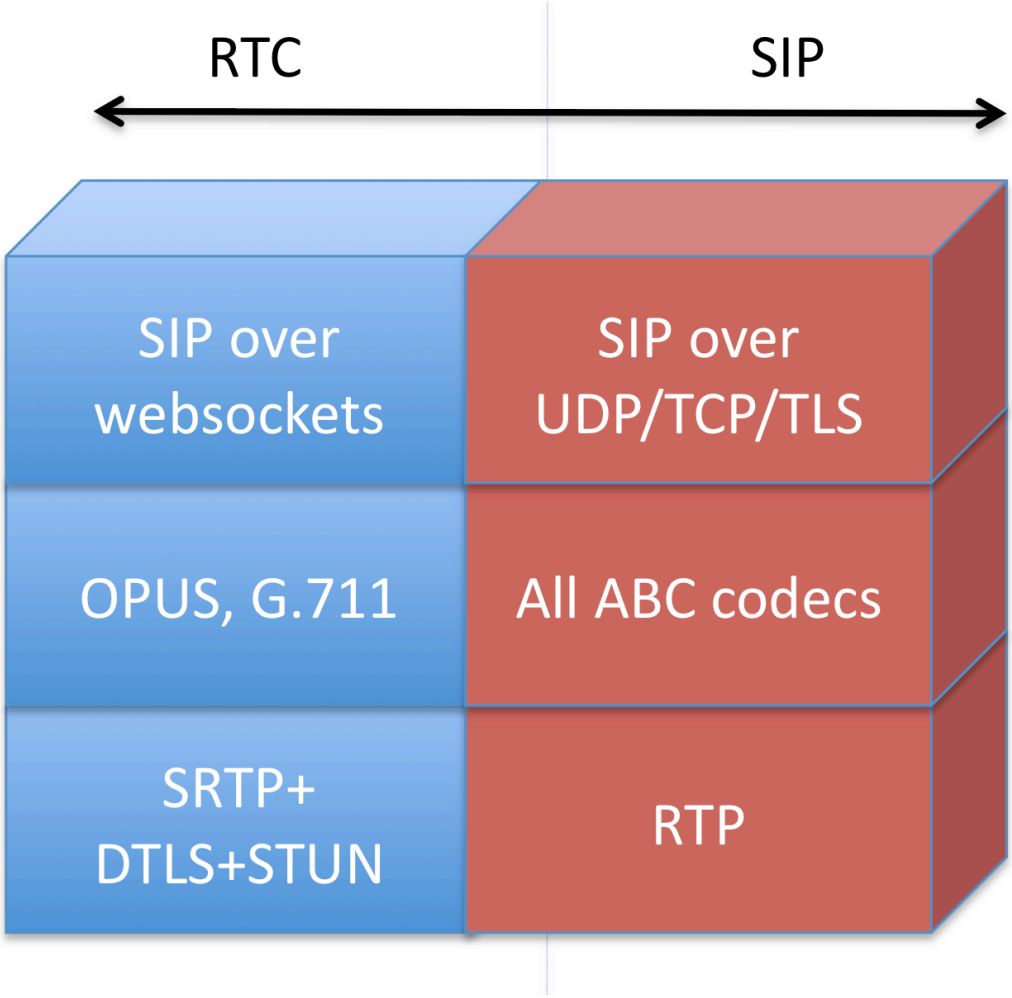

Figure 116: WebRTC Gateway Protocol Stack

Integrating a gateway in a SIP network is fortunately straight-forward. When a SIP-WebRTC gateway is installed and configured to connect to an existing SIP services (PBX, public SIP service), WebRTC clients can immediately reach and be reached from the SIP service. The existing SIP service does not need to be modified at all – it treats WebRTC traffic from behind the gateway as regular SIP traffic.

The rest of this section is split in the following parts: brief introduction to the WebRTC protocols and network architecture is given in Section WebRTC Network Architecture and Protocols. Configuration of the gateway is explained in subsequent sections: WebRTC Network Configuration, WebRTC Credentials Configuration, and WebRTC Rules Configuration. Eventually we provide guidelines for starting an RTC gateway using the Amazon Elastic Cloud services in Section Amazon Elastic Cloud Configuration Cookbook. We offer several methods using either predefined configurations or using manual configuration, and starting a single gateway or a whole failsafe cluster. We also provide recommendations for starting a geographically-dispersed service.

If you plan to start the RTC gateway service in front of an existing SIP service rapidly, best proceed directly to the Section Amazon Elastic Cloud Configuration Cookbook.

WebRTC Network Architecture and Protocols

The WebRTC protocol suite for telephony specifies use of the following protocols:

- G.711 and OPUS (**RFC 6716**) for audio codecs. Opus is a lossy compression, low-delay, codec with constant and variable bitrate ranging from 6kbps to 510 kbps. G.711 is legacy PSTN audio codec at 64 kbps.
- VP8 (**RFC 6386**) for video codec. VP8 is an irrevocably royalty-free codec.
- SRTP (**RFC 3711**) for secure real-time media transmission.
- DTLS (**RFC 4347**) for keying. - SIP over Websockets (**RFC 7118**) as one of possible protocols for signaling. It is slightly aligned SIP using websockets as transport. It is particularly easy to translate to and from legacy SIP.
- ICE (**RFC 5242**), STUN (**RFC 5389**) and TURN (**RFC 6062**) for NAT traversal. STUN is a probing protocol that allows clients to detect how it is reachable over NATs. TURN is a STUN-based protocol that allows a client behind NAT to allocate a publicly reachable IP address from a server and tunnel traffic from and to it. ICE is methodology for finding the best combination of IP addresses to communicate between clients.

At the time of publication of this handbook, Firefox (version 23 and above) Chrome (version 28 and above), Opera (version 20 and above) and Safari (Preview, June 1017) were supporting this protocol stack and have demonstrated mutual interoperability. Several JavaScript applications[4] emerged that implemented signaling using SIP over websockets.

In the simplest scenario, two browsers can use the protocol stack to interconnect with each other. Most of this document is however concerned with the case when one party is using a WebRTC capable browser, and the other party is using a SIP phone or a PSTN phone behind a SIP gateway. This is the most complicated and also critical scenario because it connects the web telephony users to existing population of SIP users. The key component in this scenario is WebRTC-to-SIP gateway which translates signaling and media between the WebRTC and non-WebRTC SIP protocol stacks.

The WebRTC clients use the protocol stack is shown in Figure RTCWeb Protocol Flows. Initially the client registers itself to become reachable for incoming calls. It does so by sending a SIP REGISTER message over websockets. It is that simple.

[4]The JSSIP application is available under MIT License and can be obtained from http://jssip.net.

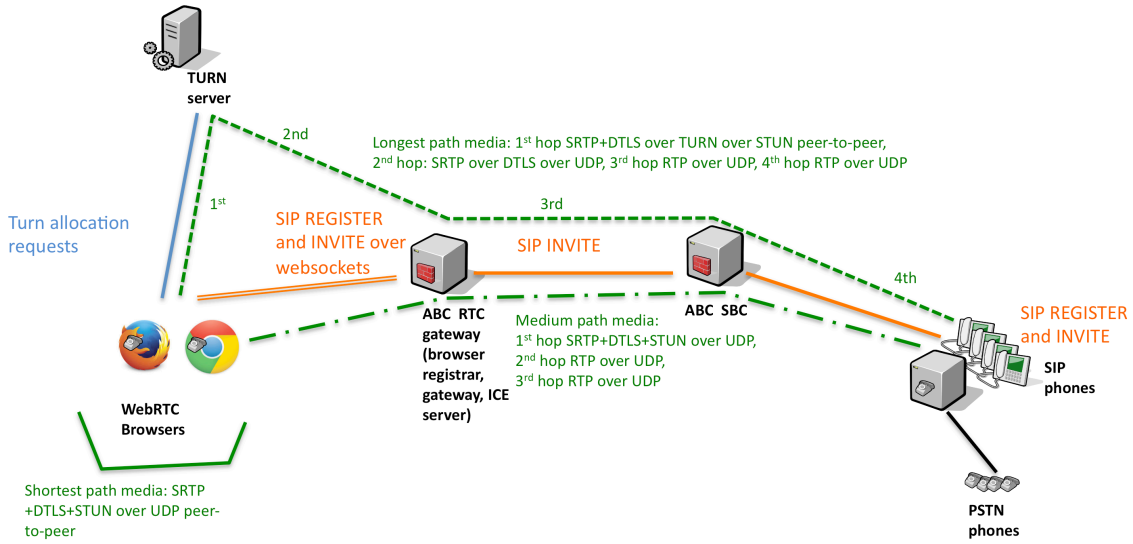

Figure 117: RTCWeb Protocol Flows

When the browser user wants to make a call, it is a more complicated process. The browser starts the ICE process in which it learns IP addresses under which it can be reached. The IP addresses include the WebRTC client's own, its own as seen on the Internet and learned using the STUN protocol, or even a completely different IP address belonging to a TURN media-relay. When the browser initiates SIP signaling, it offers all IP addresses learned in the previous phase. After the called party answers the call, the client probes the IP addresses against the caller to chose the IP address with best IP connectivity. When the "best" IP address is chosen, an encryption session-key is generated using DTLS and media is exchanged using SRTP.

The actual media call-flow can vary depending on how the WebRTC application is configured and the actual call-by-call result of ICE connectivity checks. In a typical scenario deploying FRAFOS gateway, media is sent directly between the WebRTC client and the gateway. This is shown in the Figure RTCWeb Protocol Flows as the green dashed-dotted line. However, the WebRTC application can be also configured to communicate using a TURN server which introduces another hop to the media path. That's the dashed green line in the Figure. It can be for example useful if one wishes to relay media using TCP protocol. It can also occur that both call parties are WebRTC clients on the same subnet and media can flow the shortest-path between them – shown as solid line in the Figure.

However, in scenarios using the gateway the most practical client configuration choice is to limit ICE process to its own IP address. That eliminates gathering the STUN and TURN choices and greatly reduces "post-pickup delay", i.e. the period of time between when the caller answers and media can be actually heard and seen.

WebRTC Network Configuration

This subsection is about what components must be placed in the network and how they must be configured to enable working WebRTC call-flows. First, the following planning questions must be answered:

- do you want to enable NAT/firewall traversal using media over TCP? This may increase the NAT/firewall traversal success rate. If so, the TURN server application must be used on the media interface.
- which client do you want to use? The RTC-capable Web-browser alone includes the RTC engine but still needs an application that uses it. There are various commercial and open-source projects implementing the VoIP functionality, such as JSSIP.
- do you want to integrate the gateway functionality in an SBC or run it on a dedicated server? We suggest to use a dedicated server unless you have a good reason for tight integration. With a dedicated server, it is easy to discriminate WebRTC-to-WebRTC calls from WebRTC-to-RTC, apply different security logic to WebRTC clients, and avoid interference with legacy-SIP configuration.
- under which IP address and port number will be the websocket interface available? To enable websocket communication, you must configure an SBC interface and create a Call Agent linked to the interface. The interface configuration dialog is shown in Figure Websocket Interface Configuration. The most important element is “Interface type” which must be set to “websocket signaling”. The Call Agent configuration is shown in Figure Websocket Call Agent Configuration. By setting its interface to the previously created websocket interface and its IP address to “All” (0.0.0.0/0), it captures every WebRTC clients communicating with the ABC SBC using websockets.

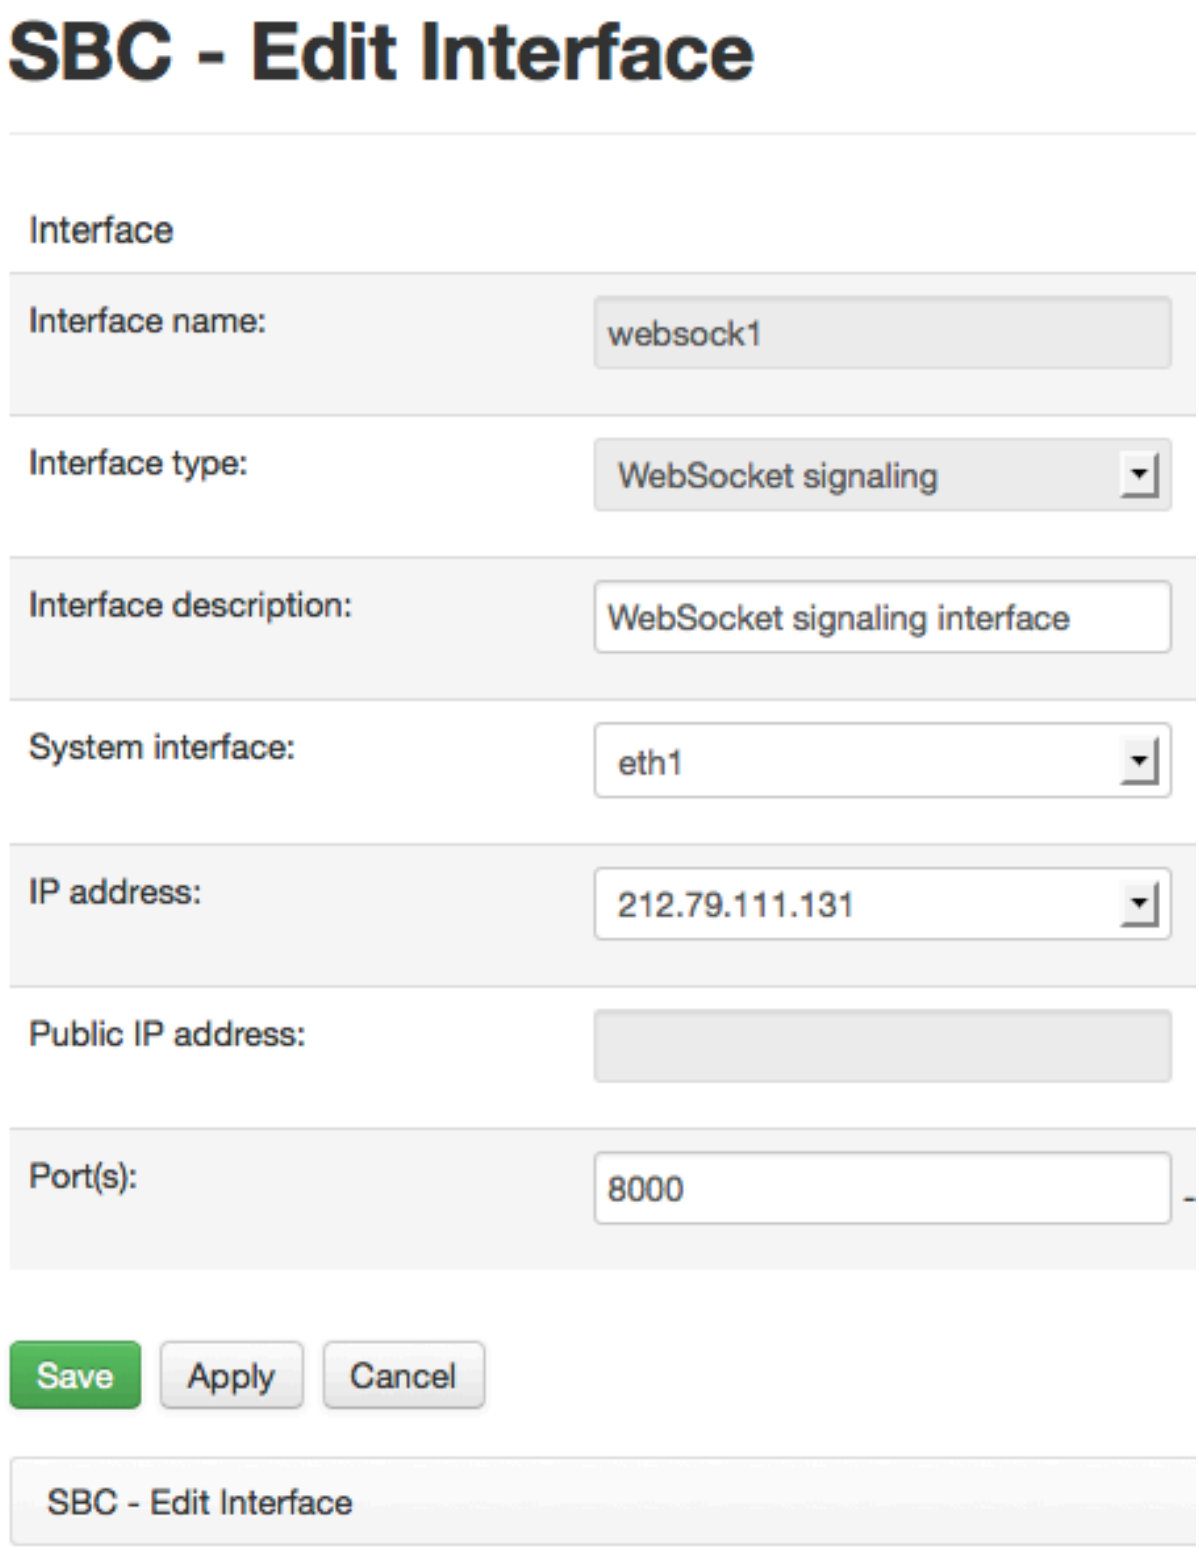

Figure 118: Websocket Interface Configuration

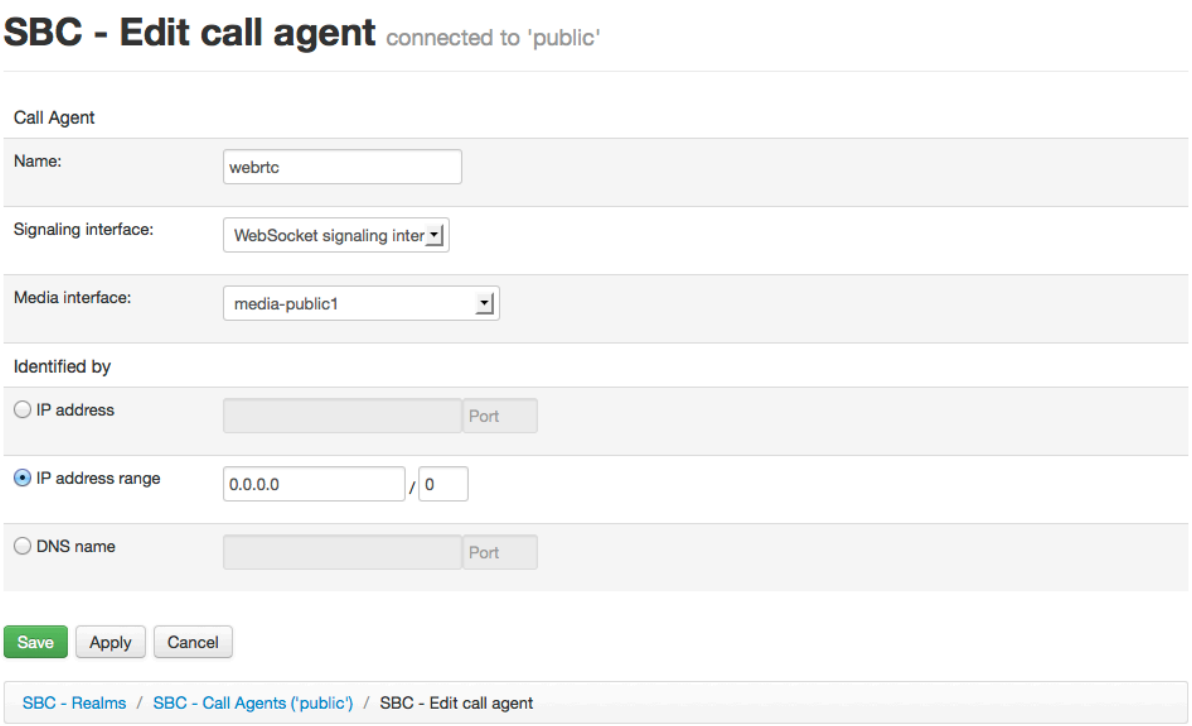

Figure 119: Websocket Call Agent Configuration

WebRTC Credentials Configuration

Confidentiality of calls by encryption is one of the major WebRTC features. Fortunately, it is rather easy to configure. DTLS-SRTP is always enabled in current version of ABC SBC (in previous versions there was possibility to disable it). All other configuration options are optional. Such configuration is shown in Figure SRTP Configuration Page.

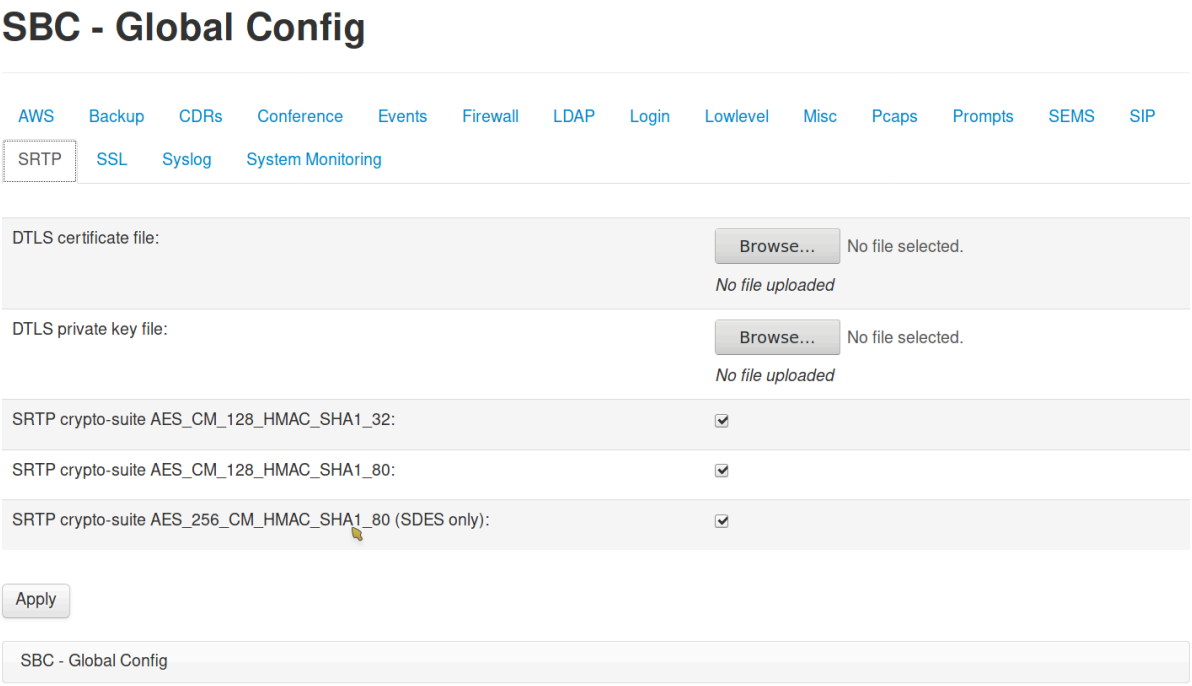

Figure 120: SRTP Configuration Page

When no further options are selected, the ABC SBC creates ad-hoc self-signed credentials. A particular advantage of these is the length of resulting DTLS-SRTP packets will be bellow 1500-bytes packet length which is almost always certain to traverse networks without IP fragmentation.

If you prefer your own certificates, you must upload them using the “DTLS certificate file” and “DTLS private key file” global config options (located under Misc tab).

Note that some credentials may result in too long DTLS-SRTP packets. If they exceed the length of 1500 bytes, they will be most likely fragmented and may result in failure to set up media channel. This is almost certain if there are NATs along the communication path.

WebRTC Rules Configuration

The configuration of the rules for SIP-WebRTC gateway must address both generic SIP processing aspects, which is routing and NAT traversal, and then specific aspects of WebRTC interworking.

In this configuration example we assume topology shown in Figure RTCWeb Protocol Flows, two types of calls: WebRTC-to-RTC and RTC-to-WebRTC, and media flowing through the ABC SBC along the dash-dotted green line.

The SIP routing flow is rather simple in this scenario: every call coming from the WebRTC Call Agent (i.e. over the websocket interface) will be routed to a SIP PBX, and reversely every call coming from the PBX will be routed to RTC browsers using websockets. The routing configuration is shown in Figure SIP-WebRTC Gateway Routing Rules.

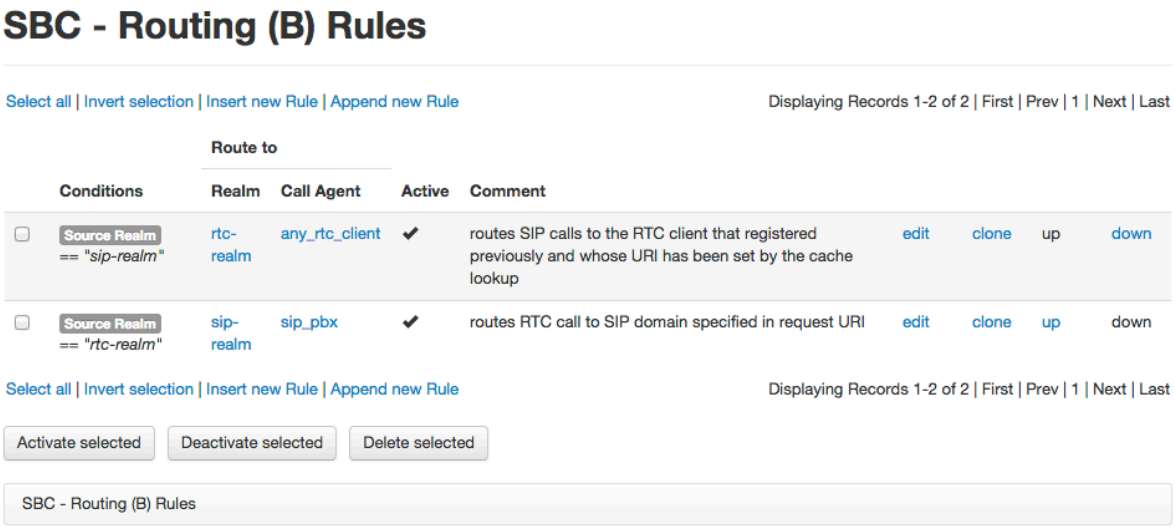

Figure 121: SIP-WebRTC Gateway Routing Rules

The task of A and C rules is to anchor media to itself and to determine when to convert calls from RTC to SIP and vice versa. Therefore we create two realms: one for RTC clients and one for SIP clients. For each of them, we create one Call Agent that captures all traffic from/to any IP address flowing through the websocket and SIP interface respectively. The actions are configured to accommodate the following policies :

Realm	Direction	Policy (Actions)
RTC	A-rules	• enforce frequent re-REGISTERs to keep persistent TCP connections for websockets alive (REGISTER throttling) • cache registrations to forward SIP calls for RTC clients properly (Enable REGISTER caching) • fix NAT bindings (Enable Dialog handling) • anchor media, offer ICE and RTC Feedback to RTC clients (Enable RTP Anchoring)
RTC	C-rules	• anchor media (Enable RTP Anchoring) • enforce SRTP using DTLS keying (Force RTP/SRTP)
SIP	A-rules	• lookup registered RTC user, decline the call if offline (Reply to request with reason and code, Retarget R-URI from cache (alias))) • anchor media, don't offer ICE to SIP callers (Enable RTP Anchoring)
SIP	C-rules	• anchor media (Enable RTP Anchoring) • enforce plain RTP on the way to the SIP Call Agent Force RTP/SRTP

We have met most of the rules in previous sections: driving re-registrations high to keep transport-layer connections alive, caching registrations, fix bindings and anchor media. Now we need to include the specifics of SIP and RTC interworking. SIP calls towards RTC clients must appear RTC-capable, i.e. they must offer SRTP encryption, ICE connectivity checks and RTC feedback. Reversely, the RTC calls to SIP must be transformed to plain RTC.

The "Force RTP/SRTP" action determines if plain RTP or SRTP is used for a call. When this action is placed in C-rules, it converts media for the called party into the enforced protocol. When SRTP is chosen, one must set an additional option: the keying protocol. Only DTLS makes sense for RTC. In our example we convert all media traffic towards SIP devices by placing "Force RTP" in SIP realm's C-rules. Analogically we convert all media traffic towards RTC clients by placing "Force SRTP" in RTC realm's C-rules. The "Force SRTP" action is using "DTLS" as the keying option because that's the keying protocol standardized for use with RTC.

One could also use the "Force RTP/SRTP" action in A-rules: here however it only determines if the caller's SDP offer complies to the enforced preference and rejects the call otherwise. We are not using this kind of admission policy in our example.

The other options specific to the RTC interworking use-case are specific to how we anchor media. We need to make sure that RTC clients relying on ICE will receive proper STUN answers for their connectivity checks towards the built-in media relay and also RTC feedback. Therefore, we turn the options "offer ICE" and "offer RTCP feedback" on in the media anchoring action for both RTC A-rules and C-rules. The A-rules make sure that incoming RTC call offers obtain ICE and RTC/F capable answers, the C-rules ensure that SDP offers towards the RTC clients will be also ICE and RTC/F capable.

The resulting configuration is shown in Figures Configuration of RTCWeb Rules for RTC Realm and Configuration of RTCWeb Rules for SIP Realm for the RTC and SIP realm respectively.

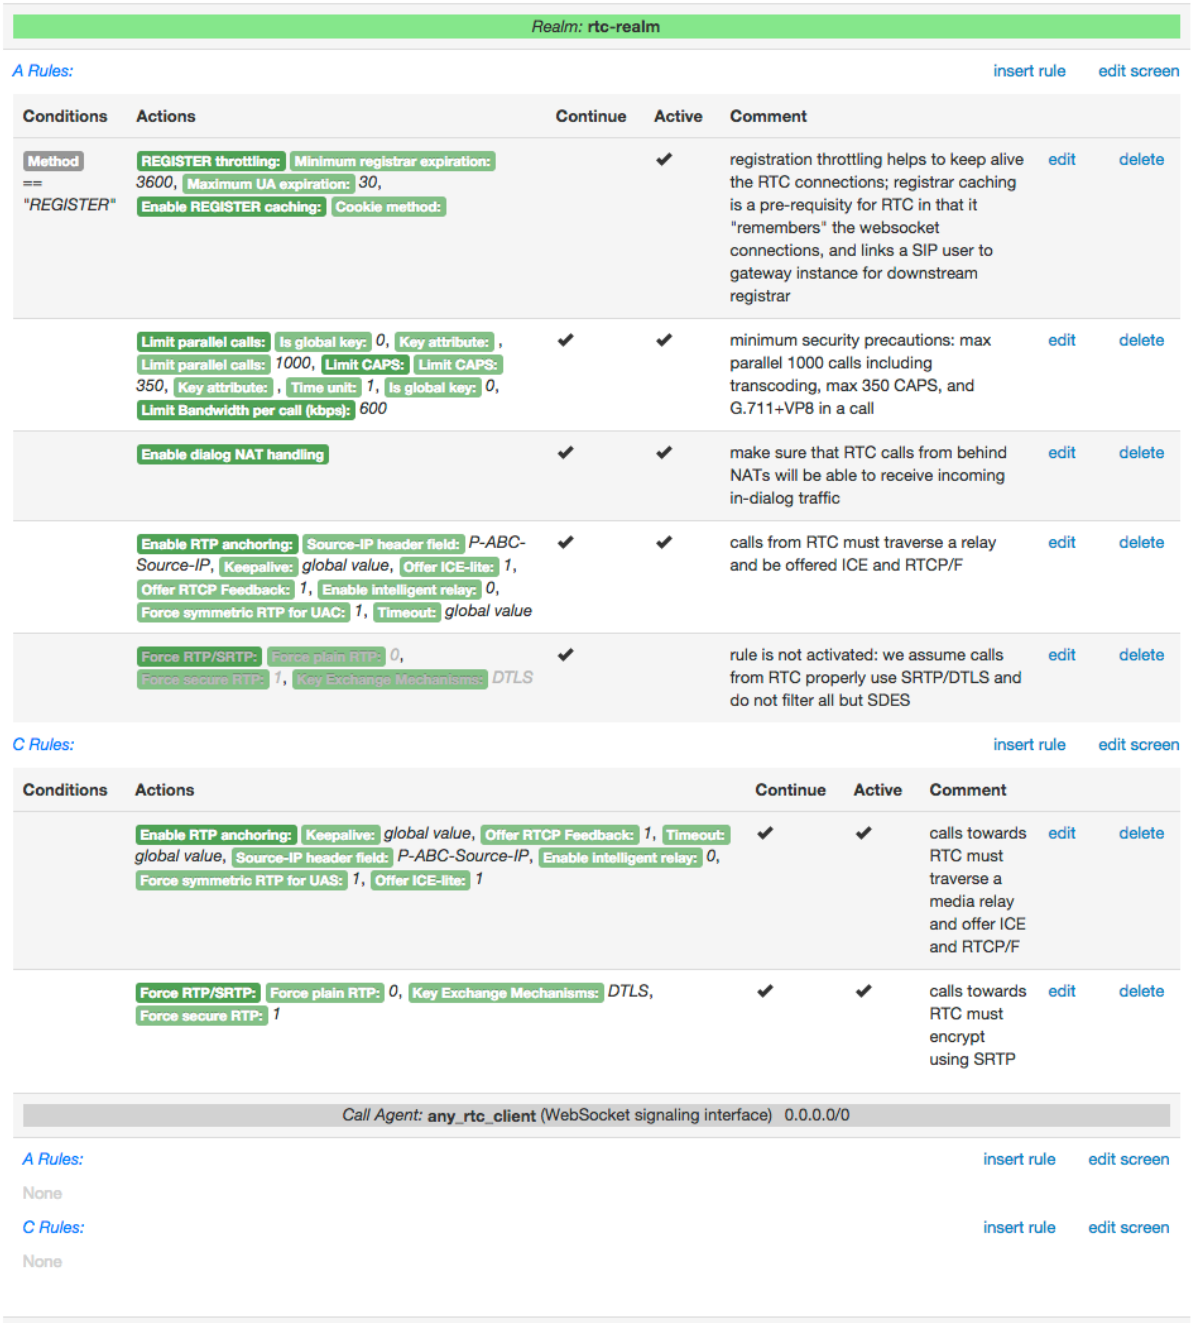

Figure 122: Configuration of RTCWeb Rules for RTC Realm

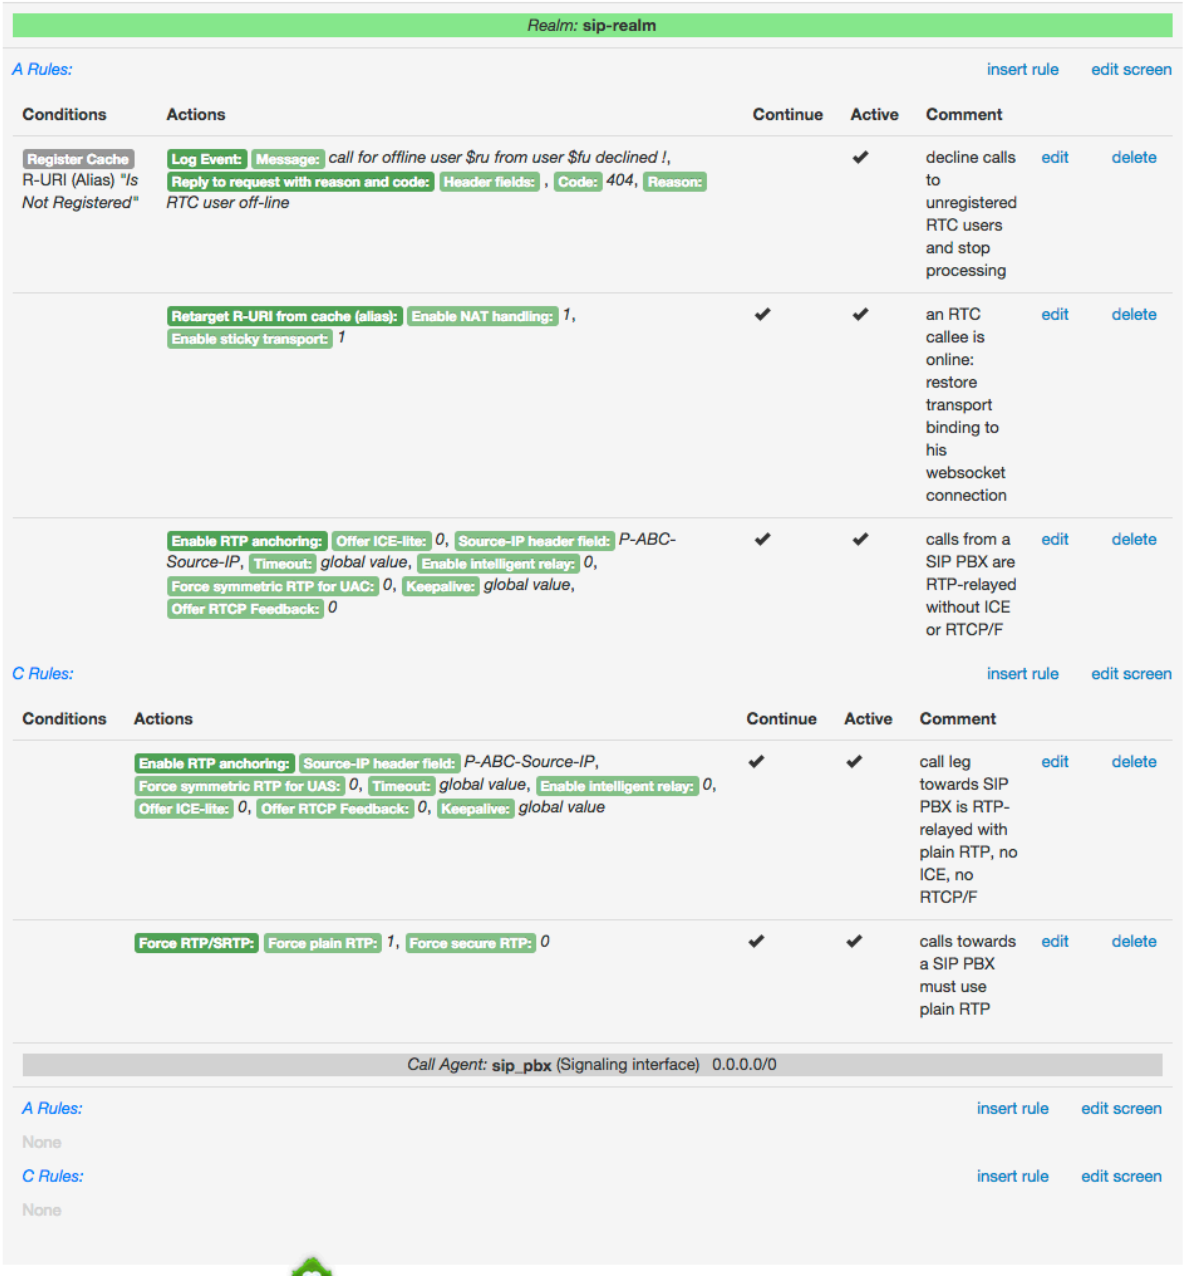

Figure 123: Configuration of RTCWeb Rules for SIP Realm

Note that this configuration works even if two WebRTC clients connect to each other through the gateway. However the WebRTC-to-RTC conversion and forwarding to the SBC still takes place resulting in an WebRTC-to-RTC-to-WebRTC loop,as shown in Figure The WebRTC-to-WebRTC Loopback.

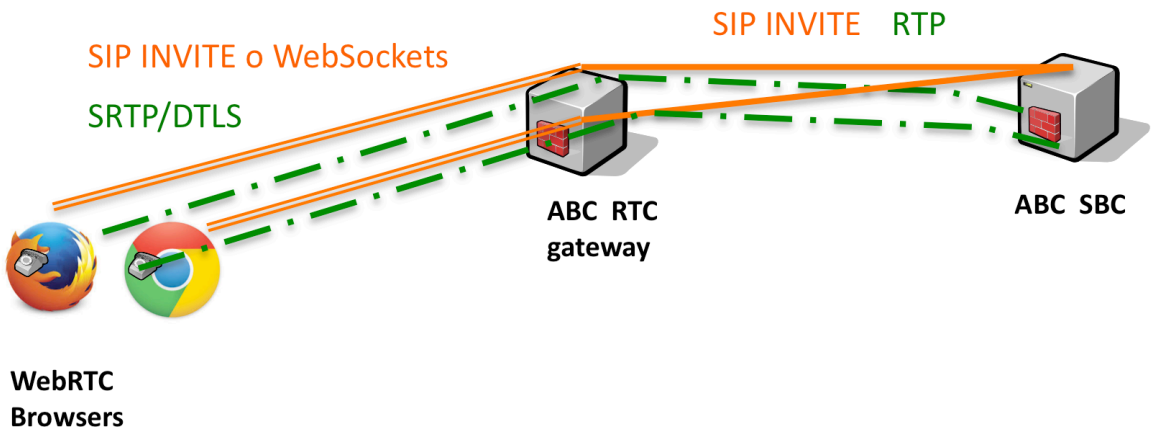

Figure 124: The WebRTC-to-WebRTC Loopback

Optionally it may be useful to manage codec negotiation. For example, one could blacklist G.711 in favor of OPUS, if there are SIP clients that can speak the codec. Or video could be stripped off, if there is no support for royalty-free VP8 codec. Note though that if codecs are stripped too aggressively, a SIP user agent may fail to interoperate and return a 488 in UAS or an immediate BYE in UAC role.

WebRTC Interoperability Recommendations

The WebRTC standard and implementations are relatively new and as result degree of interworking largely depends on network configuration and used client. Unfortunately interoperability is still changing with every new version of WebRTC stack and the clients built upon it.

Network complications typically arise when there is a "middlebox", an Application Layer Gateway (ALG) or an HTTP proxy in the path. This sort of network equipment manipulates HTTP traffic in a way that may impair interoperability. If the middlebox cannot handle the websocket extension of the HTTP protocol, signaling connection will fail. Therefore the default transport protocol for SIPoWebsockets is TLS.

WebRTC application complications typically arise when the application has imperfect support for the SIP protocol running on top of websockets, and/or changes its behavior with a new software version. ***We urge our customers to test extensively the client application before initial deployment of a WebRTC service AND during an update to a newer version.***

The most "fluid" interoperability difficulty is continuous changes to the WebRTC protocol stacks hidden insider the browsers. Almost with every browser release, some minor changes appear that impair interoperability. Until the environment becomes more stable, typical reaction is reverse analysis of the new interop behavior and using ABC SBC mediation features to address it. For example, Chrome browsers Version 39.0 and higher are known not to handle "early media" correctly. The ABC SBC configuration allows to mediate "183 early media" into regular "180 ringing" as shown in Figure WebRTC Mediation Example.

Header: X-Browser begins with *"chrome"*	Add Header: *xx: 183-supressed*, Drop SDP from 1xx replies: *183*, Translate reply code: *183 -> 180 ringing*	✔	✔	interop hack: suppress early media to chrome -- new versions cause "Bad Media"

Figure 125: WebRTC Mediation Example

In summary, while the industry is converging to a solid level of interoperability, thorough effort during initial and regression tests is highly recommended.

4.10.16 Amazon Elastic Cloud Configuration Cookbook

An easy way to start and run an RTC gateway is "off a cloud", i.e. on a hosted platform, without purchasing and operating own physical infrastructure: computers, racks, disks and IP connectivity. A single click is enough to start a service, of course as long as you keep paying for the cloud services. Whereas there are many "cloud platforms", we focus on running the RTC gateway on the Amazon Web Services (AWS) platform in this section.

The AWS platform is a mature system that allows, among many other useful things, to start and run pre-built virtual machines, load-balance traffic among them, monitor their health and scale the infrastructure up and down to be on par with user load. Applications, RTC-to-SIP gateway in our case, come pre-installed ready-to-start in form of a virtual machine image, called AMI (Amazon Machine Image).

The FRAFOS RTC gateway installation is pre-configured to address a simple yet useful scenario: add RTC connectivity to a running SIP PBX service. Once started, the RTC gateway passes RTC registrations and calls coming from RTC clients down to the PBX(s). Reversely, the gateway routes calls coming from the SIP PBX(s) to the previously registered RTC browsers. No further configuration is needed.

The following subsections describe how to start the RTC-to-SIP gateway service off the amazon platform. We offer you several ways to do the same: the easiest is launching a cluster using Cloud Formation template. This way you create a load-balanced scalable infrastructure by pressing a button without any further knowledge of how the components must be configured. If you want to understand in more detail how the gateway works, you can launch a single-instance service and/or configure it in detail step by step.

We suggest you explore our demo site https://go.frafos.com. It includes additional information about use of AWS and WebRTC technologies, including live services and ready-made demo AMIs and Cloud Formation Templates. These can be launched by a single click without any need for further configuration. Note that the demo versions have a 90-seconds limitation to maximum call duration.

Before you Start: Prerequisites and Important Warnings

Before you start, you shall have the following:

- Amazon Web Services (AWS) account. Note that the accounts come with several service plans charged at different levels, and credit card number and a telephone must be ready to verify identity and payment. Go to http://aws.amazon.com to sign up.
- AWS Elastic Cluster SSH keypair. This is important to be able to administer the virtual machines remotely. If you haven't created or uploaded one, do so under "EC2▸Keypairs". If you want to start the services in multiple regions, make sure that you have a keypair for every region before you start.
- Amazon Machine Image (AMI) with the RTC-2-SIP gateway from FRAFOS. You will find the right one for your geographic region on our experimental web page, https://go.frafos.com/.
- RTC-enabled browser for testing. Latest version of Chrome has been tested by FRAFOS to play well, yet there are other implementations as well.
- Optional: Publicly available SIP service and a SIP account. You need to have a SIP URI and password with a SIP service to be able to make calls through the RTC-to-SIP gateway. Otherwise you can only make anonymous calls.
- Optional: a DNS name under which your RTC-to-SIP gateway will be reachable

To begin visit our experimental web page https://go.frafos.com/. The web page contains predefined links to available AMIs that allow you to launch quickly.

Info

IMPORTANT: USE OF AMAZON WEB SERVICES WILL INCURE ADDITIONAL COST. ALL DATA CREATED AND STORED ON AN INSTANCE SUCH AS PROVISIONED TABLES, ABC RULES, CONFIGURATION PARAMETERS, LOG FILES AND MORE REMAINS ON THE INSTANCE AND WILL BE LOST UPON INSTANCE TERMINATION.

Quick Start Using Cloud Formation

The ultimately fastest way to launch your service is using amazon's Cloud Formation. The Cloud Formation amazon.com service is used to quickly start a whole network based on a description included in a template. The template includes information about virtual instances, how to scale them up and down, how to spread the load across them using a load-balancer, and what firewall policy to use to filter IP traffic: quite some work if administrator was setting all of this up manually.

FRAFOS has created a starter template to be used to start a fail-safe cluster of one-to-four gateways behind a load-balancer. The template is available on our site, https://go.frafos.com.

During the process you will be prompted for very few parameters. Their scope can change as we keep developing the template and for most cases they are best served by leaving them to their default values. The only required parameter you must set is the name of your SSH key. Once you start the cloud formation process, it takes several minutes until it completes. After the stack is launched, you will have one load-balancer and one to four gateways running behind it. A URI shown upon completion of the cloud formation process will allow users to download a demo JavaScript application and start using the service. Sometimes you may need to be patient for a couple of minutes until the service is really "warmed up".

When trying to place your first phone call, you may for example try to call sip:music@frafos.net. When opening the web-page, allow the browser to accept self-signed certificate and use your microphone and camera.

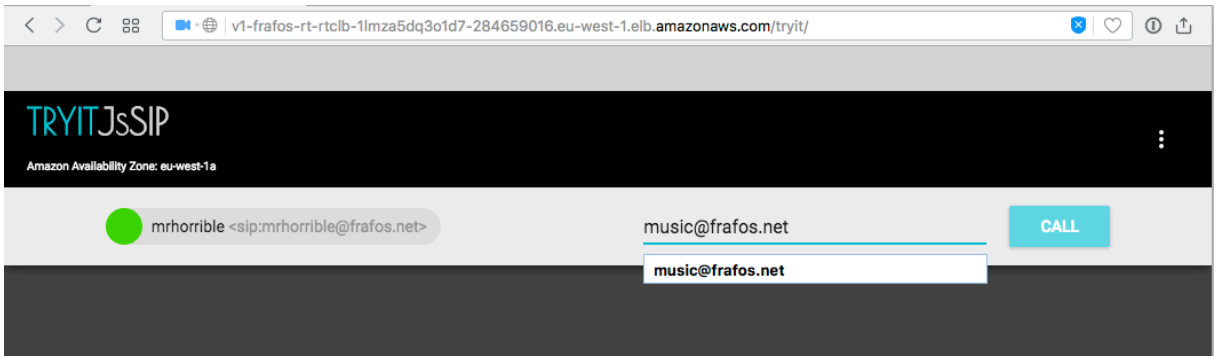

Figure 126: Screenshot: First Browser Call to music@frafos.net

Also you can try out the built-in audio conferencing bridge by dialing an 8-digit number prefixed with *. Anyone calling the same address will appear in the same conferencing room.

As the next steps, you can follow the links that show in the Cloud Formation Output window: a WebRTC web telephony application and the ABC Monitor (use sbcadmin username and default password). You can also administer the actual instances by going to their web address "https://IP/", username "sbcadmin" and password equal to instance ID. For example, you can review rules that remove video streams between WebRTC and legacy SIP to allow at least audio where video signaling often fails, or look at the dialing plan for the on-board conferencing.

Quick Start: Launch Single Instance

If beginning with a cluster may appear too heavy start, one can also start a single RTC gateway instance instead. This can be done also from our site, https://go.frafos.com.

During the process you will be prompted for instance type, detail, used storage and security group. Choose an instance type with at least 2GB RAM and leave everything else except the Security Group to default. The security group must be set to permit the following flows from 0.0.0.0/0:

- TCP/5060-5069 — SIP service
- UDP/5060-5069 — SIP service
- TCP/443 - web user interface
- TCP/22 — secure shell
- UDP/10000-11000 RTP media

Eventually chose an existing or create a new key-pair and store the private key securely.

Once the virtual machine is up and running, you can access administrative interface using the https://PUBLIC_IP/. The administrative username is "sbcadmin", the password is the ID of your amazon instance. You can also access remote shell if you login using the private part of the AWS SSH key:

```
$ ssh -i .ssh/frafos-aws-keypair.pem -l ec2-user 54.171.123.109
```

If you would like to use additional AWS features that the instance supports, CloudWatch and System Manager, you must enable an instance role that permits these. An easiest way to do so is to create an AWS/EC2 role with predefined permissions "EC2 Role for Simple Systems Manager" (AmazonEC2RoleforSSM) and attach it to the instance.

Updating License

On Amazon Cloud there is an easy way to install centrally a license file that is then used by all newly started ABC SBC instances. This is practical when you upgrade to a feature-richer license and do not want to configure the license individually in every new instance. The license is then used by both instances that are newly started individually as well as via Cloud Formation and AutoScaling. You only need to make sure the license file matches the AMIs you are using.

After obtaining the license file from FRAFOS support, all you need to do is to enable instance's access to Systems Manager (see http://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-access.html#sysman-configuring-access-role) and put the license in a parameter with a well-known name in the Parameter Store. The Parameter Store is located under the EC Dashboard under "System Manager Shared Resources ▸ Parameter Store". The parameter name must be "/abcsbc/license" as shown in the screenshot bellow.

Note that setting this parameter does not affect running instances, only applies to the AWS Region for which you provisioned it, and must include a license specific to the AMIs you are using.

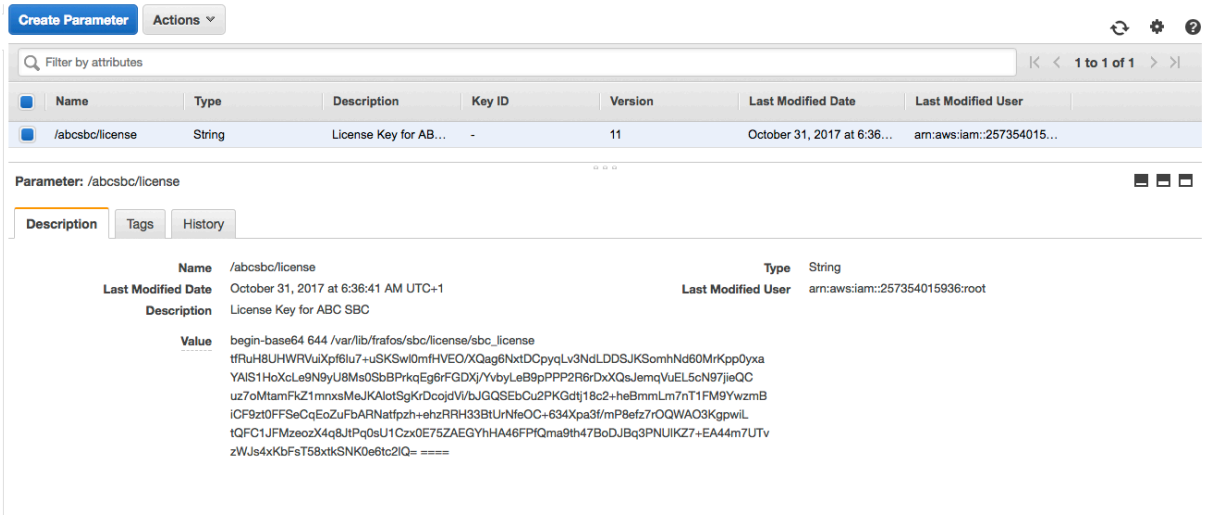

Figure 127: Screenshot: Setting License in Amazon Parameter Store

Introducing Geographic Dispersion

Introducing geographic redundancy and dispersion may be useful to become resilient against regional disasters and/or decrease VoIP latency. Latency may have major impact on quality of service. For example if an American user accesses a European RTC gateway to reach an American SIP PBX, media will travel across the Atlantic back and forth, resulting in noticeable latency and QoS degradation.

Fortunately there is an easy-to-manage way with AWS to build up geographic redundancy for both individual instances and whole clusters. All that needs to be done is creation of the instances or whole stacks as described in previous subsections multiple times in different regions, and linking their addresses to a single latency-routed DNS name. That is a particular feature of Amazon Route53 DNS service, that returns the lowest-latency IP address associated with a DNS name.

We experimented with this amazon feature and confirmed significant latency savings. In our example, we created two instances, one located in Ireland, the other in California. We create CNAME records “”eu.areteasea.com” and “us.areteasea.com” for them. Eventually we created the latency-routed global DNS name entries “world” for both regions, as shown in Figure Screenshot: Creating DNS Latency-based Routing Records.

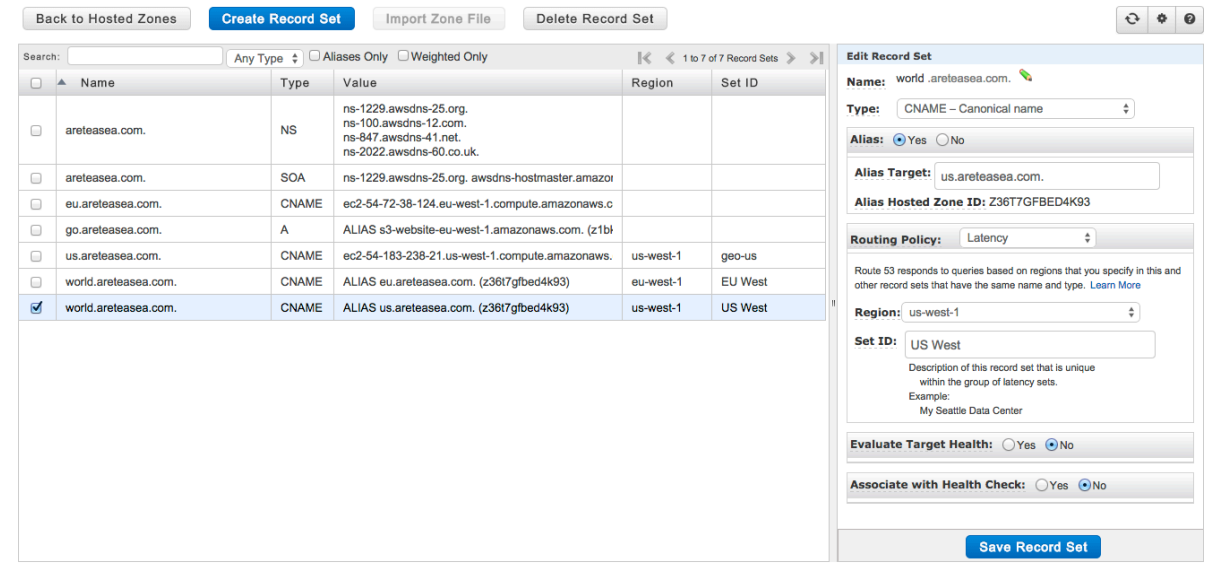

Figure 128: Screenshot: Creating DNS Latency-based Routing Records

Clients trying to open up a connection to “world.areteasea.com” resolve this DNS name to different IP addresses depending on where they ask from.

One can easily verify the outcome by using services like Cloud Monitor. (http://cloudmonitor.ca.com) Results shown in Figure Latency Measurements for Multiple Sites Served by Route-53 Latency-Routing prove that proximity makes a difference. Clients in geographic proximity of the two sites feature minimum latency bellow 50ms: US from California to Illinois show 30 to 50ms, Western Europe shows 24-37 ms, Ireland 8 ms. Clients located out of served continents have significantly higher latency, starting with 180ms for Australia, slightly above 200ms for Argentina and Egypt, and peaking with 329 ms in China – values that make VoIP quality poor.

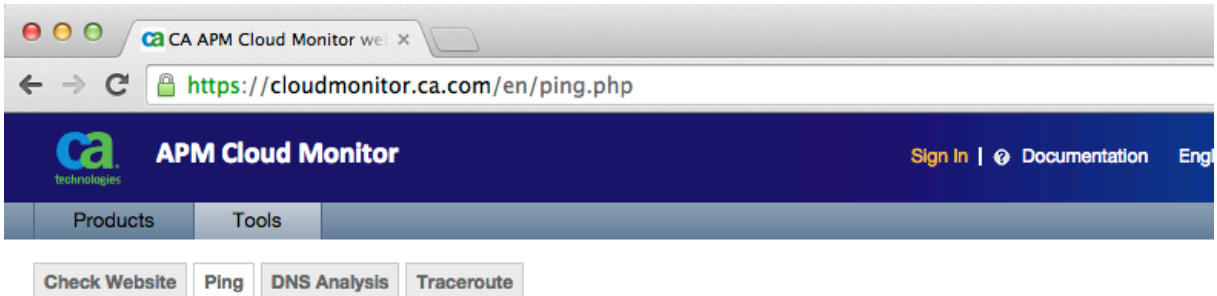

Ping to: world.areteasea.com

Checkpoint	Result	min. rtt	avg. rtt	max. rtt	IP
Australia - Perth (auper01)	Okay	210.809	216.162	233.296	54.183.238.21
Australia - Brisbane (aubne01)	Okay	182.918	195.687	218.158	54.183.238.21
Argentina - Buenos Aires (arbue01)	Okay	203.796	214.552	227.346	54.183.238.21
U.S.A. - Atlanta (usatl02)	Okay	62.042	66.020	70.947	54.183.238.21
Australia - Sydney (ausyd02)	Okay	157.751	161.100	173.726	54.183.238.21
Brazil - Sao Paulo (brsao03)	Okay	230.865	242.161	262.523	54.72.38.124
Brazil - Porto Alegre (brpoa01)	Okay	258.186	262.903	267.292	54.72.38.124
Brazil - Rio de Janeiro (brrio01)	Not available				
Canada - Vancouver (cavan02)	Okay	23.204	25.231	36.537	54.183.238.21
Belgium - Antwerp (beanr02)	Okay	23.952	25.235	33.214	54.72.38.124
Bulgaria - Sofia (bgsof01)	Okay	52.589	52.811	53.151	54.72.38.124
India - Bangalore (inblr01)	Okay	155.741	159.433	163.390	54.72.38.124
U.S.A. - Boulder (uswbu01)	Okay	30.142	35.903	54.235	54.183.238.21
U.S.A. - Boston (usbos01)	Okay	76.957	83.841	101.428	54.183.238.21
Canada - Toronto (cator01)	Okay	80.689	85.491	104.307	54.183.238.21
China - Hangzhou (cnhgh01)	Okay	319.052	327.840	335.757	54.72.38.124
U.S.A. - Chicago (uschi04)	Okay	49.816	50.264	51.917	54.183.238.21
U.S.A. - Santa Clara (usscz01)	Okay	49.516	52.902	69.543	54.183.238.21
Switzerland - Zurich (chzrh01)	Okay	34.025	34.184	34.320	54.72.38.124
U.S.A. - Charlotte (usclt01)	Okay	66.776	73.957	80.898	54.183.238.21
China - Shanghai (cnsha02)	Okay	329.831	335.842	344.591	54.72.38.124
Costa Rica - San Jose (crsjo01)	Okay	123.763	125.345	128.386	54.183.238.21
Canada - Calgary (cacal01)	Okay	37.088	43.842	65.961	54.183.238.21
Czech Republic - Prague (czprg01)	Okay	31.476	31.718	31.948	54.72.38.124
Germany - Munich (demuc01)	Okay	36.318	37.166	38.934	54.72.38.124
Germany - Cologne (decgn01)	Okay	22.085	22.321	22.468	54.72.38.124
Germany - Frankfurt (defra02)	Okay	26.457	26.587	26.831	54.72.38.124
Germany - Berlin (deber01)	Okay	33.570	34.021	36.331	54.72.38.124
Ireland - Dublin (iedub02)	Okay	7.637	19.883	33.541	54.72.38.124
Egypt - Cairo (egcai01)	Okay	212.333	224.375	241.187	54.72.38.124

Figure 129: Latency Measurements for Multiple Sites Served by Route-53 Latency-Routing

After we had disabled the European site, all sites began to be served by the Californian server and we observed increase in minimum latency of European clients up to 160-180 ms, i.e. by about 130 ms! Therefore we recommend anyone serving global user population to consider establishing presence in multiple amazon's availability zones.

Monitoring the Autoscaling Cluster Using CloudWatch

Once the cluster is up and running, it may be worthwhile to experiment with its autoscaling behavior and monitor how the cluster reacts to varying load. There are various ways how you can observe the status of the cluster using Amazon's CloudWatch facility. The CloudWatch facility collects data from all related instances and load-balancers, aggregates it for the whole autoscaling groups and triggers alarms if some critical values are exceeded. The collected data, how it is aggregated and when it triggers autoscaling alarms is part of the CloudFormation template definition, so if you started the cluster using the template it is already in place. By default, the autoscaling alarms add a new instance when it the average CPU load in the cluster exceeds 80% for several minutes, and remove an instance if it drops bellow 60%.

The interesting data you can observe include the event-by-event history under "EC2 ▸ Autoscaling Group ▸ Scaling History", details of autoscaling alarms in the CloudWatch Console, and graphs showing the cluster changes along a timeline are also found in the CloudWatch console. The rest of this section shows typical autoscaling situations and how you can inspect them using these monitoring facilities.

The first Figure Screenshot: Scaling History shows example of scaling history. We interpret it in time order from bottom up. Initially when the Autoscaling process started it launched the first instance at 12:34. Because we kept the machine busy, some seven minutes later at 12:40 the Autoscaling process chose to reinforce the cluster. It increased the desired capacity to 2 and launched a new instance. Then we started reboot of an instance to simulate a failure. The ELB checks detected the unresponsive instance at 12:48, terminated it, and started a new one. Eventually we relaxed the load, the low-CPU alarm was triggered in response to which the Autoscaling process reduced cluster size back to one.

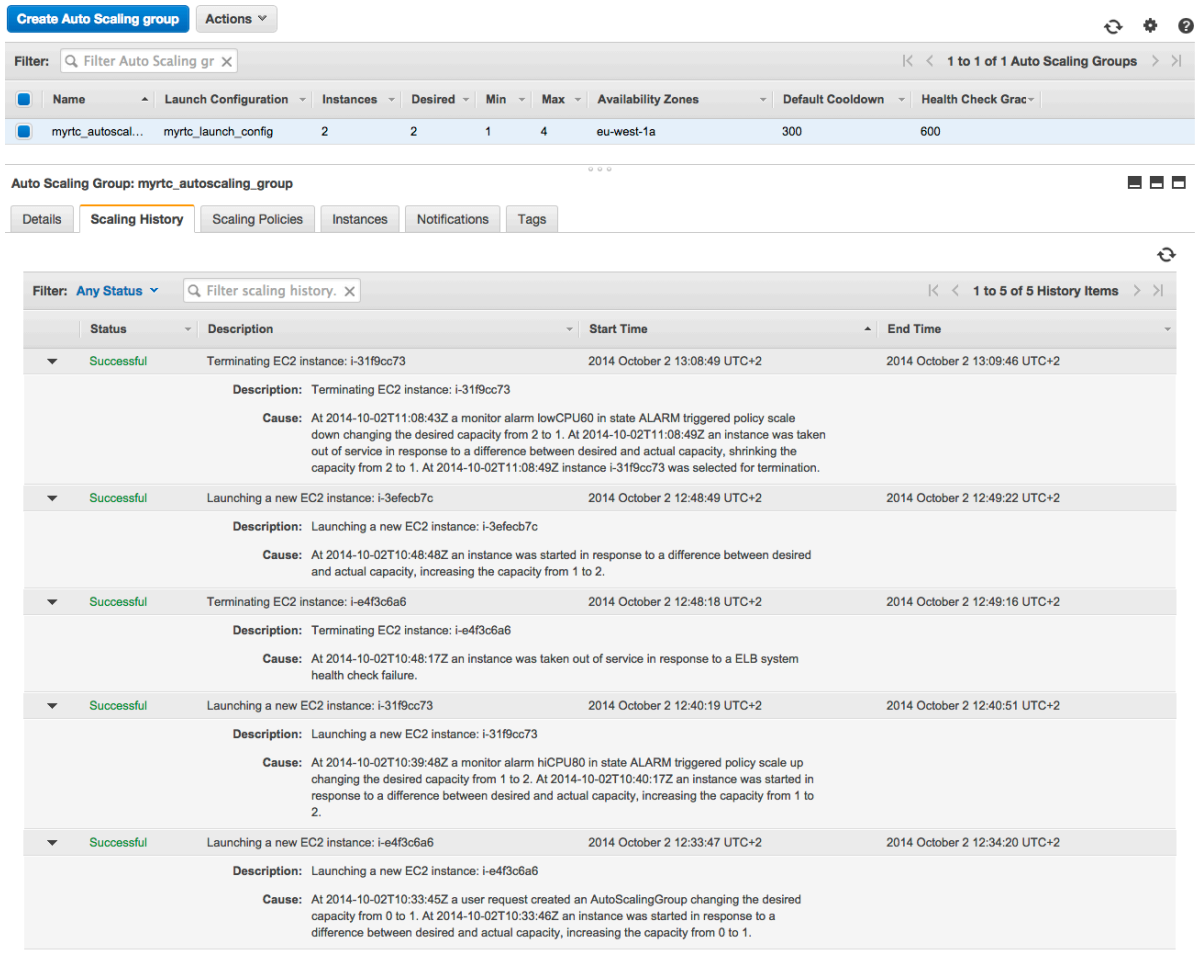

Figure 130: Screenshot: Scaling History

The next Figure CloudWatch Screenshot: Low Load Alarm shows details of a CloudWatch autoscaling alarm. It displays a situation when cluster began to be idle after a period of congestion and an alarm is raised to scale the cluster down. The autoscaling process will remove an instance in response to this alarm.

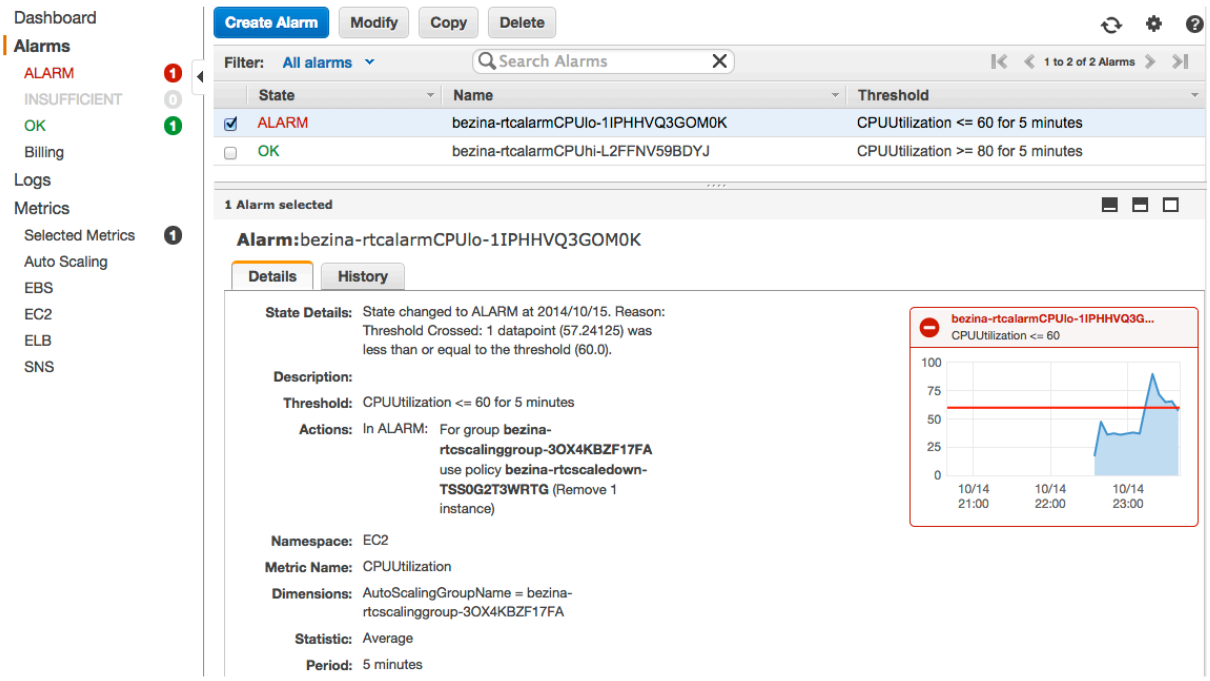

Figure 131: CloudWatch Screenshot: Low Load Alarm

Development can be show in different time-scales using CloudWatch graphs. Figure CloudWatch Graphs: Correlation of Cluster CPU Load and Autoscaling shows how detection of overload and idle conditions affect cluster size along the time axis. There are three lines in the graph: the orange line shows average CPU load in the cluster. The autoscaling assessment of needed capacity is shown using the blue-line and the actual number of available instances is shown using green line. The CPU-load-line leads the changes: it must remain for a period of time above the threshold of 80% until the auto-scaling process determines to increase the target capacity. It then takes some time again until the capacity is ready: a new instance must be launched, detected as ready and included in the load-balancer's distribution list. Therefore the green line legs behind the blue-line, and the blue-line always legs behind the orange-line.

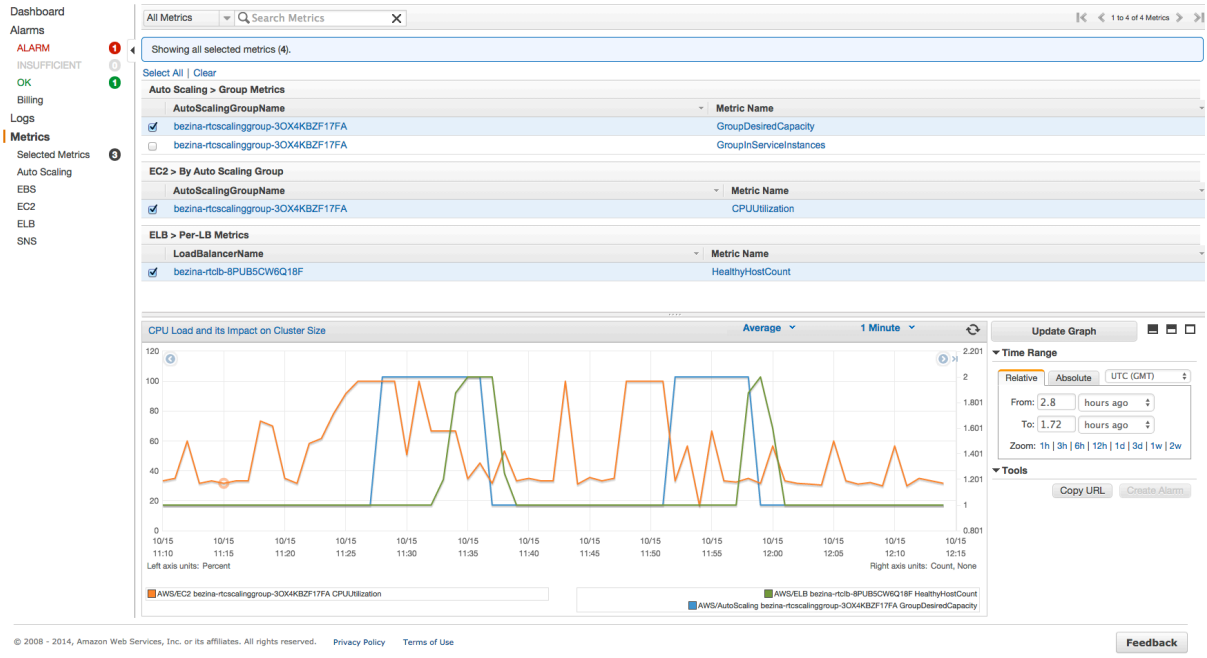

Figure 132: CloudWatch Graphs: Correlation of Cluster CPU Load and Autoscaling

Performance Recommendations

In virtualized cloud environments, performance can vary significantly due to the "sharing" nature of these environments. It is therefore advisable to choose properly dimensioned computing instances. Amazon offers several types of "instance types" that very in various performance aspects. The instance type vary by region and change over time. Current offering is available on the amazon web page http://aws.amazon.com/ec2/instance-types/.

For minimum density trials, the 2GB RAM T2.small instance type is sufficient. This instance allows very little CPU capacity in short bursts. However if the allowed burst is exceeded, the virtual machine will slow down to an extent that it stalls. Experiments with on-board conferencing have shown that a single conference with more than three participants already brings the machine to stalling.

For predictable performance, you will need a Fixed Performance Instance (FPI) type.

In the mainstream case, when media anchoring is enabled and there is neither transcoding nor encryption taking place, the critical parameter is the number of parallel calls (PC). Our lab measurements in this configuration have shown the following capacity for the following instance types available on the Amazon Marketplace: (the instance parameters are from https://aws.amazon.com/ec2/instance-types/)

Instance Type	PC	vCPU	Mem (GiB)	Networking Performance	Notes
m3.medium	180	1	3.75	Moderate	CPU-constrained
m4.large	372	2	8	Moderate	network-constrained

In the less usual case that SIP is processed without RTP, number of call attempts becomes the critical parameters. This can be the case when the SBC is used as a signaling-only load-balancer. Then choosing a CPU-strong instance type makes sense. Our tests have shown that the m3.xlarge instance type can deliver 40 Calls Per Second signaling rate, c3.8xlarge delivers about 500 CPS.

Note that OS-reported CPU-load values may be misleading on virtualized machines. CPU time may be "stolen" by virtualization hypervisor and system tools may or may not accurately report the status. The more accurate method to determine actual utilization of the virtual instances is CloudWatch. We recommend that CloudWatch-observed CPU utilization shall not exceed 80% – if deployed in an Elastic Cluster, this should be the threshold value triggering autoscaling cluster growth.

4.10.17 Template parameters

When configuring multiple SBC nodes, it might be needed to use different value in a rule or config variable for certain nodes or config groups. Template parameters can be referenced in various configuration parameters and thus allow for building different values for specific nodes or complete config groups. Template parameters are referenced by their names enclosed by a percent characters ("%"). For example: **%my_param%**. Each instance of a template parameter is replaced by its value when the configuration for a specific node is generated.

Definition of Template Parameter

There two ways to define new template parameter:

Define parameter directly in input field

Just enter its name enclosed by "%" characters into any input field allowing template parameters and click on the 'Create' link.

If a template parameter needs to be used in a drop-down or checkbox field in the GUI, just simply click on the pencil icon next to the input field. It allows for entering free form text into the input field.

As of now the template parameters are supported in Rules, Global Config and in interfaces/node TLS profile assignment.

Define parameters on the "Cluster config parameters" screen

The screen located at "Config->Define cluster config parameters" lists all the defined parameters. The screen allows for creating new parameters as well.

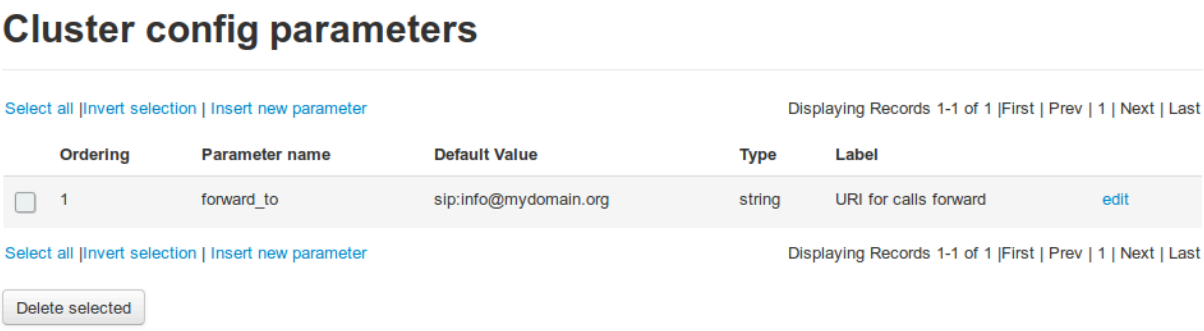

Set specific values for Template Parameters

Once the parameters are defined, their value can be customized per node or config group on the "System -> Config Groups" screen.

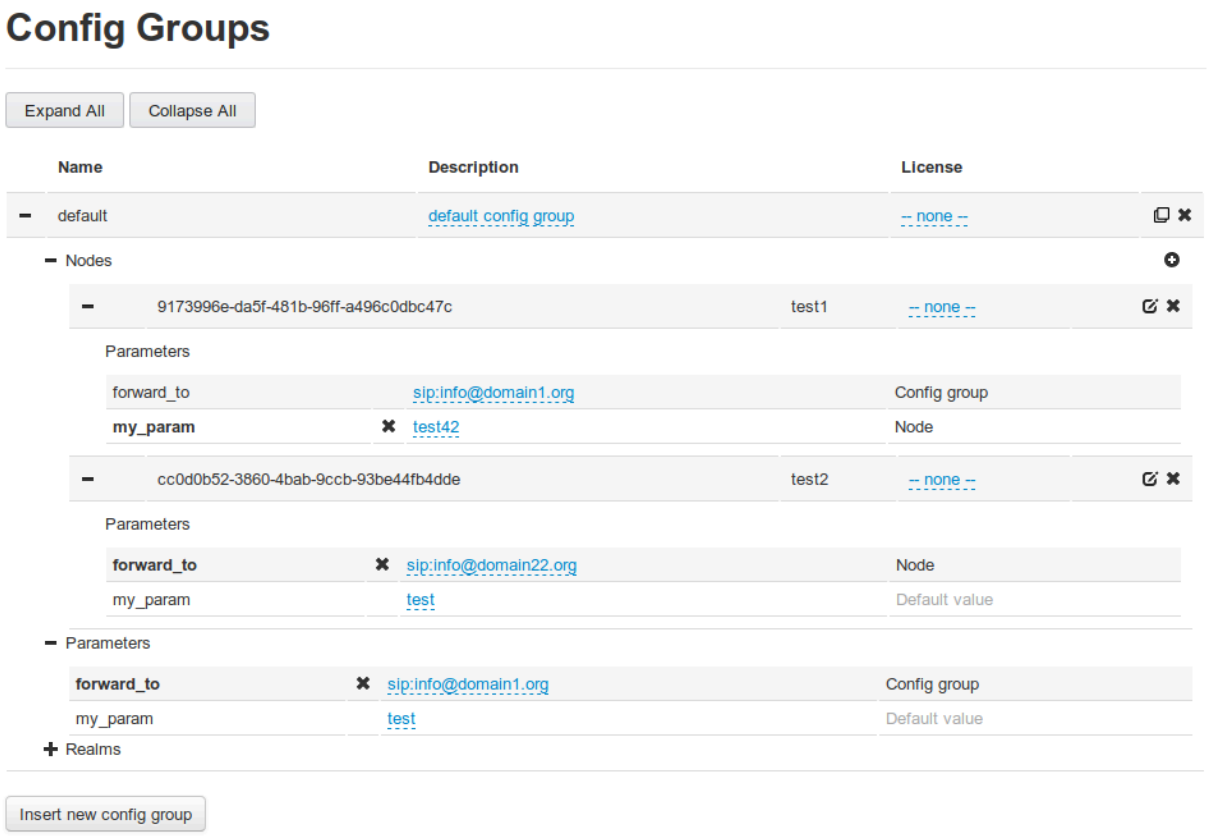

If the value for a defined parameter is not customized, its default value is used. Otherwise it is replaced by the config group value or node value (node specific value takes precedence over config group value).

4.11 ABC SBC System administration

This Section describes the administrative tool available on the ABC SBC. There is also the CLI reference, see Command Line Reference.

4.11.1 User Management

There are two ways how to administer User Accounts and granular access control for the ABC SBC: using GUI and using CLI. The primary method is via GUI, the CLI method is restricted in functionality. Both methods are described in the following subsections.

In opposite of other SBC configuration the changes in user accounts take immediate effect. They do not require activation of SBC configuration.

The access control concept is based on the notion of group membership. Groups define at a granular level permissions to perform specific actions, such as GUI access, viewing and/or modifying the ABC SBC topology, monitoring various aspects of the ABC SBC, accessing RPC, firewall administration, etc.

A user gains all associated privileges by being assigned to a group. A user can be member of multiple groups.

The system comes with preconfigured user accounts as described in the Section Default User Accounts.

GUI User Management

Access to the administrative GUI can be managed using User Management (menu: *"CCM ▸ Users"*) and Group Management (menu: *"CCM ▸ Groups"*).

User Management allows to add, delete and modify users – their passwords and group membership.

Group Management allows to enable/disable the respective permissions for a group. Once set and applied, it applies to all group members.

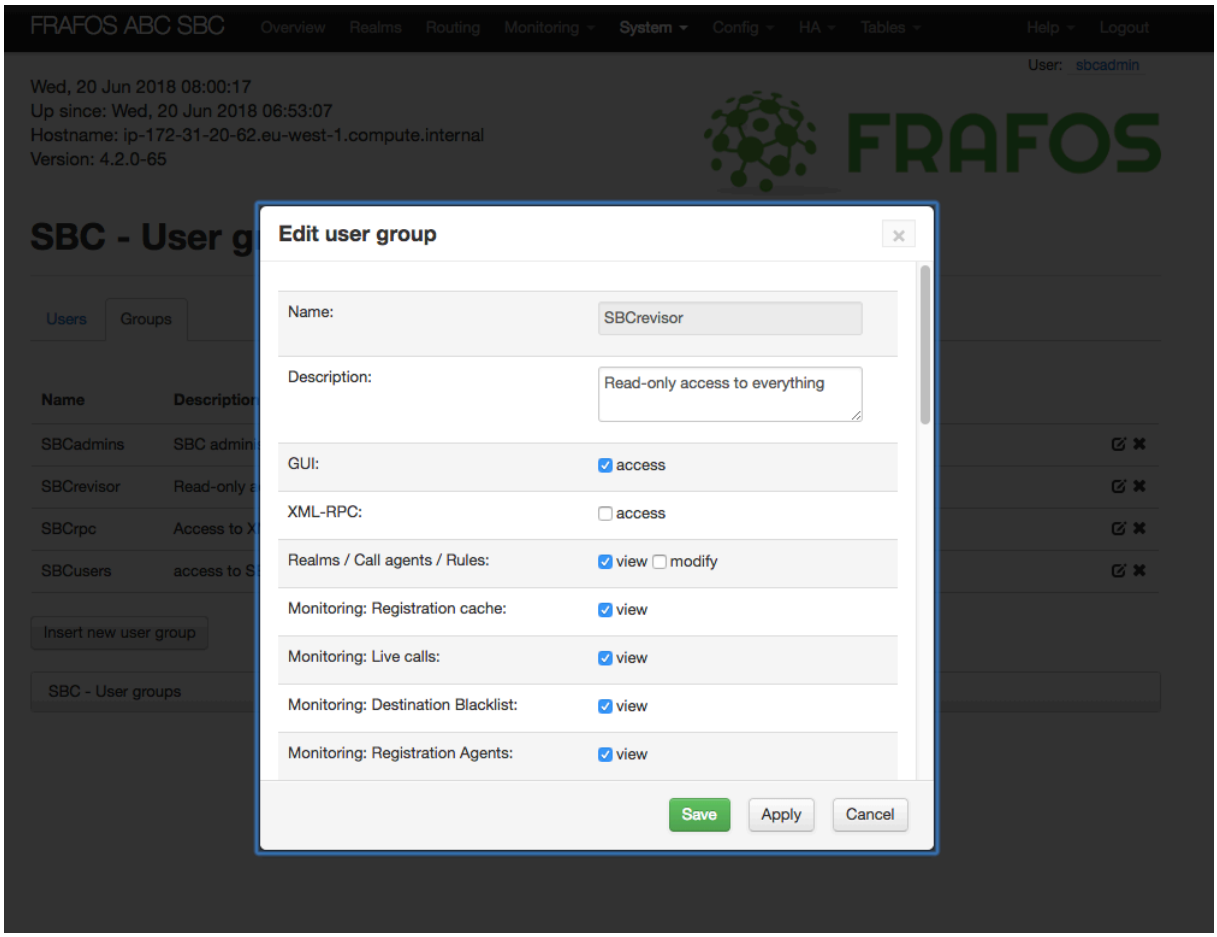

Figure 133: Example: Group Management

CLI User Management

The CLI User management permits to create new users and assign them to a group. This subsection lists available commands.

To add a new user, use the CLI the following way:

```
% sbc-add-user '--password=DoyoulikeH.323?' admin
```

The new user comes without any privileges and must be assigned to a group. To assign the new user to the administrative group with all permissions use the following command:

```
% sbc-add-user admin SBCadmins
```

To unlock locked account due to unsuccessful login attempts use this command:

```
% sbc-user-passwd -u admin
```

Additional commands include:

- **sbc-del-user** to delete a user
- **sbc-list-groups** to show existing user groups

- **sbc-list-users** to show existing users
- **sbc-user-passwd** to change a user's password

GUI Two Factor Authentication

Starting with the 5.5 ABC SBC release, administrators can enable two-factor authentication (2FA) for enhanced security. The ABC SBC supports one-time passcode (TOTP) authentication method for which an external authenticator application (Microsoft Authenticator, Google Authenticator, etc.) is required.

To enable 2FA for a new user, navigate to the `CCM --> Users` and click on `Insert new user` button. In a newly opened pop-up fill all necessary details and check the `Enable two factor authentication with authenticator app` checkbox and click on save button.

To enable 2FA for an existing user, navigate to the `CCM --> Users` and click on edit icon for selected user. In a newly opened pop-up check the `Enable two factor authentication with authenticator app` checkbox and save your changes.

During next login a user still provides his username and password but once the `Login` button is clicked, a screen for 2FA using one time passcode is opened. When a user logs in for the first time using 2FA, he needs to add this account into his authenticator application. To do that, just scan the code with your authenticator app or type the secret key into it. Once this is done, fill the 6 digit generated code and proceed with the login. After successful login, user will be prompted just about the 6 digit code next time (once valid credentials were provided.

TOTP is also supported for LDAP users. If this functionality is required, ensure that the Cluster Config Manager is correctly configured for LDAP. Once configured, the user must first log in with their credentials. After this initial login, the user's name will appear on the `CCM → LDAP users 2FA` screen. To enable TOTP, check the checkbox next to the user's name. TOTP will then be required for the user's next login attempt.

Passwordless authentication

Starting with the 5.5 ABC SBC release, administrators can enable passwordless authentication for any configured account. This allows administrator to login into Cluster Config Manager without existing credentials using passkeys authentication.

To enabled passwordless authentication some mandatory configuration is required:

- A valid TLS certificate must be configured in `CCM --> CCM Config --> Management access tab`.
- To access CCM a hostname must be used in a browser. IP address is not allowed.
- If the CCM hostname is set properly then the `Relying Party ID` (`CCM --> CCM config --> Login tab`) can be left untouched and hostname value is used instead. If a different ID is needed then it can be modified accordingly. Important is the Cluster Config Manager TLS certificate must be valid for `Relying Party ID` that means its value must appear in `Subject Alternative Name` section of a certificate.

Once all required configuration is ready, login with a targeted user and click on user name in the right top corner, select `Passwordless Authentication`. On newly opened screen, click on `Add a passkey` button. If this is a first time you use your token, you might be prompted to choose a security PIN. Otherwise you might be asked to provide your security PIN. Once this is done you will finish a passkey creation by touching your token.

When login with such a user, click on `Login with a passkey` button. You will be prompted to provide a security PIN for your token and then touching the token to finish the login process. After successful verification you will be logged in into Cluster Config Manager.

The passkey authentication is also available for LDAP users.

> **Info**
>
> A user account with enabled passwordless authentication can use both methods to login into Cluster Config Manager - either using standard login credentials or using passwordless authentication. Both methods coexist next to each other.

4.11.2 Server Administration

If a maintenance of the server running ABC SBC is needed, it is possible to shutdown the container from the GUI. The following buttons may be available on the Nodes screen, which can be accessed via the 'System > Nodes' link under the hamburger menu icon.

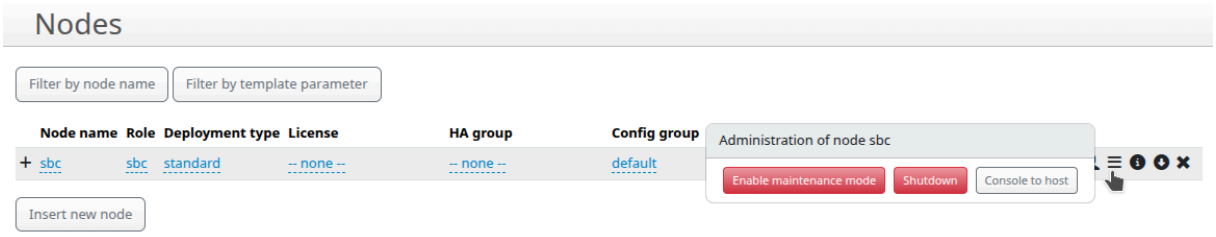

Figure 134: Example: Server administration

- Shutdown: performs soft shutdown of the container.
- Force HA offline: displayed only if node is running as HA pair, it puts forcibly the node into HA FAULT mode. In this mode the HA resources (VIP addresses, routes) and signaling application are stopped on the node, and should be moved to the other node which becomes new HA MASTER (under the condition that the other node is up.) It can be used when upgrading ABC SBC, to make sure the node being upgraded is not processing traffic. The same can be achieved also from command line using the "sbc-ha-offline" command on the specific node.
- Un-force HA offline: displayed only if node is running as HA pair, reverts the forcibly set HA FAULT mode that was set using "Force HA offline", meaning it puts the node back to normal mode, in which a node can be either HA BACKUP or MASTER depending on negotiation with the other node. The same can be achieved also from command line using the "sbc-ha-online" command on the specific node.
- Console to node: open SSH connection to the server. The SSH connection is opened by `xterminator` user from CCM to `root` user on SBC. So if you need to setup automatic login using SSH keys, just put SSH public key of `xterminator` user on CCM to the SBC nodes. SSH is looking for the keys in `/data/sbc/xterminator/.ssh` on CCM and in `/data/root/.ssh` on SBC. Note: This functionality requires SSH application to be enabled on one of SBC interfaces. By default, the SSH is disabled. Note: The direct SSH connection to the SBC node is available only for 5.4 nodes and older.
- Console to host: open an SSH connection to the host where the SBC container is running
- Enable maintenance mode: If the "maintenance mode" is activated, the SBC answers 503 to any request. The maintenance mode can also be enabled from command line. See Maintenance mode for more details.
- Disable maintenance mode: Once "maintenance mode" is disabled, the SBC starts operate normally again. The maintenance mode can also be disabled from command line. See Maintenance mode for more details.

Note: Current HA status can be checked either from "System status" screen or using "sbc-ha-status" command line command on the specific node.

SSH to host

The SSH to Host feature allows you to access the host of the ABC SBC container via SSH connection from the GUI. This feature is only available if the host information has been configured for the SBC node.

Node hosts can be configured on the 'System > Node hosts' screen. The node host must be assigned to the node on the 'System > Nodes' screen. The node hosts can be configured using the following parameters:

- Name: The name used to reference the host in the GUI.
- IP or hostname: The IP or hostname of the host to which the ABC SBC container is running on.
- SSH Port: The port used to access the host. Default is 22.
- SSH Username: The username used to access the host.

The SSH connection is opened by the `xterminator` user from the CCM. If you need to set up automatic login, using SSH keys, simply add the SSH public key of the `xterminator` user on the CCM to the node hosts. SSH looks for keys in `/data/sbc/xterminator/.ssh` on CCM and usually in `~user/.ssh` on the host.

4.11.3 Backup and Restore Operations

ABC SBC Configuration Management

The ABC SBC configuration is stored in a local MariaDB database on the configuration master node. When the administrator applies the changes using the **"Activate SBC configuration"** link, an automatic snapshot of the configuration database is created and is labeled as „Automatic Snapshot" in the list of available snapshots.

The SBC administrator can manually trigger the generation of the configuration DB snapshot from the GUI. When the configuration snapshot is created, it is recommended to write a short comment to note what exactly has been modified. Optionally also content of provisioned tables database can be included in the snapshot.

These configuration DB snapshots can be accessed using the **"System ▸ Config Management"** screen, see Fig., Managing the ABC SBC configuration backups. From the GUI, the administrator can create new snapshots or change or add a comment to an already existing snapshot. To restore a saved configuration the administrator can use the **"Load config"** link of the desired configuration snapshot, or the **"Load provtables"** link to load the content of provisioned tables (if the snapshot contains it).

> **⚠ Warning**
>
> The content of provisioned tables is integral part of the configuration. It should be part of every backup unless there is a good reason to skip it; for example size of the database with provisioned tables.

> **⚠ Warning**
>
> Neither configuration nor content of provisioned tables can be loaded from a snapshot created on newer Cluster Config Manager release. For example, it is forbidden to load configuration or provisioned tables on a 5.2.x Cluster Config Manager if the snapshot was created on 5.3.x. This means that before upgrade to a new release it is really important to create a DB snapshot on the old Cluster Config Manager to allow for smooth downgrade in case something goes wrong.

Snapshots may also be downloaded and uploaded from the same GUI page. The only supported format is *.tar.gz*. Filename doesn't matter is case of upload, but for usability the default file name in case of download is: *sbc-backup-<date>_<db version>_<sbc version>_<snapshot name>.tar.gz*

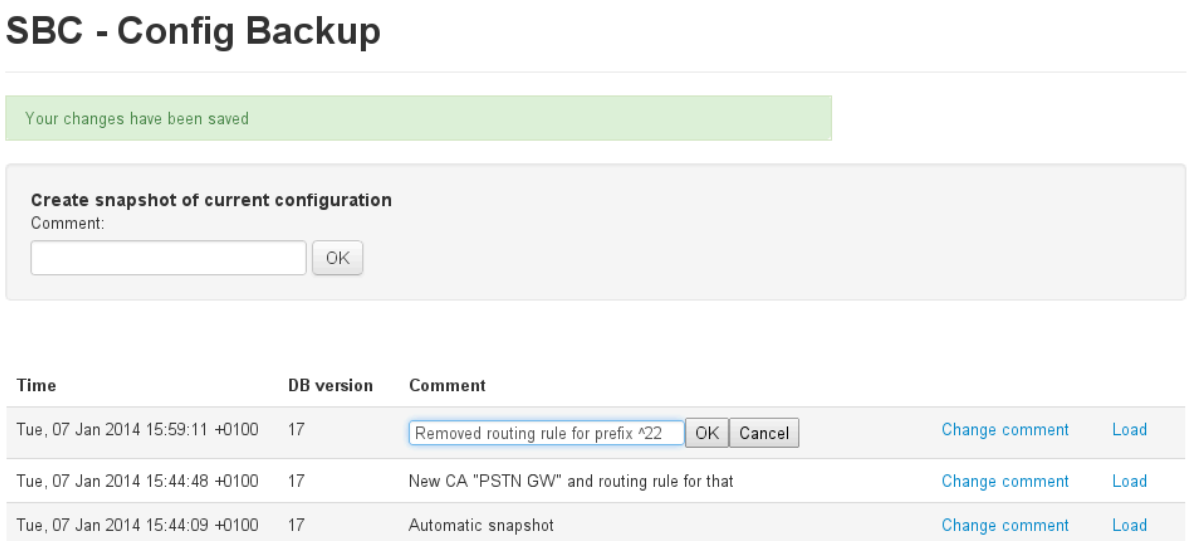

Figure 135: Managing the ABC SBC configuration backups

Info

When the configuration DB backup is loaded, the configuration is NOT automatically applied. The administrator should check if the restored configuration is the correct one and then has to manually apply it using the "Activate SBC configuration" link.

ABC SBC Configuration Backup

Apart from the above configuration snapshots, it is also possible and recommended to use automatic daily ABC SBC backups, which can be enabled under Config / Global Config / Backup tab. The following options can be set there:

- **Equivalent settings as for CCM** - if enabled, the settings on this Backup tab will not be used on Sbc nodes, but the same settings as configured for CCM node (under CCM / CCM Config / Backup page) will be applied on Sbc nodes.
- **Create daily Sbc configuration backups** - this enables the daily backup.
- **Include provisioned tables in daily backups** - when enabled, also content of whole provisioned tables database will be included in the daily backup. It is enabled by default and recommended.
- **Number of days to keep backups** - sets the retention period for the daily backups. On each backup run, all backup files older than the specified number of days will be deleted. Use 0 to disable any automatic removal.
- **Destination directory for backups** - sets the directory, to which the daily backup files will be created. Default setting is "/data/sbc/configs". The partition holding this directory should have enough space for the daily backups.
- **Full path to extra files or dirs to include in backup, separated by comma** - it is possible to include custom files or dirs into the backup. The paths to files or directories has to be full path. Directories will be included recursively. It is also possible to use wildcard "*". The path must not contain comma character.

It is highly recommended to enable the daily backups and include the backups destination directory to customer off-server backups to external device, or at least to copy the backup files to external device after important changes done on ABC SBC configuration, to be able to recover SBC node in case of hardware failure. Note that to minimize external backup impact on ABC SBC performance, a solution allowing to use only idle I/O and CPU should be used.

The daily backup files are gzipped tarball archives and contain the following data, which can be used (directly or as a reference) when recovering ABC SBC configuration: the main ABC SBC configuration database dump (backups from master node), optionally also whole provisioned tables database dump, versions of important RPM packages installed, local configuration files templates (if existing), system network interfaces configuration files, system hostname, system hosts file, MariaDB server configuration file, node UUID info, optionally also root user SSH authorized keys files.

ABC SBC Recovery Procedure

In case ABC SBC server dies, it is possible to recover it using the following steps. In case more ABC SBC nodes are used, this procedure differs depending on if recovering main configuration master CCM node or not.

Steps to be done when recovering configuration master CCM node:

- Locate latest ABC SBC backup file from the daily ABC SBC backups.
- If needed, prepare new server to host container(s), check system network interfaces, routes, hostname. Pay attention e.g. to possibly changed system interface names.
- Install CCM container, following normal installation steps, see Sec. Container Installation. Use the same major release line (like 5.0.x) that was there before.
- Restore ABC SBC configuration from the backup, either using GUI or by calling the following command, where the <backupfile> is the daily backup gzipped tarball file.

```
% ccm-restore --rest-all --bckfile <backupfile>
```

- Access ABC SBC administration GUI and check the restored configuration.
- Activate the configuration using “Activate Sbc configuration” link, which is available at bottom of Overview GUI page.
- Check if the ABC SBC node(s) pulled new configuration using Monitoring / System status.

Steps to be done when recovering a node not being configuration master, where the configuration master node is still working:

- If needed, prepare new server to host container(s), check system network interfaces, routes, hostname. Pay attention e.g. to possibly changed system interface names.
- Install ABC SBC container, following normal installation steps, see Sec. Container Installation. Use the same major release line (like 5.0.x) that was there before.
- Configure the new container to point at the existing configuration master node by passing the appropriate environment variables (`MASTER`, `CONFIG_USER`, `CONFIG_PASS`, …) when creating it; see Configuring the SBC container. Important: when recovering a node, it is highly recommended to set `NODE_UUID` to the UUID that was used previously on the node being recovered. It can be found under “Details” of the node on the “System status” page or on the “Nodes” page under the “System” menu on the ABC SBC config master GUI. If the previous node UUID is not provided, the node can still be recovered, but in case there was some configuration specific for that node (like system interfaces assigned to that particular node), it will have to be fixed to apply to the newly created node in the GUI. (Alternatively, the `sbc-init-config` script can be run from inside the container.)
- Access ABC SBC administration GUI on configuration master CCM node and check configuration. In case system interfaces names differ on the new server, update the logical to system interfaces mapping under System / Interfaces.
- Activate the configuration using “Activate Sbc configuration” link on configuration master node, which is available at bottom of Overview GUI page.
- Check if the recovered node pulled new configuration using Monitoring / System status on configuration master node GUI.

Manual Backup of the Complete SBC Configuration

It is also possible to manually backup the important files:

- **SBC database: administrator can manually create a full SBC**

  configuration backup using **"ccm-backup"** command which creates a backup files into *"/data/sbc/configs"* directory, where also automatic DB backups are stored.

The following command line options can be used:

```
% --prov
```

Optionally also content of provisioned tables db can be included.

```
%  --comment <comment>
```

A backup comment can be specified.

```
%  --bckfile
```

If used, the backup files will be put also to gzipped tarball. Filename will be automatically generated from date, version and comment.

```
%  --bckdir <dir>
```

Specifies the directory where to save the backup file, if –bckfile option is used. Defaults to /data/sbc/configs.

```
%  --filename <file>
```

If used together with –bckfile, save the backup under specified full file pathname instead of automatically generated name.

```
%  --remove
```

If the –bckfile option is used, by default the backup files are left also in the default directory (/data/sbc/configs/). When this option is used, they will be deleted from that directory after creating the gzipped tarball file.

```
%  --incl-system
```

Enables inclusion of system stuff - hostname, hosts and root user ssh authorized keys and password.

```
%  --incl-dbsnaps
```

Include also configuration database snapshots. This can be used only on config master node.

```
%  --incl-all
```

Enables inclusion of all provisioned tables, system and ssh settings and keys at once.

```
%  --excl-tls
```

Do not backup TLS profiles. This option can be used only when creating backup of config master node.

```
%  --cfg-changes CFG_CHANGES
```

Use a file as a source for config-changes.json.

```
% --status STATUS
```

Write status information into a file.

```
% --force
```

Force the backup even if the container init is not finished.

```
% --quiet
```

Print only warnings and errors, no info messages.

```
% --verbose
```

Verbose mode.

- **CDRs:** the customer's billing system should regularly download CDRs generated by the SBC which are stored on the SBC for 93 days by default. CDRs are stored to the *"/data/cdr"* directory.
- **Logs:** log files are stored in *"/var/log/frafos"* directory
- **Traffic logs:** traffic logs created by the **"Log received traffic"** action are stored in the *"/data/traffic_log"* location.
- **Recording files:** recording files are stored in the *"/data/recordings"* location.

Manual Restore of the Complete SBC Configuration

For manual backup restore, the command "ccm-restore" can be used. Be careful when using it, as it can overwrite also various system files. The following options can be used:

```
% --bckdir <dir>
```

The backup from specified directory will be restored.

```
% --bckfile <filename>
```

The backup from specified backup gzipped tarball file will be restored.

```
% --prov
```

Restore also provisioned tables, if those are included in the backup.

```
% --provonly
```

Restore only provisioned tables, do not restore the main SBC configuration.

```
% --rest-ccm
```

Restore CCM configuration.

```
% --rest-system
```

Restore all system files (ssh root user authorized keys and password, hostname, hosts). Use with caution.

```
% --rest-dbsnaps
```

Restore config database snapshots, if included in the backup tarball. Can be used only on config master and only when restoring from tarball.

```
% --rest-all
```

Restore everything included in backup (Sbc config, provisioned tables, ssh, system, sbc pull/push config).

```
% --excl-tls
```

Do not restore TLS profiles, if those are included in the backup.

```
% --quiet
```

Print only warnings and errors, do not print info messages.

```
% --verbose
```

Verbose mode

4.11.4 How to setup a Semi-redundant CCM on ABC SBC

This section describes steps needed to setup a "semi-redundant" cluster config master (CCM) node for Frafos ABC SBC.

With current ABC SBC release, there is no official support for redundant CCM node. But with help of this document, a backup CCM node plus automatic transfer of configuration backup to it can be set up, which will be ready to take over the role of cluster config master for SBC nodes in case of main CCM node failure.

Note: the official full support for geo-redundant CCM is planned for future ABC SBC release.

We refer here to the main active CCM node as "primary CCM" and to the backup node as "backup CCM".

The procedure is based on standard ABC SBC backup / restore using configuration snapshots, as described in ABC SBC handbook Sec. Backup and Restore Operations. Please refer to that for details.

Note that the switch to backup CCM still requires manual intervention.

Setup primary CCM node

Start and configure the primary CCM node the standard way.

After initial configuration is done, note the primary CCM node "node UUID" value, which can be found on GUI "System" / "Nodes" screen, value from the row for the CCM node itself.

Setup backup CCM node

Start clean fresh CCM container for the backup CCM (on some other physical host), using the same SBC release and the same installation way as the primary CCM, but do not configure anything in GUI.

Configure configuration snapshot backups

On primary CCM GUI, navigate to "Config" / "Global config" screen, and there select the "Backup" tab.

Enable the "Create daily Sbc configuration backups" option.

Make sure that also “Include provisioned tables in daily backups” option is enabled.

The option “Destination directory for backups” sets directory, to which the daily configuration snapshots are saved. Default directory is “/data/backups” local directory on the primary CCM node. It can be also possibly pointed to a specific directory which is mounted externally to the primary CCM node, e.g. from external network share device or via NFS.

Activate new configuration from CCM GUI.

Optional step: in case the backup done daily would be too low frequency, it is possible to change that according to customer need, e.g. to be done every hour. The command which performs the daily backup is “sbc-daily-backup” and by default it gets started from “/etc/cron.daily/sbc-backup”. This can be customized by administrator e.g. by moving the file to cron daily directory, to make it run every hour:

```
% mv /etc/cron.daily/sbc-backup /etc/cron.hourly/
```

But please consider the fact that the configuration snapshots contain also all provisioned tables data, so they can be big. Also the change like this, to set more frequent backups, won’t survive possible CCM container replacement with newer one.

Setup configuration backups transfer to backup CCM node

The configuration snapshot backup files need to be transferred from the primary CCM node to some other safe location, to ensure they do not get lost in case of primary CCM node failure, or can be transferred directly to the backup CCM node.

The recommended way is to manage the files transfer from other external customer server, not from the primary CCM node itself, to avoid loosing that functionality when e.g CCM container is replaced with newer one.

Depending on customer deployment, possible ways to achieve this are like:

- If the daily backups on primary CCM node are created to some externally mounted directory, customer can setup some regular way of copying those files to the backup CCM node “/data/backups” directory, e.g. using “scp” secure shell copy, or “rsync” program. In this case please also pay attention to available space on target location on backup CCM and possibly delete old files (rsync option `“--delete”` can be used for that).
- Another option is that the latest configuration backup can be copied from external location to the backup CCM only when primary CCM failure happens.
- If the configuration snapshot backups are created only to “/data/backups” local directory on primary CCM, it is possible to setup e.g. “rsync” command to be run periodically, which will transfer whole “/data/backups” directory content in a efficient way from primary CCM node to backup CCM node directly, using SSH as underlying protocol. The transfer can be set up to be initiated either from primary CCM side or from backup CCM side. We recommend to initiate it from backup CCM side.

Example of rsync command to synchronize the configuration backups (assumes ssh on port 24 on primary CCM node, 192.168.0.1 is the IP address of the primary CCM), to be run on backup CCM node:

```
% rsync --delete -r -e 'ssh -p 24' root@192.168.0.1:/data/backups /data/
```

This command can be run periodically e.g. using system “cron” service, by creating a file like “/etc/cron.d/rsync” with the following content on the backup CCM node, to make it run every hour (at 10 minutes past every hour):

```
10 * * * * root rsync --delete -r -e 'ssh -p 24' root@192.168.0.1:/data/backups /data/
```

Note: if ssh is being used as underlying protocol for rsync, it is possible to make it work from backup CCM to primary CCM without passphrase using the following commands on the backup CCM (192.168.0.1 is the IP address of primary CCM):

```
% ssh-keygen
% ssh-copy-id -p 24 root@192.168.0.1
```

After the setup of configuration backups transfer is done, make sure that the backup files are really being transferred automatically to "/data/backups" directory on the backup CCM node.

Check also the backup size and available space, and tune global config setting "Number of days to keep backups" on primary CCM GUI (Backups tab). Note that if using the rsync command (with the `"--delete"` option), the files deleted on primary CCM node directory will be also deleted automatically by rsync from the backup CCM node directory.

Steps to make the backup CCM available in case of primary CCM node failure

In case of primary CCM node failure perform the following steps:

- Find the latest configuration backup file that was transferred to backup CCM directory "/data/backups", or if using external backup location copy it to backup CCM "/data/backups" directory.
- Restore the backup using command like this on the backup CCM:

```
% ccm-restore --prov --bckfile
  /data/backups/sbc-backup-2020-10-22_11-36-50_42001027_4.2.22-77_daily_backup.tar.gz
```

Note: if system stuff like ssh keys or hostname was included in the backup too, and restore of that is needed, add also the following option:

```
--rest-system
```

- Access the backup CCM GUI and review the loaded configuration.
- Activate the new configuration from backup CCM GUI, to make it available for SBC nodes.

Steps to be done on SBC nodes to start using new CCM

Once the backup CCM is available and the configuration snapshot backup was loaded on it, and if the backup CCM uses a different IP address from the previous main CCM, re-configure all SBC nodes to use the new CCM address. The recommended way is to update the `MASTER` environment variable on each ABC SBC container (and recreate the container so the new value takes effect); see Configuring the SBC container for the full list of variables. Alternatively, the `sbc-init-config` script can be run from inside the container:

```
% sbc-init-config
```

Alternatively, a DNS hostname can be used as CCM node address on all SBC nodes. In that case it is recommended to use a DNS record with short TTL value, which allows then easy central change of the CCM address just by updating the DNS record, without need to update it on all SBC nodes. (Note: but avoid using more A records under one DNS name, pointing to more IP addresses).

Additional steps and checks

Access the backup CCM GUI "Monitoring" / "System status" screen. Check if all SBC nodes have pulled new configuration from the new CCM.

There may be a duplicate "System status" record shown for the CCM node itself (coming from the node UUID update done initially), but this older CCM node status (which can be identified by the "Last report" column) can be safely ignored, or deleted if using newer CCM which allows that.

Note: in a specific case, when the configuration snapshot that was restored on the backup CCM was not the latest one, and if the SBC nodes were not re-pointed at the backup CCM (whether by updating the `MASTER` environment variable on the containers or by re-running `sbc-init-config`), the nodes will not pull the configuration from the backup CCM after switching to it automatically, because the "configuration version" number used to detect new configuration will be lower on the backup CCM than what the nodes already expect. This should not happen if following this procedure correctly and latest configuration snapshot was restored on the backup CCM. But in case it happens, which can be seen on backup CCM GUI screen "Monitoring" / "System status" by the nodes configuration versions higher than the "Latest config version", it is possible to manually forcibly increase the configuration version of configurations exported from CCM using "sbc-set-confversion <version>" command.

4.11.5 Upgrade Procedure

FRAFOS regularly releases a new version of ABC SBC. New features, modifications and bug fixes are described in a "Release notes" section of SBC handbook for every new release.

If ABC SBC is deployed in the non-HA mode then it is expected a service disruption during the upgrade process. For that reason it is strongly recommended to perform upgrade in the service maintenance window.

If HA is used, before the upgrade is started, both cluster nodes have to be online and all required services running. The administrator should also create a configuration snapshot, see Sec. Backup and Restore Operations.

Please upgrade the CCM node first, and continue with SBC node(s) upgrade only after new CCM is verified to be correctly functional.

The following upgrade procedure applies to ABC SBC container installation.

Container ABC SBC upgrade

When the ABC SBC is deployed as a container, there is no "online upgrade" of the existing (and running) container, but the whole container is replaced by newer version.

It is highly recommended to use separate directory "mounted" to the container for "/data" path, as described in the container install section, which keeps data that is expected to be persistent and makes the container replacement easier. If it is not used, it is possible to manually copy or move the /data content of old container after stopping it to new container before starting it for the first time. If doing so, please pay attention to keeping files permissions and ownership.

When replacing container, please follow these steps:

- Create a ABC SBC backup. Note: the backup file is needed when replacing CCM node container, but might be needed also in case of troubleshooting possible issues, so create it on both CCM and SBC nodes. Use command like this on container:

```
% ccm-backup --incl-all --bckfile
```

It will create backup under "/data/sbc/configs" directory by default. Note the created backup tarball filename.

- Stop the container.
- Backup directory with the old container, by renaming the directory like:

```
% cd /var/lib/machines
% mv <name> <name_backup>`
```

- Create new directory (using the same name as before) and unpack the new container image to it, similar way like listed also in the install section - example:

```
% mkdir /var/lib/machines/<name>
% tar --xattrs -p --numeric-owner -C /var/lib/machines/<name> \
% -xzf frafos-abc-sbc-4-6-1-481.tgz
```

- If the externally mounted "/data" persistent directory is not used, as mentioned above, copy or move content of old container "/data" sub-directory to new container "/data" sub-directory, example:

```
% cp -a /var/lib/machines/<name_backup>/data/* /var/lib/machines/<name>/data/
```

- Start the new container.

Note: Using the same name for the directory means that SBC hostname visible in the GUI will not change. Of course it is also possible to keep the original container's directory name and unpack the new container into a new folder. In that case the SBC hostname which is used in SBC GUI will change.

If the container is CCM node:

- Access the CCM GUI, review the configuration and activate it.

If the container is SBC node:

- If persistent /data directory is not used, recreate the container with the appropriate environment variables (`MASTER`, `CONFIG_USER`, `CONFIG_PASS`, …) so that it can find the Cluster Config Manager; see Configuring the SBC container. Alternatively, the `sbc-init-config` script can be run from inside the container.
- The SBC should automatically pull configuration from the CCM node.

Access the CCM node GUI and check Monitoring / System status page, check for any errors reported by SBC nodes.

4.11.6 Migration from 4.5/4.6 to 5.0

The migration from 4.x ABC SBC product line requires some additional work. First of all it is necessary to prepare one or more host servers which will be serving ABC SBC containers. As a minimum installation CCM and SBC containers need to be deployed. It is possible to use a single server to host both containers or to use separate servers.

In case of HA deployment, two host servers are required to host two ABC SBC node containers.

Frafos recommends to use Debian 13 stable as OS for host server however it should be possible to use any other recent Linux OS.

It is also possible to re-use the existing CentOS 7 server as a host server however we don't recommend this as CentOS 7 is end of life and is no longer supported.

The following upgrade procedure applies to ABC SBC installation.

ABC SBC migration procedure

The migration is very similar to upgrade procedure described in Container ABC SBC upgrade. Please refer to upgrade section for more details.

First of all, it is necessary to do the backup of existing servers. The backup of CCM is mandatory, backup of SBC is an optional but recommended:

```
% sbc-backup --incl-all --bckfile
```

Copy all backup files to secure location, to have them ready, if needed.

Now deploy a new CCM and SBC container(s) as described in XXX. Once ready start all of them. Navigate to the CCM GUI and on the initial login screen use the upload option to upload a backup file which was generated on 4.x CCM.

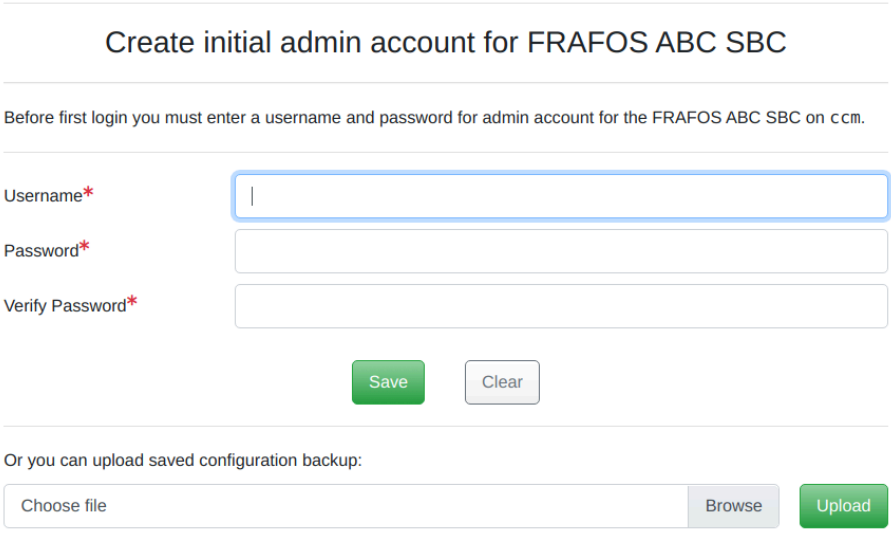

Figure 136: Initial GUI login screen

Once the configuration is restored, you should see the following message in the pop-up window:

```
% Sbc configuration restore finished.
```

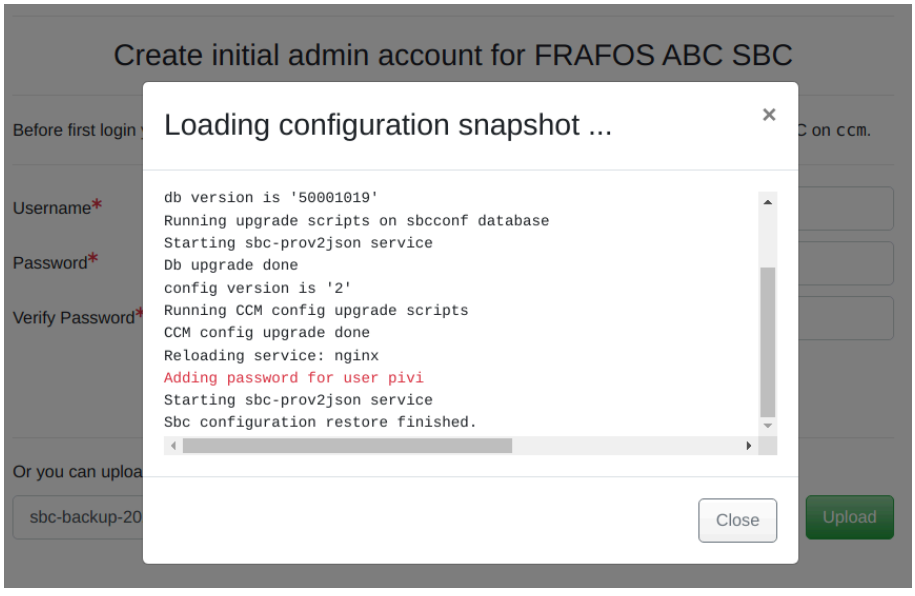

Figure 137: Successful configuration restore

Close the pop-up windows and navigate to the login screen. Use you login credentials from your 4.x installation. Once logged into the GUI, there is a warning about pending configuration changes which need to be activated. At this point there is no active configuration which could be downloaded by SBC nodes. Before configuration activation please **double check** all your configuration and do necessary

changes if they are required. Please pay attention to system interfaces and applications configured on interfaces. Problematic parts can be:

- different system interfaces names as you used on your 4.x setup,
- SSH configuration,
- all hard coded IP addresses which might now be different (for interfaces, interface applications, routing rules or A/C rules),
- all 4.x CCM related interfaces can be removed as they are no longer needed in 5.0,
- there is no XMI interface in 5.0.

If your configuration is OK, then activate it.

Once the configuration is activated, every ABC SBC node must be pointed at the new CCM. The recommended way is to (re)create each ABC SBC container with the appropriate environment variables — at minimum `MASTER`, `CONFIG_USER`, `CONFIG_PASS` — and to set `NODE_UUID` to the UUID previously used by that node. The full list of supported variables is documented in Configuring the SBC container. The previous UUID can be found under the "Details" pop-up of the node on the System –> Nodes page in the Cluster Config Manager GUI.

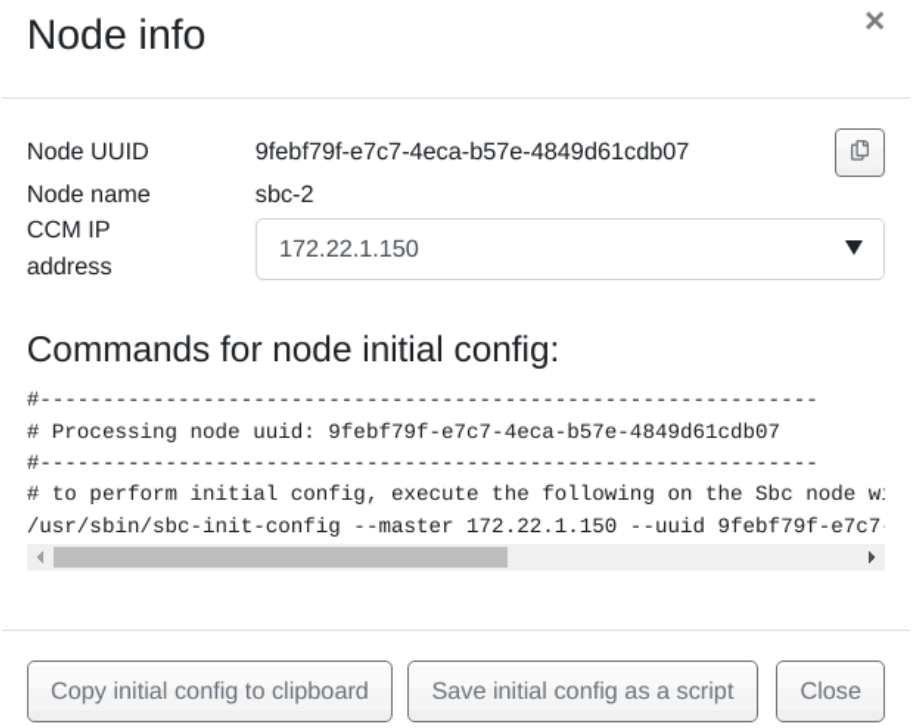

Figure 138: Node info pop-up

Alternatively, the `sbc-init-config` script can be used to configure the container interactively from inside. To do so, SSH to the host server and log into the ABC SBC container, then run `sbc-init-config` and provide the node UUID when prompted (or copy the prepared command from "Copy initial config to clipboard" on the node info pop-up shown above).

Repeat this for every SBC node which you would like to restore.

Expected things which might be surprising

In the 5.0 there is no default root password set for containers. Also SSH is disabled by default. This can cause some unexpected surprises as during migration from 4.x to 5.0 we do not migrate **any** system accounts.

If you were using SSH, you will not be able to use it after migration until you create all necessary accounts again or you upload the SSH authorized_keys file manually into the container. Please note if you create some new system user accounts inside the container then those accounts will be lost during next container replacement while upgrading to a newer version.

During the migration only configuration related information is transferred. However there might be need to migrate also other files like CDRs, audio recordings, traffic logs, prompts. If this is the case, please, transfer all necessary files from 4.x server to 5.0 manually. There is no script which would do this automatically. All above mention data are stored in /data partition in corresponding directories. Please note, starting 5.0 CDRs were moved to /data partition as well.

Data type	4.x location	5.0 location
CDR	/var/log/frafos/cdr	/data/cdr
recordings	/data/recordings	/data/recordings
traffic logs	/data/traffic_log	/data/traffic_log
prompts	/data/prompts	/data/prompts
pcaps	/data/pcap	/data/pcap

Table 15: Data directories mapping

Expected things which might be surprising

In the 5.0 there is no default root password set for containers. Also SSH is disabled by default. This can cause some unexpected surprises as during migration from 4.x to 5.0 we do not migrate **any** system accounts.

If you were using SSH, you will not be able to use it after migration until you create all necessary accounts again or you upload the SSH authorized_keys file manually into the container. Please note if you create some new system user accounts inside the container then those accounts will be lost during next container replacement while upgrading to a newer version.

4.11.7 SBC Dimensioning and Performance Tuning

This section provides background information on typical traffic patterns, its performance implications, and performance tuning possibilities. This information can help to make a more educated estimate than provided in Section Capacity planning. However, confidence can only be achieved by measurement of the target ABC SBC configuration against actual traffic on the used hardware.

The reference hardware we used is Sun SunFire X4170 with the following configuration:
- 2 x Intel Xeon X5570 @ 2.93GHz CPUs, each 4 cores with hyper-threading enabled
- 2 x on-board Intel Gigabit Ethernet adapter
- 8 GB RAM

As an alternative, we measured on a Dell R410 with the following configuration:
- 2 x 4-core Intel X5550 CPU 2.6GHz
- 2 x Broadcom NetXtreme II BCM5716, 8 IRQs/queues
- 12 GB RAM

The alternative results are shown in parenthesis.

On the reference hardware, the maximum performance limits of the ABC SBC have been measured as follows:
- 5000 parallel G.711 calls with media anchoring (3600)
- down by factor of five when transcoding is used,
- call rate of 480 calls per second, without media anchoring
- registration rate of 9900 registrations per second.

Actual use-cases may have significantly different traffic characteristics and therefore the resulting performance may be driven by different limits. In this section we look at the most typical cases: a trunking case with and without transcoding and a residential deployment scenario. All the use-cases assume a single-pass SBC traversal for both SIP and RTP. In the next section, we also summarize the critical configuration aspects that need to be checked when tuning the system for the highest performance.

Trunking Use Case

This case is characterized by handling many calls, both signaling and media, from rather few sources. SIP traffic is not NATed and does not include REGISTER transactions. In this case, the most demanded functionality is media forwarding and the most significant bottleneck is the packet rate of the Ethernet card. Small packets as common with VoIP saturate Ethernet card nominal capacity much earlier than large HTTP packets would. A reasonable Ethernet card shall deliver at least 400 thousand packets per second in each TX and RX direction.

With such a packet rate, the following number of parallel calls can be achieved for the respective number of calls:

codec/packetisation	number of calls
G.711/20ms	5,000 (3,600)
G.729	6,000 (4,680)

Trunking with Transcoding

Transcoding is typically deployed as an additional feature in the trunking case. However, as transcoding is computationally expensive the bottleneck shifts to CPU. The following numbers are achievable on the reference platform:

transcoding	number of calls
G.711-to-G.729	1,000 (750)

Traffic Estimates for Residential VoIP

With residential VoIP, the deployment sees many challenges: the clients are connecting from unmanaged networks over NATs and variety of SIP client types causes interoperability issues. Addressing NAT traversal by enforcing media anchoring (see Section Media Anchoring (RTP Relay)) and frequent re-registration (Section Registration Handling Configuration Options) causes substantial increase in overhead. The ABC SBC keeps the heavy SIP traffic off the infrastructure behind it, however the bandwidth impact on the incoming side must be considered.

Without frequent re-registrations NAT address bindings would expire and SIP devices behind NATs would loose incoming traffic. While other more light-weight methods (STUN, CRLF) exist, re-registrations are safe in that they work with every SIP client and create traffic keeping any NAT bindings alive. The penalty is quite high resource consumption. The “background SIP traffic” is even higher in public SIP services than one could infer from baseline calculation based on re-registration period. Alone use of digest authentication doubles number of REGISTER transactions, many clients send additional traffic to check voicemail status (SUBSCRIBE), announce their online status (PUBLISH), and get over NATs on their own (OPTIONS). As a result, the number of SIP request roughly quadruples against base-line.

The following table summarized empirical impact of driving re-registration traffic to 180 seconds period for population of 1000 subscribers:

	Rate per second (incoming interface)
REGISTER requests (w and w/o digest)	20 pps

	Rate per second (incoming interface)
all requests	40 pps
all requests and answers	80 pps
bandwidth (TX and RX about 1:1)	380 kbps

This load is noticeable both in terms of bandwidth and CPU impact.

The following table present impact of media-relay on bandwidth for 1000 subscribers in peak periods. The underlying assumption is that in peak periods, there is one call for every ten active subscribers.

number of subscribers	1000
parallel calls (10:1)	100
G.711 bandwidth (RX and TX)	197*2*100= 39,400 kbps
SIP bandwidth (RX and TX)	380 kbps
Total bandwidth	~40 Mbps

Performance Tuning

The performance of the system can be increased by proper configuration of the hardware, operating system and the SBC. The following paragraphs list configuration suggestions that are known to bring the greatest performance benefits.

Hardware has expectedly profound impact on system performance. With networking applications, is network cards that shall receive particular attention. Our experience has been that Intel Ethernet cards are at least on part with cards of other vendors and often overperform them. Note that it is necessary that the kernel is using specific card drivers: performance of generic Ethernet drivers is noticeably lower. Hints to hardware specific configuration option are provided in Sec. Hardware Specific Configurations.

Key card driver parameters that don't come preconfigured with the ABC SBC are those specific to Ethernet interfaces. If available, tune the following parameters:

- enable Receive Packet Steering
- increase coalesce and ring buffer size
- bond statically NIC RX queues to CPU cores

Furthermore, you shall also make sure that your ABC SBC configuration is not causing unnecessary load. Configuration options that can considerably increase overhead are especially media relay and registration processing. Media relay shall be avoided if not needed. If an SBC connects networks that are mutually routable, anchoring media may be entirely unnecessary. Also in many cases, when signaling passes the SBC twice on the way in and out, you may need to pay attention to configure the media to pass the SBC only once. Registration processing is primarily driven by the NAT keep-alive interval. We recommend a period of 180 seconds. Shorter intervals will not dramatically improve NAT traversal and will cost performance degradation. Longer intervals could result in expired NAT bindings for NATs that expire too rapidly.

4.11.8 Removing SBC Node

Before SBC node is removed from the system, please make sure it is stopped. Then you could remove it from the system in GUI: *"System ▸ Nodes"* screen.

Removing node in GUI just remove it from the list of nodes for which CCM generate configuration. This does not perform any action on the node itself (like stopping it).

If alive node is removed from the system, it might re-appear again if node auto adding is enabled (see Miscellaneous Parameters).

4.11.9 High Availability administration

This section describes various High Availability advanced topics.

High Availability statuses

Each node in the high availability (HA) setup can have one of the following statuses:

- Unknown: The current status of the node could not be determined, or its role (MASTER or BACKUP) has not yet been verified.
- BACKUP: The node is in standby mode, ready to take over if the master fails.
- MASTER: The node is actively serving as the master in the HA setup.
- FAULT: keepalived has detected error on this node, like if any SIG/Media interface is down or the SBC has been manually taken offline (sbc-ha-offline or via the CCM GUI), or configuration error detected
- STOP: keepalived is in the process of shutting down.

High Availablility switchover

A VRRP failover, i.e., an SBC node in BACKUP state becomes MASTER, always occurs if there is no other node (not in FAULT state) with the same/higher priority.

The priority is influenced by:

- whether the node was previously MASTER (weight 1; to prevent sporadic failover with the same priority)
- whether the sems process (SBC process) is running (weight −2 in case of an error)
- if configured: hardware monitoring (temperature, PSU status; configurable weight)
- if configured: the gateway heartbeating (configurable weight)

External track point for lowering HA priority

Note: this applies to ABC SBC release 5.5 and higher.

It is possible to lower a Sbc node HA priority from external host side. It is intended for use cases like monitoring host hardware health and triggering HA switchover proactively when some serious hardware problem is detected.

The container has a “track file” prepared, using the following path inside: /data/sbc/ha/ha_track

From host side, it is possible to write a numerical value as text to this file. On each file modification, the Sbc HA priority is lowered by the provided value. Under normal circumstances, writing value “1” loweres the HA priority, but does not cause HA switchover yet, while writing value “2” causes HA switchover (if that Sbc node isn’t already HA backup). Writing “0” returns the HA priority to normal.

When non-zero value is written to the track file, the Sbc node status is set to warning on CCM “Nodes status” screen too.

It is also possible to write optionally a text reason (comment) to a file “/data/sbc/ha/ha_track_reason”, which is then shown in the Nodes status details, if the status is set to warning.

To access the track file(s) from host: if /data is provided to container via separate host directory, the file can be simply written directly under the corresponding host path. If the /data is provided using a volume, the dir with track file(s) can be passed to the container create command using “–mount” podman (or similar for e.g. docker) option - example:

```
% --mount type=bind,source=/some_path_on_host/ha,destination=/data/sbc/ha
```

4.11.10 API tokens

API tokens are long-lived credentials used for authenticating automated clients without interactive login. Manage them in *CCM –> API tokens*.

The token is sent in the standard Authorization header as a bearer token:

```
Authorization: Bearer <token>
```

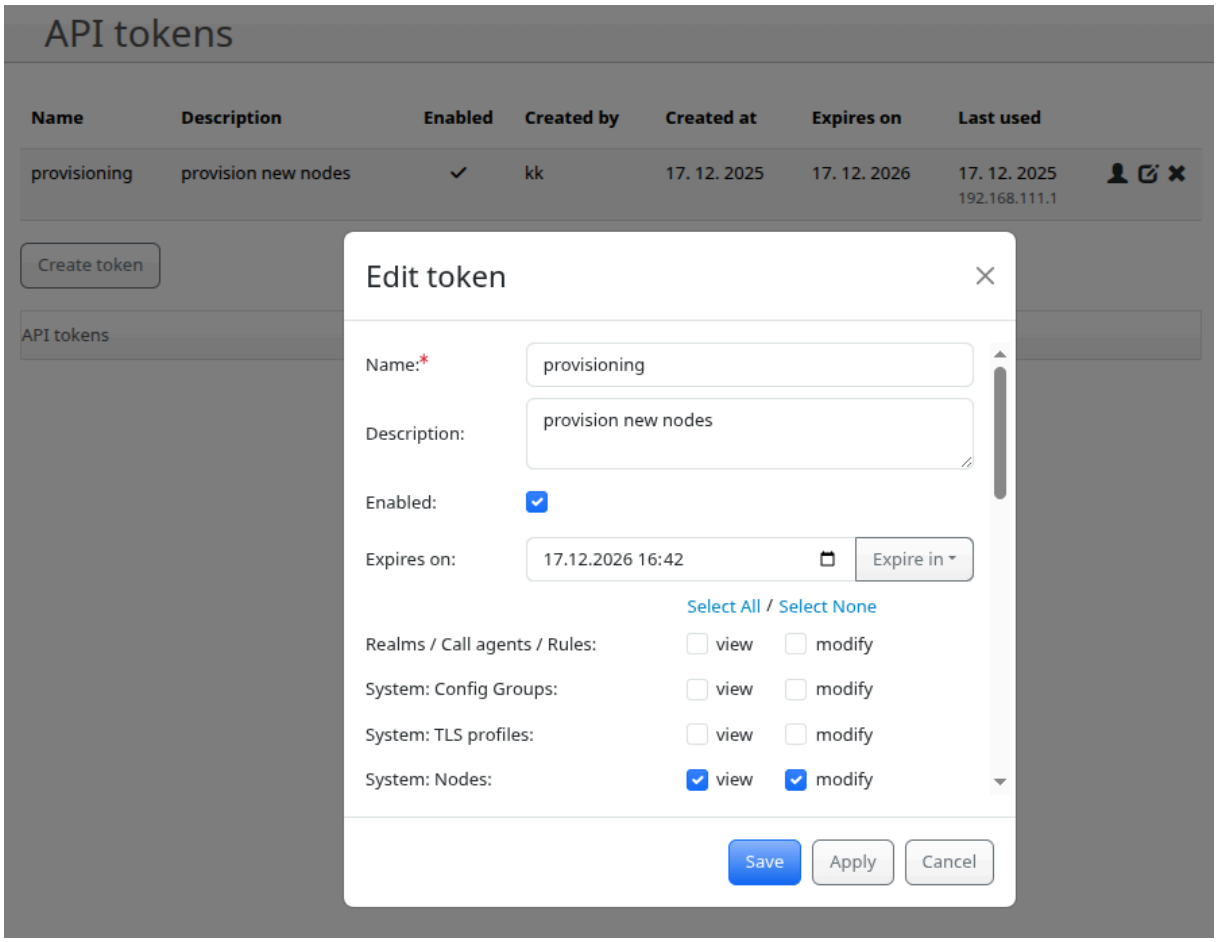

Figure 139: Example: API token management

There are two token types:

- **CCM API tokens** authenticate against the Cluster Config Manager REST API. Their capabilities correspond to the same permission model used for Cluster Config Manager users and groups.
- **SBC API tokens** authenticate against the ABC SBC node API. They use a separate set of capabilities specific to SBC node operations (see SBC API token capabilities). SBC tokens are part of the Cluster Config Manager configuration and are embedded in the ABC SBC JSON configuration that nodes receive via pullconf or pushconf. This means that any change to SBC tokens — including restoring a Cluster Config Manager configuration backup — requires the nodes to receive an updated configuration.

Token permissions and lifecycle

When a user creates a token, the selectable capabilities are limited to the user's current permissions. After creation, the token is independent of the user account and is not affected by later user changes. If the user is removed or their access is reduced, existing tokens remain valid with their original permissions until they expire or are revoked.

Token permissions are defined on the token itself and are not linked to any user or group. You can inspect a token to see its effective capability set.

An expiration date is proposed by default, but it can be cleared if a non-expiring token is required.

The token value is shown only once at creation time. Store it securely, otherwise you will need to create a new token.

Tokens can be revoked from the *CCM –> API tokens* screen.

Auto-generated SBC API token

The Cluster Config Manager automatically generates an SBC API token for its own communication with ABC SBC nodes. This token is stored in the Cluster Config Manager configuration file and is distributed to all ABC SBC nodes as part of their JSON configuration (alongside any user-created SBC tokens). It does not appear in the token management screen.

To renew this token, go to *CCM –> CCM config –> SBC security* and click **Renew**. A confirmation dialog will appear and warn that all ABC SBC nodes must receive the updated configuration. After confirming, the Cluster Config Manager generates a new token. Nodes configured in push mode receive the updated configuration immediately via an automatic push. Nodes configured in pull mode will pick up the new token on their next configuration pull.

> **Info**
>
> After renewal, the Cluster Config Manager keeps the previous token as a fallback. When communicating with an ABC SBC node, the Cluster Config Manager first tries the current token; if the node rejects it (HTTP 401), the Cluster Config Manager retries with the previous token. This ensures that nodes which have not yet received the updated configuration remain accessible on subsequent attempts.

Best practices

- Set a short expiration period.
- Grant only the minimum capabilities needed.
- Revoke unused tokens promptly.

4.12 Monitoring and Troubleshooting

4.12.1 Overview of Monitoring and Troubleshooting Techniques

The ABC SBC and its accompanying monitoring product, ABC Monitor, are designed to provide real-time insight into service health and user behavior for sake of troubleshooting, trending and security. Routine monitoring and troubleshooting is a key part of a SIP service life-cycle. It is also a complex one: the amount of traffic an SBC must handle is enormous and finding abnormal patterns in such quantity is not entirely easy. This is especially true when the service is exposed to a larger user population and is running on the public Internet. Also varying degree of SIP compliance of attached devices often causes unexpected behavior.

Any abnormal service patterns can have a variety of reasons including unusual traffic caused by a security attacks or broken devices, or administrative shortcomings such as a incorrect rule-base or an under dimensioned system. Even if an abnormal situation does not impact a SIP service as whole but only a particular user it is important to find out what is happening.

Identifying presence and root causes of abnormal situations therefore requires solid data about the operation of the service. Here a virtue of the ABC SBC comes in play: it produces a lot of data reporting on the status of operation. In fact the number of bytes produced for monitoring typically exceeds the number of bytes used for the actual SIP signaling. What may seem disproportional is *the* recipe for the capability to understand and keep the status of operation smooth at any time. Good operational decisions can only be made with reliable intelligence.

In the following chapters we will discuss various methods how to monitor an ABC SBC-powered SIP service operation.

The most detailed and therefore powerful method to monitor the operation is using the **events** produced by the ABC SBC (if the event license is installed). The ABC SBC "documents" what SIP users are doing by issuing a report called event on every important user activity: registering, unregistering, failing to authenticate, completing a call, and so on and so forth. An administrator can even produce his own custom events. The events provide a history of user activity which can be looked backed at and analyzed. In a way, it tries to act as secret police would: it holds "files" on the observed subject that include an exhaustive gap-free activity history. At the same time, the overall collection of events also provides aggregated insights into the overall service health and can be used for example to see how the service usage varies in course of a day.

The events do indeed come in a quantity that may make nailing down a problem or identifying a trend a tedious task. Therefore the ABC Monitor is available from FRAFOS to aggregate and filter the events. In addition to user events, the ABC Monitor also shows the utilization of the system. If a situation requires, the ABC Monitor collects even traffic bits: SIP or even RTP data passing the ABC SBC. For more information about ABC Monitor, please refer to ABC Monitor documentation.

The next chapter, Measurements and Monitoring shows how to monitor the overall system health using SNMP. SNMP is the industry standard for monitoring system health and is supported by many third-party monitoring tools, both commercial and open-source. The FRAFOS ABC SBC reports various OS-related and SIP-related counters using SNMP and can also report custom-based ones.

Additional diagnostic information is available directly in the SBC GUI. There is real-time GUI view of established calls and cached registration entries described in Section Live ABC SBC Information. There is also a possibility to review most recent traffic at IP layer as described in Section User Recent Traffic.

Additional methods for determining service status data are eventually described in the Section Additional Sources of Diagnostics Information.

4.12.2 Live ABC SBC Information

The Frafos ABC SBC allows to inspect its internal state in the administrative GUI.

Registration Cache

Registration Cache plays a significant role in off-loading registers, see Section Registration Caching and Handling for more details. The actual Content of the ABC SBC registration cache can be inspected using the web interface under the **"Monitoring ▸ Registration cache"** link, see Fig. Registration cache.

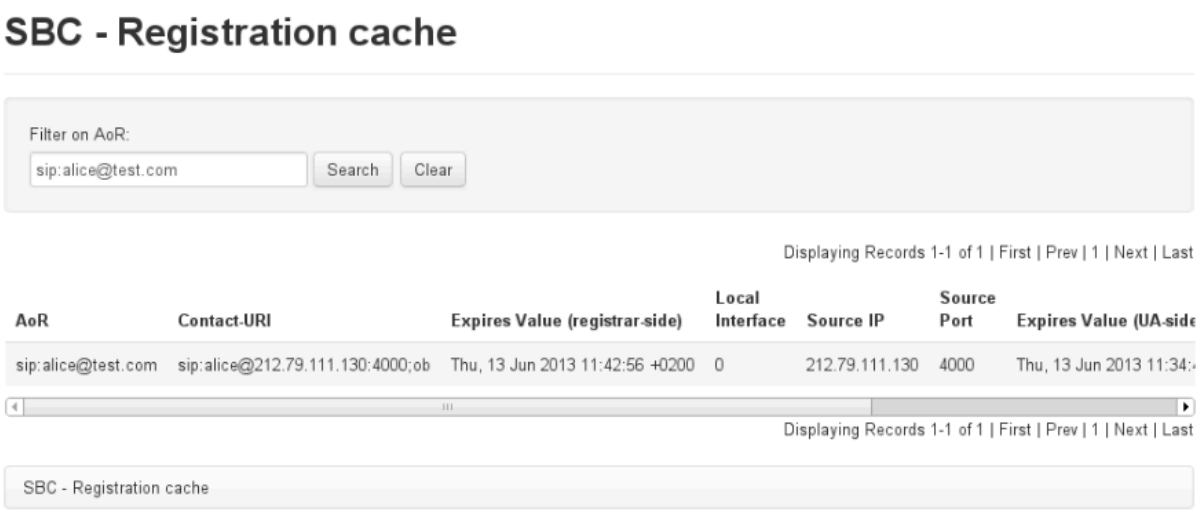

Figure 140: Registration cache

The following information is displayed for each entry:

- *AoR* - Address of Record. SIP URI address that is associated with none, one or more user Contacts by the SIP registration procedure.
- *Contact-URI* - Contact registered by the user agent and associated with an AoR.
- *Expires Value (registrar-side)* - registration expiration at registrar side. This is the time when both the downstream registrar and the ABC SBC will let the contact expire.
- *Expires Value (UA-side)* - registration expiration at client (UA). This is the time when the ABC SBC expects the client to re-register. Failures to re-register timely are ignored to keep the client reachable even if it its re-registration procedure doesn't work accurately. Because of REGISTER throttling feature (see Section Registrar off-load) the actual value may be different (earlier) from *Expires Value at registrar-side*.
- *Source IP* - IP address where the REGISTER was received from
- *Source Port* - port where the REGISTER was received from
- *User Agent* - user agent identity (content of User-Agent header in REGISTER message)

Live Calls

"Monitoring ▸ Live calls" shows list of active calls, i.e. calls that have been forwarded and established. The calls appear there from the time when a 200 SIP response is received from a downstream SIP element, till the call is terminated. Calls that are in so-called "early media" or "ringing" status do not show, neither are locally processed calls shown (e.g. calls processed using Onboard Conferencing).

Since the ABC SBC acts as a SIP B2B user agent, two call legs are shown for each established call:

- A leg (originating leg) - SIP dialog established with caller
- B leg (terminating leg) - SIP dialog established with callee.

Information displayed for each call leg include:

- Source IP - IP address where the REGISTER was received from
- Source Port - port where the REGISTER was received from
- Call-id - SIP dialog identifier
- Remote party - URI of remote party (equals the *From* URI for A leg and the *To* URI in case of B leg)
- Remote target -Contact of remote party
- Local party - Local URI.
- Dialog state - Current state of the SIP dialog
- Call start time - Time of call setup

The administrator can manually terminate the call using the **"kill"** link and inspect call status details using the **"Call Status Information"** link.

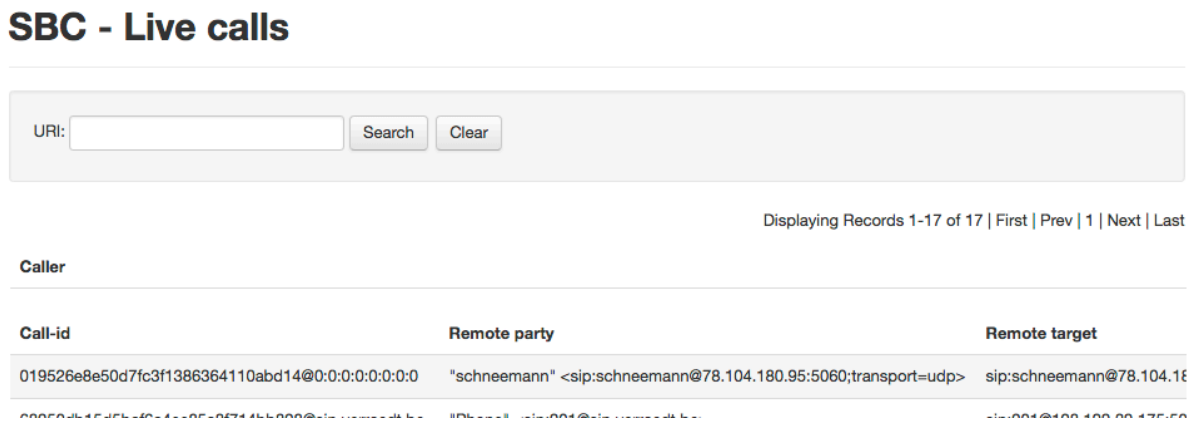

Figure 141: Live Calls

Destination Blacklists

"Monitoring ▸ Destination Blacklists" shows IP addresses that have been found to be unresponsive. See section IP Blacklisting: Adaptive Availability Management to find out how to configure the ABC SBC to handle routing to unresponsive SIP destinations.

The Figure Destination Blacklists shows the user-interface for monitoring the unavailable IP addresses. It shows a single IP address and time-to-live to remain on the availability blacklist.

The TTL field specifies the time interval (in seconds) for which the destination is placed on the blacklist. When this interval passes, the destination is automatically removed from the blacklist. If '-1' value is used for TTL or if "Valid forever" checkbox is checked, the destination is placed on the blacklist forever or until it is removed manually.

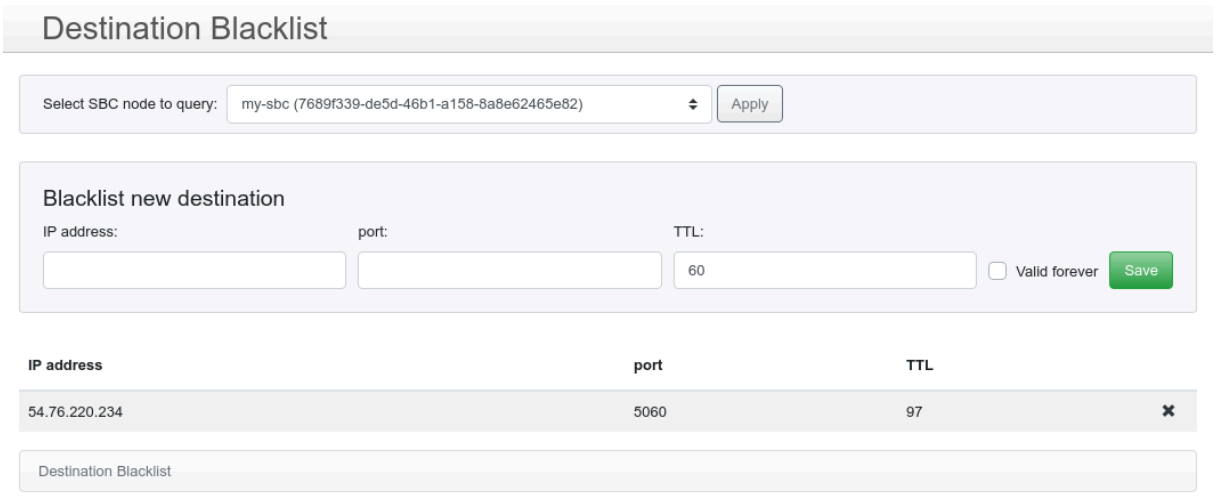

Figure 142: Destination Blacklists

Registration Agents

"Monitoring ▸ Registration Agents" allows monitoring of connections to third-party services where the SBC has registered. The Figure Registration Agents displays the successful registration status. If needed, the administrator can manually trigger re-registration either for all registration agents using the "Re-Register All" button or individually using the "Re-Register" button for a specific registration agent.

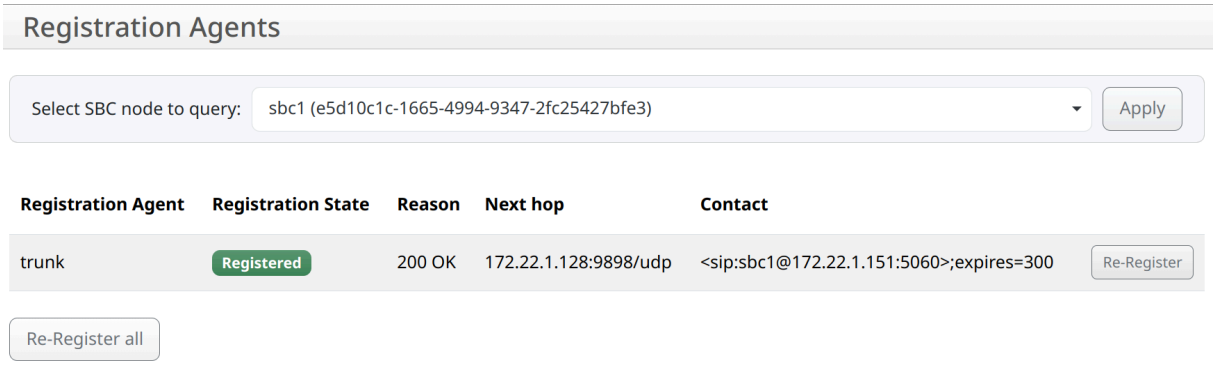

Figure 143: Registration Agents

The following information is displayed for each entry:

- Registration Agent - name of the registration agent.
- Registration State - one of the following values: “NotRegistered”, “Registering”, “Registered”, “Refreshing”, “Deregistering”, “DnsError”, “Error”.
- Reason - SIP response code sent by registrar.
- Next hop - address of the destination where the registration agent is registered to. If the registration agent is not registered, content of this field is undefined.
- Contact - registered contact of the registration agent. If the registration agent is not registered, content of this field is undefined.

Call Agents status

“Monitoring ▸ Call Agents status” allows to monitor availability and blacklisting status of call agents.

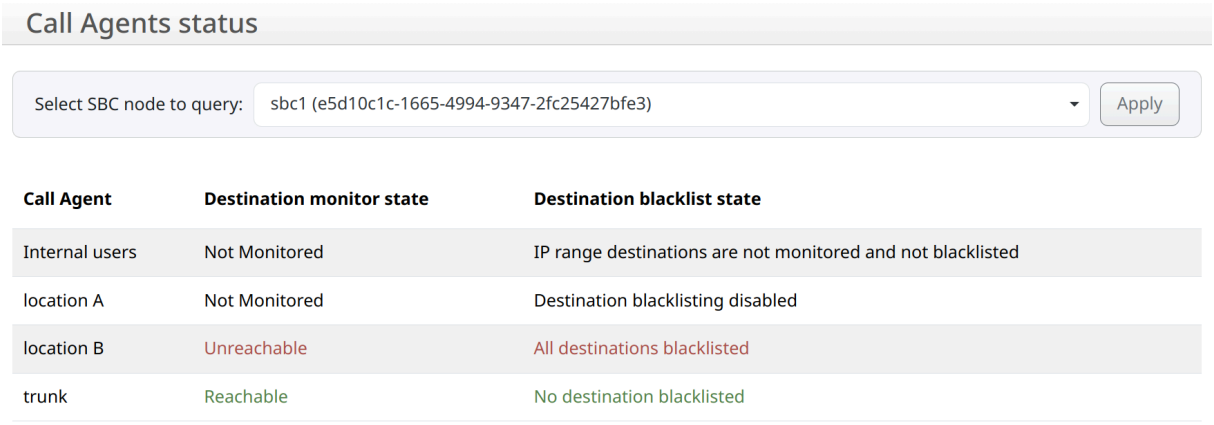

Figure 144: Call Agents status

The *Destination monitor state* shows call agent reachability if destination monitoring is enabled for the appropriate call agent. It can have one of the following values:

- Not Monitored - Destination monitoring is disabled.
- Reachable - All IP addresses associated with a given call agent are reachable.
- Partially reachable - Some IP addresses associated with a given call agent are reachable, while others are not.
- Unreachable - All IP addresses associated with a given call agent are unreachable (request was sent by SEMS but no response arrived in time).
- Sent - No destination is reachable yet, and the system is still processing the status of at least one destination.
- Overloaded - No destinations are reachable. The last destination’s error occurred because the OPTIONS probe received a reply code listed in the “Unavailable on reply code” configuration (503 if the configuration is empty).

- Resolved - For call agents identified by “Domain or Hostname” with destination monitoring disabled, this status indicates that the last destination was successfully resolved via DNS.
- Dns Error - For call agents that are identified by “Domain or hostname”, no destinations are reachable and failed to dns-resolve the last destination.
- Unknown - The initial state before any information about the given call agent is available.
- Error - This state appears when SEMS is unable to send a request due to network level error(s) or signaling interface problem.

The *Destination blacklist state* shows call agent’s destination blacklist status, if enabled for the appropriate call agent. It can have one of the following values:

- Destination blacklist disabled - The destination blacklist is disabled for a given call agent.
- No destination blacklisted - No IP addresses associated with a given call agent are blacklisted.
- Some destinations blacklisted - Some IP addresses associated with a given call agent are blacklisted.
- All destinations blacklisted - All IP addresses associated with a given call agent are blacklisted.

Hovering over the current status with the mouse pointer displays a tooltip with additional information, which IP addresses are reachable or unreachable.

Note: If a call agent is defined by IP range, neither destination monitoring nor destination blacklisting can be configured.

Conference Rooms

“Monitoring ▸ Conference Rooms” displays a list of existing conference rooms, identified by their room ID, along with a list of active participants for each conference room.

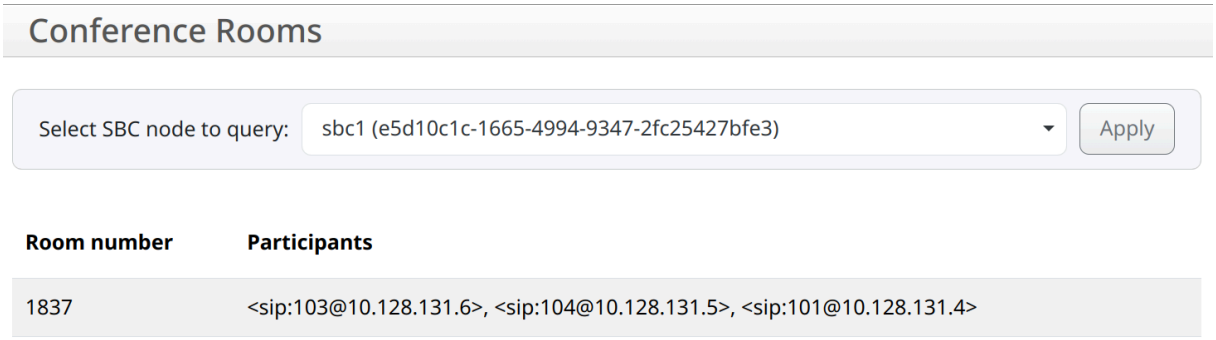

Figure 145: Conference Rooms

System status

“Monitoring ▸ System status” is divided into two sections. The first section, related to the CCM, provides a list of all active warnings and errors related to the CCM itself. The second section focuses on all SBC nodes associated with the CCM. For each node, key information is displayed regarding its configuration status and high availability (HA) status.

CCM related information:

- Time - The timestamp when the problem occurred.
- Severity - The severity level of the problem:
  - WARN - A minor issue has occurred with no immediate impact on functionality or operation. However, administrator should assess the potential risk and resolve the issue if necessary.
  - ERR - A major or critical issue has occurred, which might have significant impact on CCM functionality.
- Message - A description providing more details about the warning or error.
- Component - The service or component affected by the issue.
- Details - Clicking this button opens a window with detailed information about the problem.

Additionally, a “CCM Mgmt console” button is available. Clicking it opens a terminal window, allowing the administrator to directly access the CCM node if needed.

Node related information:

- Node name - By default the hostname of SBC node is used, but this value can be modified.
- Hostname - The hostname of the SBC node.
- Release - The SBC version currently running on a node.
- Configuration - If the node is using the latest configuration, the value is set to “Up to date”. Otherwise, it is marked as “Outdated”. Hovering over this value displays a tooltip with information on which configuration version is being used by the node.

> **Info**
>
> A node may show “Up to date” but still be using a different configuration version than the one indicated by the “Latest config version” label, which is located below the nodes table. This can happen when the latest changes have not affected that node.

- Status: The current status of the node:
  - OK - No issues detected.
  - WARN - A minor issue has occurred with no immediate impact on functionality or operation. However, administrator should assess the potential risk and resolve the issue if necessary.
  - ERR - A major or critical issue has occurred, which might have significant impact on SBC functionality.
- Config pull/push - Indicates whether a config PULL or config PUSH mode is configured for the node.
- Last report - The last time the node reported its status to CCM.
- Last pull - The last time the node pulled a new configuration.
- Signaling up since - The timestamp when the signaling interface became available.
- Master address - The IP address or hostname of the CCM by which the SBC node is managed.
- HA state - If the node is a part of high availability (HA) group, this field can display High Availability statuses. If HA is not configured, this field remains empty.
- Prov tables - Clicking the ‘versions’ button opens a window displaying a list of all provisioned tables on the node, along with their versions.
- Interfaces - Clicking the ‘show’ button opens a window listing all interfaces configured for the node.
- Details - Clicking the ‘details’ button opens a window with detailed status information about the selected node.
- Mgmt console - Clicking the ‘Mgmt console’ button opens a terminal with an SSH connection to the host where the node is running.

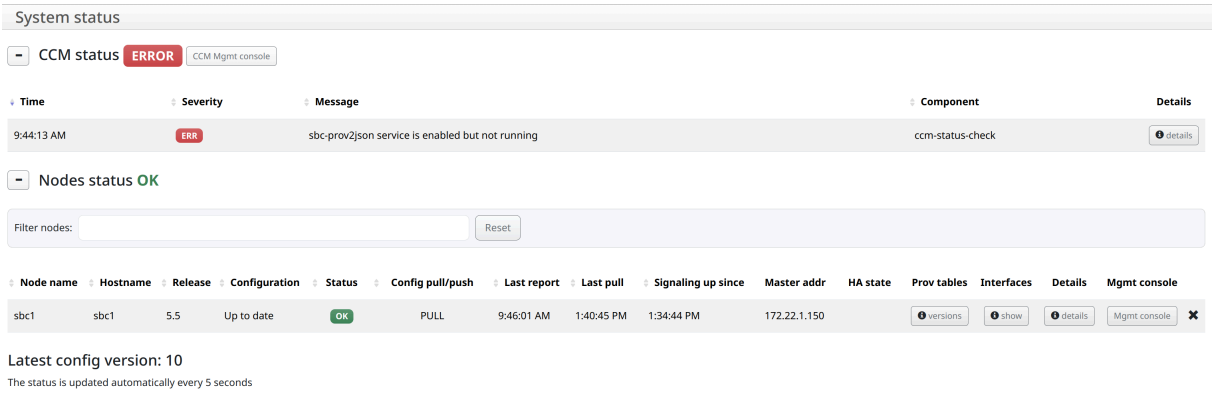

Figure 146: System status

User Recent Traffic

The ABC SBC can be configured to keep track of the most recent SIP traffic. This is particularly useful when a problem is identified which doesn't occur anymore and needs to be troubleshooted retro-actively. Another reason to look-back in this stored traffic is it includes even IP packets that are filtered at IP layer (see Section Police: Devising Security Rules in the ABC SBC) and cannot be troubleshooted at higher layers.

By default, this feature is turned off, but can be turned on by changing the number of PCAP files to keep on "Config ▸ Global Config ▸ Pcaps" page to a non-zero value.

The file size and number of files to keep should be tuned according to available disk space.

These files are rotated: traffic starts to be written into a new file once the desired size of current one is reached. Once the configured number of files is written, writing starts into the first file again.

Note: the files use extensions ".pcapXX", where the "XX" part corresponds to the file number. If the global config option to set number of files to keep is modified, all existing traffic.pcap* files are deleted on the SBC once the configuration change is activated.

The administrative page "Monitoring ▸ User Recent Traffic" allows administrators to retrieve SIP traffic for a specific IP address. Also a secondary IP address can be included in which case packets matching either IP address will be retrieved. The retention policy for the stored traffic can be configured as shown in the Section PCAP Parameters.

To retrieve the SIP traffic, a user must choose the time interval within the available retention period (configuration of which is described in Section PCAP Parameters), IP address, and press the "Get PCAP file" button. Processing can take up to several minutes depending on the time interval chosen. The traffic comes in an archive in PCAP format along with TLS session keys that can be used to decrypt SIP traffic that came over TLS connections.

Wireshark can be used to inspect encrypted TLS traffic using the session keys, see the following link for a detailed HOWTO: https://jimshaver.net/2015/02/11/decrypting-tls-browser-traffic-with-wireshark-the-easy-way/

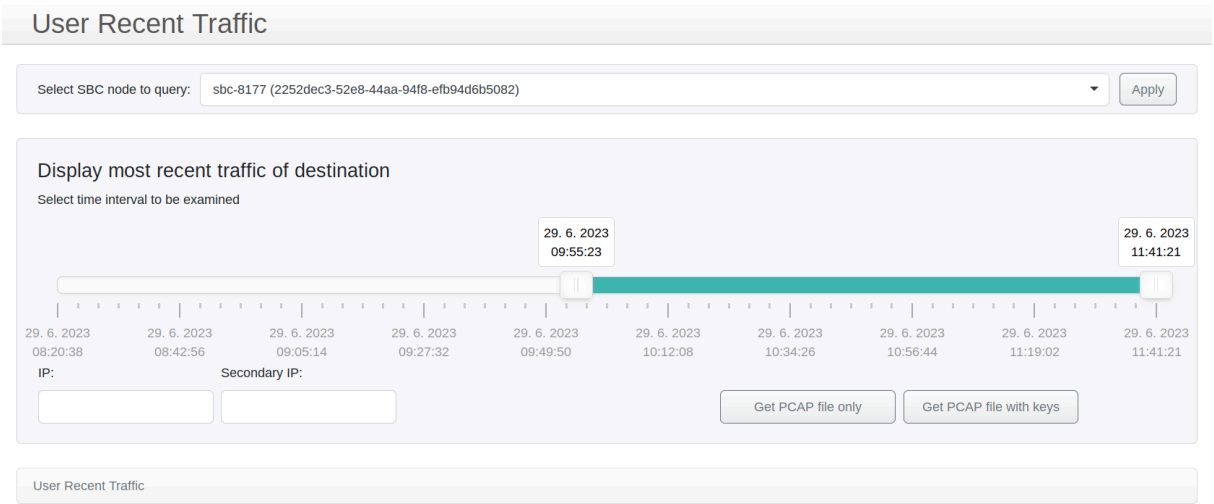

Figure 147: User Recent Traffic

View Logs

"Monitoring ▸ View logs" allows access to ABC SBC logs. The lnav program is used as the log viewer. For details on its control and usage, please refer to its documentation (https://docs.lnav.org/en/latest/).

The log viewer allows access to log files located in the /var/log directory and its subdirectories. The ABC SBC logs are split into following categories:

- container_init - Logs related to container bootstrap initialization.
- ha - Logs specifically related to high availability.
- mgmt - Logs related to periodic jobs (tasker services).
- monitoring - Logs related to `file_uploader` services.
- sbc - The default logging location for all other services.
- status/config.log - Logs from the last config activation.
- status/config.status - Status of the last config activation.
- syslog-ng - Logs specifically related to syslog-ng.

> **Info**
>
> Log files follow the S6 format and rotation conventions. Each logging directory contains a current file with the latest log entries. When logs are rotated, old log entries are archived in a file named @timestamp.s, where timestamp is a TAI64N timestamp representing the absolute time of the rotation.

4.12.3 Measurements and Monitoring

The ABC SBC provides SIP and RTP traffic related counters (general, per Realm/Call Agent and user defined measurements). These measurements are exposed to external monitoring tools such as Prometheus or SNMP.

Starting ABC SBC version 5.5, FRAFOS recommends to use the Prometheus protocol for monitoring. The SNMP is considered deprecated however to ease transition from SNMP to Prometheus, it is still possible to integrate SNMP support with the SNMP daemon running on the host. Please note, all SNMP related global config options are used for ABC SBC nodes up to 5.4 version only.

Prometheus Configuration

In ABC SBC 5.5 the Prometheus pull service became part of *Unified SBC management service*. All necessary configuration is done in Prometheus directly.

- In *targets* use SBC IP address or hostname for which Unified SBC management service is configured. Use `4224` as a port. For example: 192.168.10.20:4224
- As a *metrics_path* use the `/api/v1/metrics/prometheus` value

TLS certificates can be used to secure a connection between Prometheus and SBC node if needed. To do that assign a TLS profile with required certificates to the IMI interface.

Metrics can also be obtained by asking the server directly using *curl* tool. For example: `curl -k https://<SBC_node_IP>:4224/api/v1/metrics/prometheus`. Please note, the port `4224` is controlled by firewall and by default only CCM is allowed to access it.

SNMP Configuration

On the host machine, an SNMP daemon must be configured and the AgentX socket directory provided by the daemon must be mounted into the SBC container.

Using *snmpd*, the following configuration file (*"/etc/snmp/snmpd.conf"*) can be used:

```
agentAddress udp:161,udp6:[::1]:161

# Frafos GmbH Enterprise OID
view   frafosStats  included   .1.3.6.1.4.1.39695

# UCD-SNMP-MIB::prTable
view   frafosStats  included   .1.3.6.1.4.1.2021.2

# NET-SNMP-EXTEND-MIB::nsExtend
view   frafosStats  included   .1.3.6.1.4.1.8072.1.3.2

rocommunity sbc_com_321 0.0.0.0/0 -V frafosStats

master agentx
agentXSocket /var/agentx/master
```

Then, the SBC container must be started with `--mount type=bind,src=/var/agentx,dst=/var/agentx` to mount the AgentX socket directory into the container. The ABC SBC will automatically detect the presence of the AgentX socket and start its agent.

General Prometheus Statistics

General ABC SBC statistics related to calls, transactions and registrations.

Explanation of terms used in following table:

- **rx** - data received by ABC SBC
- **tx** - data sent by ABC SBC

Metric Name	Description
`abcsbc_sessions`	Total number of active sessions. Type: gauge Attributes: status: active
`abcsbc_calls`	Total number of established calls. Type: gauge Attributes: direction: all
`abcsbc_call_starts_total`	Total number of call starts: the metric sums both successful and rejected calls. Type: counter Attributes: direction: all
`abcsbc_bits_total`	Total number of RTP data transferred in bits. Type: counter Attributes: direction: all
`abcsbc_sip_messages_total`	Total number of SIP messages. Type: counter Attributes: direction: rx/tx, type: request/reply
`abcsbc_sip_retransmissions_total`	Total number of SIP re-transmissions. Type: counter Attributes: direction: tx, type: request/reply

Metric Name	Description
abcsbc_transactions	Number of ongoing UAC/UAS transactions. Type: gauge Attributes: type: uac/uas
abcsbc_open_ports	Number of media ports used. Type: gauge Attributes: type: media
abcsbc_high_availability_status	HA node status • 0: Unknown • 1: BACKUP • 2: MASTER • 3: FAULT • 4: STOP For more information see High Availability statuses. Type: gauge Attributes: uuid: ID of HA group, name: HA group name
abcsbc_register_cache	Number of SIP registrations. Type: gauge Attributes: status: active

Table 16: General metrics

Prometheus Statistics per Realm / Call Agent

Measurements listed in tables below are counted for each Realm and Call Agent separately.

Explanation of terms used in following table:
- **uuid** - unique identifier either of realm or call agent
- **realm name** - Realm name which a Call Agent belongs to (shown for call agent only)
- **rx** - data received by realm or call agent
- **tx** - data sent by realm or call agent

Metric Name	Description
abcsbc_realm_calls	Number of calls established to/from call agents belonging to the appropriate realm. Type: gauge Attributes: direction: rx/tx, uuid: realm's UUID, name: realm's name
abcsbc_realm_call_starts_total	Total number of call starts to/from realm. This metric sums both successful and rejected calls per realm. Type: counter Attributes: direction: rx/tx, uuid: realm's UUID, name: realm's name

Metric Name	Description
abcsbc_realm_bits_total	Total number of RTP data in bits relayed to/from realm. Type: counter Attributes: direction: rx/tx, uuid: realm's UUID, name: realm's name

Table 17: Realm metrics

Metric Name	Description
abcsbc_call_agent_calls	Number of calls established to/from a call agent. Type: gauge Attributes: direction: rx/tx, uuid: CA's UUID, name: CA's name, realm: CA's realm
abcsbc_call_agent_call_starts_total	Total number of call attempts to/from call agent. This metric sums both successful and rejected calls per call agent. Type: counter Attributes: direction: rx/tx, uuid: CA's UUID, name: CA's name, realm: CA's realm
abcsbc_call_agent_bits_total	Total number of RTP data transferred in bits to/from call agent. Type: counter Attributes: direction: rx/tx, uuid: CA's UUID, name: CA's name, realm: CA's realm
abcsbc_call_agent_status	Call agents status: • 0: Unknown • 1: Sent • 2: Resolved • 3: Reachable • 4: Unreachable • 5: Partially Reachable • 6: Overloaded • 7: DNS Error • 8: Error For more information see Call Agents status Type: gauge Attributes: destination: IP or hostname, uuid: CA's UUID, name: CA's name, realm: CA's realm

Table 18: Call Agent metrics

Prometheus Statistics per Interfaces

These measurements are exported from the transport layer.

Explanation of terms used in following table:
- **rx** - data received by ABC SBC
- **tx** - data sent by ABC SBC

Metric Name	Description
`abcsbc_interface_sip_messages_total`	Total number of SIP messages per interface. Type: counter Attributes: interface: name, direction: rx/tx, type: request/reply
`abcsbc_interface_sip_retransmissions_total`	Total number of SIP re-transmissions per interface. Type: counter Attributes: interface: name, direction: tx, type: request/reply
`abcsbc_register_cache_registrations_per_interface`	Register cache per interface statistics. Type: gauge Attributes: interface: name, endpoint: node IP:port

Table 19: Interface metrics

Prometheus Statistics from HA redis

These measurements are exported from the redis instance used for storing call and registration state. The payload will change depending on the HA status of the node.

For more explanation about the ways the values are calculated, please refer to redis documentation about replication and for `INFO replication` command.

On a master node, the following metrics are available:

Metric Name	Attributes	Type	Description
redis_replication_connected_replicas		gauge	Number of connected replicas (should be >= 1 for healthy replication).
redis_replication_replica_offset	replica: ID	gauge	Replication offset for each replica. Where *N* is the replica index, starting at *0*.
redis_replication_replica_lag_seconds	replica: ID	gauge	Replication lag in seconds for each replica. Where *N* is the replica index, starting at *0*.
redis_replication_master_repl_offset		gauge	Master replication offset.

Table 20: redis replica metrics on HA master node

On a backup node, the following metrics are available:

Metric Name	Type	Description
redis_replication_master_link_up	gauge	Master connection state (1 = up, 0 = down).
redis_replication_master_link_down_seconds	gauge	Seconds since master link went down (-1 if up).
redis_replication_replica_repl_offset	gauge	Current replication offset on this replica.
redis_replication_master_sync_in_progress	gauge	Sync in progress indicator (1 = yes, 0 = no).

Table 21: redis replica metrics on HA backup node

User Defined Counters

User defined counters can be created and increased using an **"Increment custom counter"** action configured in inbound or outbound rules. This action increments a user-defined custom counter by a given value. As parameters the counter name and the counter increment are given, see Fig. User defined counters.

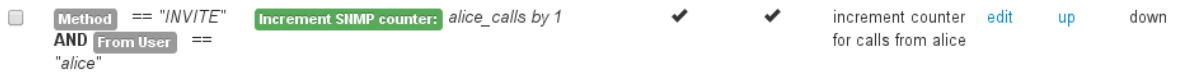

Figure 148: User defined counters

Metric Name	Description
`abcsbc_custom_counters`	Custom counter defined in GUI. Type: gauge Attributes: name: variable name

Table 22: General metrics

4.12.4 Additional Sources of Diagnostics Information

The following additional sources of management data may be also used:

- traffic monitoring and event tracking described in Section Overview of Monitoring and Troubleshooting Techniques,
- remote monitoring described in Section Measurements and Monitoring,
- logging concealed with call log in file as described in Section Monitoring and Logging.

When trying to understand some unexpected network or SBC behavior the following facilities can be also helpful:

- Audio can be recorded as described in Section Audio Recording.
- CDRs, as described in Section Call Data Records (CDRs), include useful information.
- The ABC SBC can be configured to send notification by email if some serious error such as exhausted disk space occurs. Configure the recipient email address under "Config▸Global Config▸Monitoring▸Email for sending alerts". Configure SMTP server to which the emails will be passed under "Config▸Global Config▸Monitoring▸Mailserver for sending alerts" to use external mail server. Configure various thresholds for alerts based on high system load, memory used, CPU waiting percentage and disk usage percentage under "Config▸Global Config▸Monitoring".

4.12.5 Viewing ABC SBC Logs

The GUI screen "Monitoring->View logs" allow access to ABC SBC logs. We use `lnav` program as the log viewer so for further details about its control, please check its documentation (https://docs.lnav.org/en/latest/).

By default all the rotated log files like for example: `syslog`, `syslog.1` and `syslog.2.gz` are displayed as single entry in the list of log files and are displayed together by the `lnav` viewer. Unchecking the `Join rotated log files` checkbox allows to display such log files separately.

4.12.6 Coredumps

It may happen that a process does not operate properly and is terminated by signal, that may cause a "coredump" to be generated. These coredumps are valuable for further problem debugging and might be asked by FRAFOS support to be able to properly investigate and fix the issue.

In previous ABC SBC versions generating coredumps for the most critical process - SEMS - was allowed by default but with containers it relies on proper host configuration that can not be influenced from the container itself.

With nowadays ABC SBC, an "alert" event is generated when SEMS process crashes regardless of the coredump settings, so the administrator is informed about the problem and may react appropriately.

Please note, that the process coredumps may be huge and writing them may significantly prolong the time necessary to restart a process upon a crash and thus service downtime might be increased.

Additionally, they consume a lot of space so it might be necessary to monitor HDD space of the destination used for storing them and possibly clean that storage up when necessary.

To allow a process in container to dump a core it is recommended to install `systemd-coredump` package on the host OS (Debian based Linux distribution) or its equivalent:

```
% apt install systemd-coredump
```

This service is responsible for managing coredumps generated on the host and in containers running there and can be configured (see `man coredump.conf`) to fulfill the particular deployment needs.

The `coredumpctl` utility contained in the mentioned package can be used to list:

```
% coredumpctl list -r
TIME                            PID UID GID SIG    COREFILE EXE            SIZE
Fri 2023-06-16 12:18:12 CEST  81029   0   0 SIGILL present  /usr/sbin/sems 1.2M
Fri 2023-06-16 11:46:04 CEST 257465   0   0 SIGILL present  /usr/sbin/sems 1.2M
Fri 2023-06-16 11:44:18 CEST 257212   0   0 SIGILL present  /usr/sbin/sems 1.2M
Fri 2023-06-16 11:42:41 CEST 105270   0   0 SIGILL present  /usr/sbin/sems 1.2M
```

and export particular coredumps:

```
% coredumpctl dump 257212 | gzip > core.gz
           PID: 257212 (sems)
           UID: 0 (root)
           GID: 0 (root)
        Signal: 4 (ILL)
     Timestamp: Fri 2023-06-16 11:44:17 CEST (2 days ago)
  Command Line: /usr/sbin/sems -P /var/run/sems/sems.pid -f /etc/sems/sems.conf
    Executable: /usr/sbin/sems
 Control Group: /machine.slice/systemd-nspawn@sbc-5.3.0.service/payload/system.slice/sems.service
          Unit: systemd-nspawn@sbc-5.3.0.service
         Slice: machine.slice
       Boot ID: 1e74911d896440818965717facd36aa4
    Machine ID: 16a8ad16f7004e0eac68aded464561d9
      Hostname: sbc
       Storage: /var/lib/systemd/coredump/
core.sems.0.1e74911d896440818965717facd36aa4.257212.1686908657000000.zst (present)
  Size on Disk: 1.2M
       Message: Process 257212 (sems) of user 0 dumped core.

                Stack trace of thread 77200:
                #0  0x00007f162dea4d36 n/a (libc.so.6 + 0x85d36)
                #1  0x00007f162dea73f8 pthread_cond_wait (libc.so.6 + 0x883f8)
                #2  0x000055a9accc4d2b n/a (/usr/sbin/sems + 0xd6d2b)
                ELF object binary architecture: AMD x86-64
```

Please note, that the “Enable coredumps” option on SEMS tab in Global config settings, that was used to allow coredumps for SEMS process running directly on the host, is not used any more with ABC SBC 5.3 and higher and is present just for compatibility with older ABC SBC versions.

4.13 Securing SIP Networks using ABC SBC and ABC Monitor (optional)

4.13.1 SIP Security Principles: Collect, Analyze and Police

Like any other Internet-based service VoIP servers can be target of fraud attempts, denial of service attacks and abnormal operational conditions such as registration storms after recovery of a failed router with a user population behind it. These have become common with prevalence of the SIP technology for telephony and need to be dealt with on a daily-basis. A key function of the SBC is to fend off such situations so that the infrastructure behind the SBC and service for the end-users remains unaffected.

Administering any service securely always consists of three steps: **collect data**, **analyze it** and **police**. Each of these steps is a necessity and always requires human judgment of an administrator. This can be challenging with the sheer amount of data to be handled and may resemble looking for the proverbial needle in the haystack. Every day a public SIP service for one thousand subscribers generates as much as 7 GB in 13 millions SIP packets! Obviously making an administrator look at every single SIP packet is not feasible and the ABC SBC FRAFOS solution comes therefore with many administrative aids.

The first two steps, gathering data and analyzing it, can be purchased using various tools. FRAFOS however strongly recommends use of its ABC Monitor because it has a unique access to internals of the ABC SBC and can report on many specifics not seen outside of it. The data gathered by the ABC Monitor known as **events** come from inside the ABC SBC and can therefore reveal information not visible to anyone else: plain-text signaling and media which is encrypted to the outside (always the case with WebRTC), internal information such as reasons why a specific SIP request has been dropped, or correlation of dialogs that are obfuscated to the outside using Topology Hiding.

The following real-world example shows a typical attack on a SIP service. The Figure Screenshot of a monitored password-guessing attack shows the course of the attack and defense against it. The attack

started on March 22, 2016 at 1AM local time from an IP address located in Guangzhou, China. It consisted of attempts to register as users with numbers beginning with "122". The site was initially not taking any effort to fend the attack off, resulting in 1000 authentication attempts per hour. While the attacker didn't succeed in registering a URI protected using well-chosen passwords in this case, endurance or a weak password could have crowned his undertaking with success. Therefore at about 10:30 local time, the administrator took an action and locked out the attacker's IP address. It took exactly one hour until the perpetrator realized his tool was receiving no responses back and started sending from a different IP address. Now the SIP service administrator found out that a static policy is not good enough and enabled a dynamic policy that locks an IP address if too many failed authentication attempts come from it. The effect came instantly: the attacks were locked at transport layer and began to appear only shortly in hourly interval: that's the period after which the attacker changed his source IP address – few attempts were then observed in the ABC Monitor until the new source IP address was banned again.

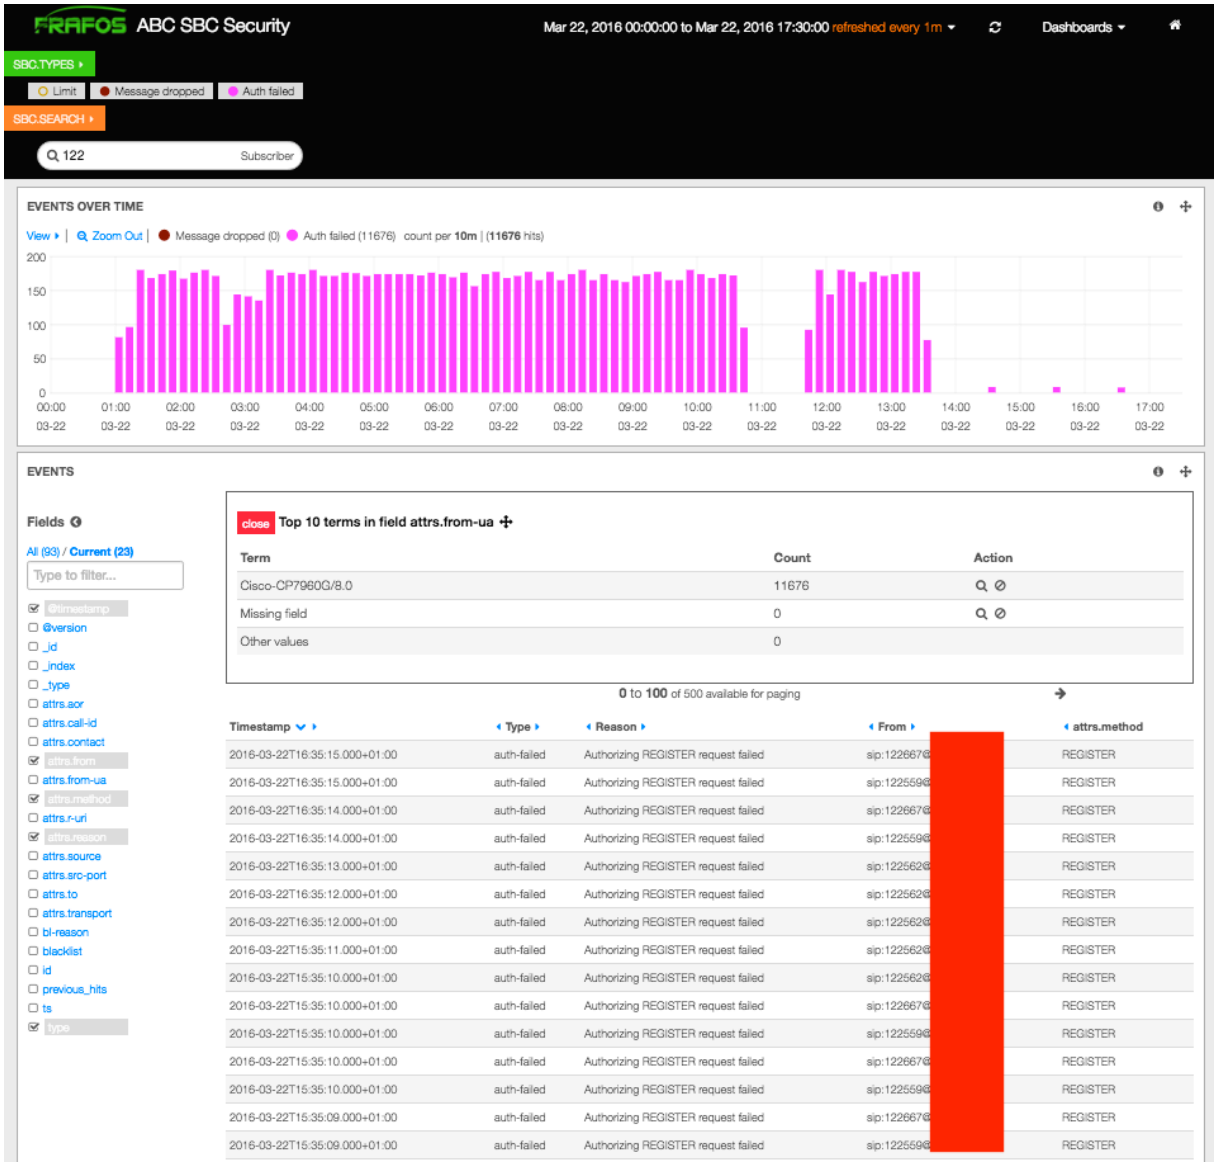

Figure 149: Screenshot of a monitored password-guessing attack

This example is re-iterating the importance of the fundamental security principle: collect, analyze and police. If the site administrator didn't have good data about what's going on, he would be literally blind and the authentication attack could have remained unnoticed. All in all, two requests per second is not an excessive amount of traffic on a multi-thousand user-site, and the way the SIP protocol is designed almost every SIP request causes a 401/407 authentication challenge. Which leads to the second, analytical point.

Usage data needs to be analyzed efficiently. The administrator needs to find out if there is anything going on at all, what are the specific patterns of an attack that can be used to fend it off, and who is the originator. The last step, fending the attack off, is the easiest once the nature of an attack is known.

These three facets of the security life-cycle are documented in the following sections. We will discuss them in the order a SIP packet encounters on its way. The first thing that happens to a freshly arrived SIP packet is it is processed by a ruleset that represent a site's security policy. We describe the available rules and practices for using them in the Section Police: Devising Security Rules in the ABC SBC.

4.13.2 Police: Devising Security Rules in the ABC SBC

There is nothing more dangerous than security. Sir Francis Walsingham, Queen Elizabeth's Principal Secretary

The objective of the policing functionality is simple to state: Filter unwanted traffic as soon as possible before it causes harm. In order to achieve this objective, reasonable policies must be administered which permit legitimate and drop harmful traffic.

The delicate challenge is to differentiate between "friend and foe". Resolving this dilemma often requires a learning period – the administrator or an automated system on his behalf need to find out the presence of illegitimate traffic and its originator. An administrator can do this by analyzing traffic. The advantage of this approach is that human assessment of the situation can capture finesses a computer fails to see. This argument is for example the reason why air traffic control has never been fully automated – computers are still not trusted a judgment about abnormal situation.

The disadvantage of relying on humans is, not surprisingly, the human factor too. Humans may fail to see an abnormality in sheer amount of traffic and keep alert 24 hours a day. That's what computers are good at: they can look over gigabytes of traffic relentlessly, find patterns they have been taught to look after, and raise alarms any time of day as soon as they appear.

Therefore we at FRAFOS suggest that highest level of security of a SIP service is given when automated traffic filtering is combined with computer-aided human judgment.

In the following list we show typical attack types and also ABC SBC policies to deal with these.

- **Intrusion attacks** are attempts to obtain unauthorized access to a system or to a SIP user's account. They come by nature as an uninvited surprise at the most inconvenient time. The challenge is therefore to counter them as quickly as possible. In the Section Automatic IP Address Blocking we are showing how to **automate prohibition of malicious traffic** even before administrators do notice.
- **Harassing traffic** may be easier to detect and yet inconvenient to deal with. Unlike with real attacks, the harassing traffic is mostly an unintended side-effect of a broken implementation or configuration of some SIP devices. It doesn't try to masquerade or surprise yet if coming in large quantities, it may have the same devastating effect as a malicious attack. The capability to filter out such "noise" helps to reduce security risk, off-load the infrastructure, and focus on the traffic that matters. We show **how to block well-known sources of harassing traffic** at both IP and SIP layer in the section Manual SIP Traffic Blocking.
- **Unprivileged traffic** is traffic that does not appear harmful yet it has not been explicitly authorized to use a SIP service. Such may not appear harmful on the first sight, yet it may be also an initial probing prelude to an actual intrusion attack. It appears therefore a wise idea to drop traffic which does not demonstrate appropriate credibility before it turns into a harm. This way **users exhibiting proper behavior are prioritized** over users that don't. The simplest and yet most powerful credibility test is that of successfully completed SIP registrations. See Section Blocking a User by his Registration Status for guidelines how to use it. Also note that the credibility-test is extended to lower-layer by a generalized technique known as **grey-listing** (Section Automatic Proactive Blocking: Greylisting).
- **Excessive traffic** may have many root causes: Denial of Service (DoS), breach of service-level agreements, or SIP network misconfiguration. Regardless of the cause the results are always the same:

quality of service (QoS) declines for legitimate uses. To prevent such QoS impairments, a site better chooses to set limits on SIP and or RTP traffic and drops traffic exceeding the limits. We show **how to shape traffic** in Section Traffic Limiting and Shaping. Traffic shaping is also important to discover some sort of attacks like SIP password guessing: if the attacking SIP device tries to masquerade as a legitimate user, the high signaling rate it needs for guessing will give it away.

- **Excessively long calls** are another irritating phenomena that needs to be dealt with in order to reduce a high-charge risk. Most often it is caused by SIP devices that do not terminate calls properly. Fraud attempts are also known that have been trying to gain maximum by running calls as long as possible. In Section Call Duration Control we explain **how to keep a SIP service robust against infinite calls**.
- **Improper content** in SIP signaling or SDP media can bring insufficiently robust SIP devices to failure. This situation doesn't happen so often because SIP devices typically do not have such processing capability like general purpose computers to be a real magnet for all kinds of viruses. Yet the situation changes as Android telephones come on the market and features offered by servers expand. Academics have already described SQL injection attacks[5][6] : They crafted SIP messages which included SQL commands, and the SIP servers passed these to backend software. When the software is not sufficiently robust, opening a web page to see a list of completed calls will also launch a potentially dangerous SQL query. If content of SIP and SDP is considered a risk, more aggressive mediation is needed. See Sections SIP Mediation and SDP Mediation for more information **how to filter SIP/SDP content**. Particularly header-field whitelisting may be instrumental for this purpose.

The ABC SBC offers several instruments for filtering undesired traffic. There are two types of: filters operating at IP/transport layer for the highest performance and filters operating at SIP layer when more sophisticated filtering criteria are needed. For example a well-known flooding attacker is best eliminated by filtering out all traffic from his IP address. On the other hand, if a single SIP user behind a SIP trunk IP misbehaves, blocking the whole trunk IP would be throwing the baby out with the bathwater. In such a case, the SIP-layer filtering would be a safer choice, albeit not that fast.

The IP-layer rules are managed from the administrative menu under **"System ▸ Firewall"**. The screen offers a search box where one can look for an IP address to see if it is present on any of the lists, and then several firewall rule lists. The lists are ordered by precedence: the top lists are more manual, have higher precedence and override the bottom-placed lists. The top lists include Manual low-level rules, Exceptions to automatic blacklists, and Manual Firewall Blacklists and are described in Section Manual IP-layer Blocking. The bottom lists are generated in an automated way using built-in security assessment algorithms without administrator's intervention, can be overridden by the manual lists in the top, and are described in the Sections Automatic IP Address Blocking and Automatic Proactive Blocking: Greylisting.

[5]Geneiatakis, Dimitris, et al. "SIP message tampering: the SQL code injection attack." Proceedings of 13th International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2005), Split, Croatia. 2005.

[6]Abdelnur, Humberto, and Olivier Festor. "Advanced fuzzing in the VoIP space." Journal in Computer Virology 6.1 (2010): 57-64.

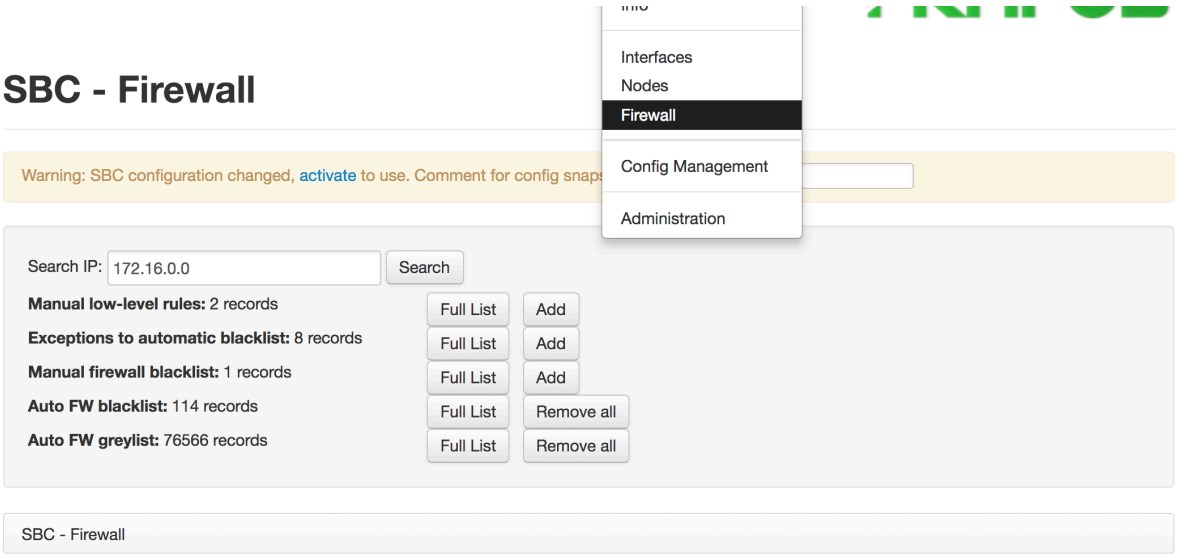

Figure 150: Firewall Rules Management

SIP layer filtering is then described in Section Manual SIP Traffic Blocking. If binary yes/no policies seem too harsh, placing quota on the traffic may be a better answer, which is described in Section Traffic Limiting and Shaping.

Manual IP-layer Blocking

In some situations, e.g. if DOS attacks are encountered, incoming IP traffic may better be blocked already on the operating system firewall (nftables) level so that CPU processing power and memory is saved as the SBC processes don't need to handle the traffic.

The ABC SBC offers a graphical user interface to configure the firewall rules under **"System ▸ Firewall"**. There are several rules list, the top-positioned rules list take precedence over the bottom rules list and are processed in the exactly same order as shown in the GUI.

If incoming packets do not match any of these rules, default rules apply. Traffic to signaling and media interfaces will be accepted if in the declared destination port range, traffic to administrative port numbers will be permitted on XMI and IMI interfaces, all other traffic will be dropped.

The top-most rules list is "Manual low-level rules" and it is a "swiss army knife" for firewall administrators. While it is the first-in-order list, we recommend to use it as the last resort due to extra complexity. Simpler rules such as "Exceptions" and "Manual blacklists" bellow are easier to manage and audit. Nevertheless the low-level rules my be still useful in situations when administrators wish to limit administrative access to well-known IP addresses or permit additional administrative protocols. These rules allow to specify IP flows using source and destination address and port numbers, and whether these flows should be accepted or dropped. That also means that attention must be paid to the order of these rules because it does affect the result. For example, the administrator can use the low-level rules block all traffic coming from the RFC1918 private IP address space as shown in Figure Manual low-level Firewall Rules. When a filtering criteria such as IP address or port number is left blank in the rule, any value in incoming IP packet matches.

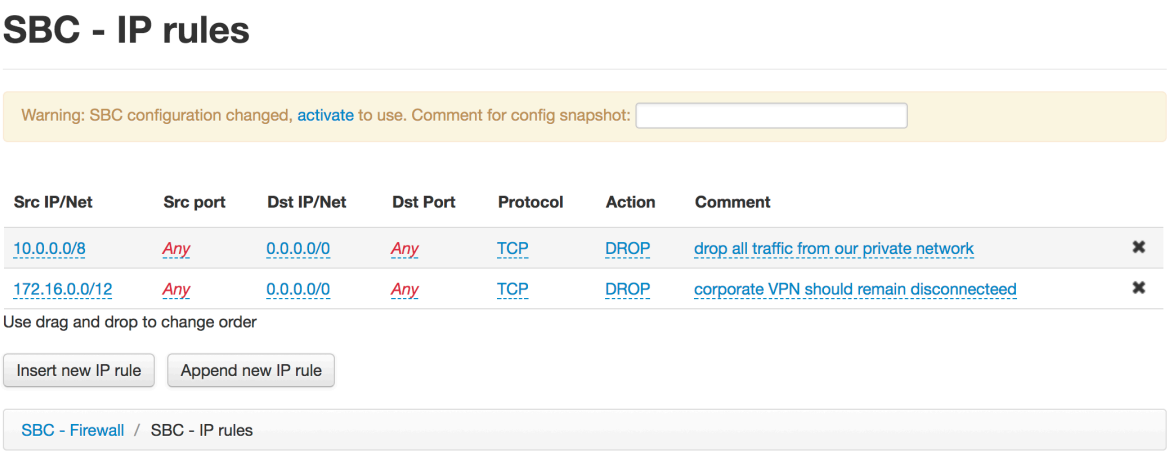

Figure 151: Manual low-level Firewall Rules

The remaining firewall rules only refer to signaling (SIP and Websocket) interfaces and are simple unordered lists of IP and subnet addresses.

"Exceptions to the automatic blacklist" are second in order and could also be called "Whitelists". They take precedences over any of the blacklists bellow. This is important to be able to override too zealous behavior of automated blacklists. This is often the case when traffic of multiple users is coming from behind a single IP address due to NATs or a peering topology. Then the automatic blacklists triggered by a single user would block all others behind their shared IP address. Similarly a SIP site administrator may want to exempt himself from being auto-blacklisted, because his signaling tests may get him blacklisted. Consequently, he would not be even able to open an SSH session to the ABC SBC.

For example a single misbehaving URI would otherwise block an IP address and all other URIs behind it. In such a case, it makes sense to exempt this address from automated blacklisting and address the problematic URI traffic at SIP layer. Example of such a "Whitelist" is shown in Figure Exceptions to the automatic blacklist.

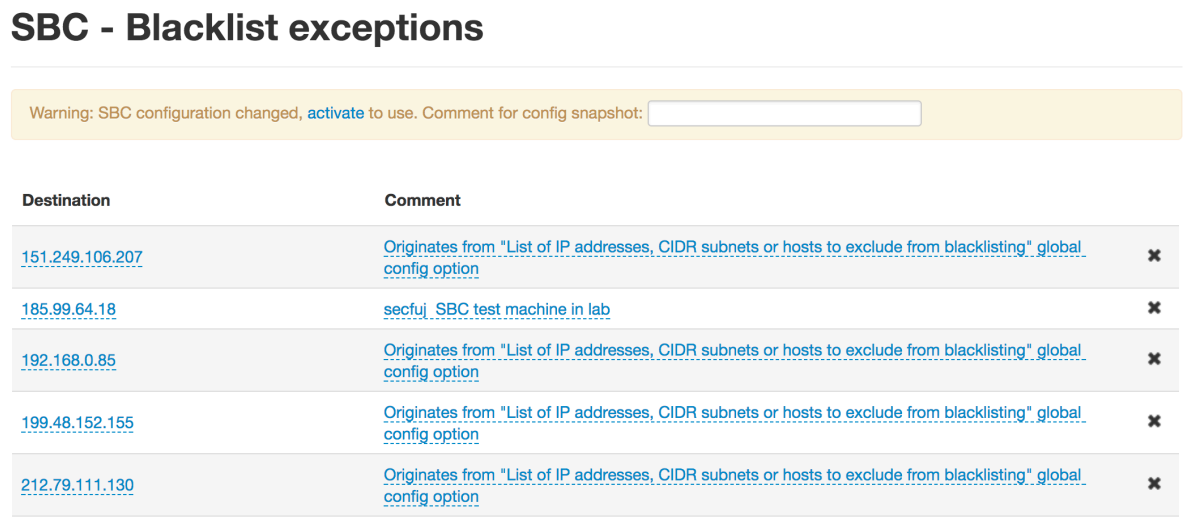

Figure 152: Exceptions to the automatic blacklist

The third in order before automatic rules is "Manual Firewall Blacklist" that can disable traffic from an IP address or subnet even before it reaches any kind of SIP processing logic. This may make sense when a DoS attacker is detected whose traffic is better disabled as early as possible. Example of such is shown in Figure Manual Firewall Blacklist.

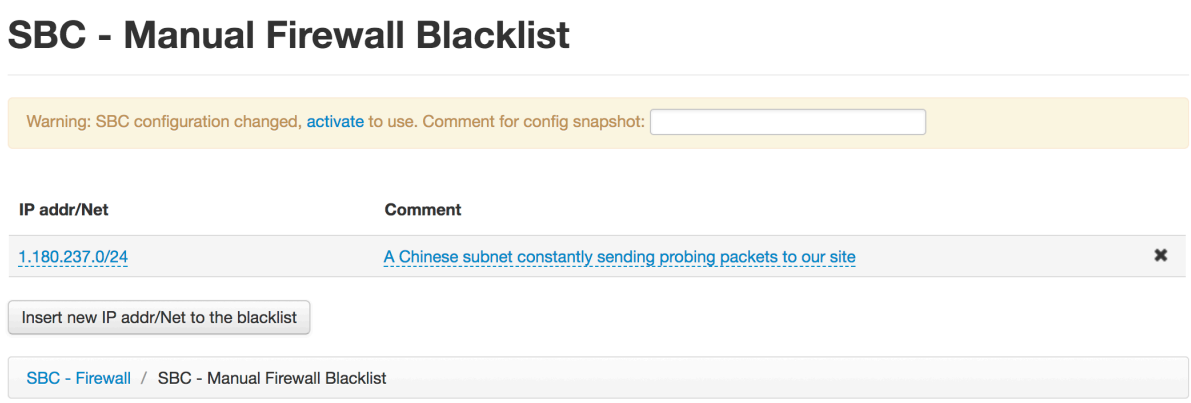

Figure 153: Manual Firewall Blacklist

The next firewall lists, automatic blacklist and greylist, are populated in an automatic way by ABC SBC, and can only be flushed by administrator. The are described in the subsequent chapters Automatic IP Address Blocking and Automatic Proactive Blocking: Greylisting.

Automatic IP Address Blocking

The ABC SBC implements an automated protection process for SIP-layer close-to-real-time detection and IP-layer elimination of offending SIP traffic. This combination provides application-aware assessment with lower-layer performance and helps to eliminate offending traffic without manual administrator intervention. This level of automation cuts the detection-reaction time to almost real-time reactiveness.

A picture tells more than thousand words: The Figure Number of Events with and without Automatic IP Address Blocking shows the profound effect of automated blocking. The event timeline begins under protection of automated blocking in a calm way with about fifty events a minute. When at 23:30 the administrator turns off the automatic protection, the offending traffic finds its way and builds up rapidly. One hour later, 2500 events are already reported every single minute, most of them failed authentication. This unfavorable status remains until the protection is re-enabled. Then, it takes less than five minutes until order is restored again.

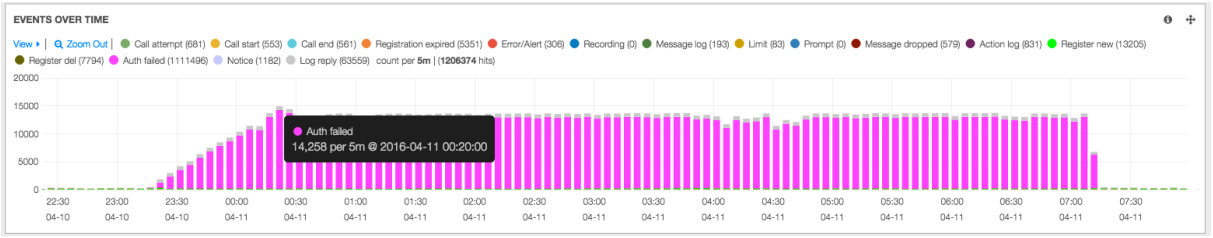

Figure 154: Number of Events with and without Automatic IP Address Blocking

Intrusion attacks, by their very definition and purpose, come uninvited. Sometimes they may try to masquerade themselves in a way that the offending traffic looks innocent: Low-pace, using names of legitimate SIP device types. A human reaction may be too slow to identify such an attack. Therefore the automated process comes in: It acts before a human administrator could.

The ABC SBC protection process is based on the following empirical observations: Offending traffic comes in abnormal quantities, which are indicated by repetitive failures, these failures are linked to an IP address, and the IP address can be blocked. In other words, when some of the security-related events come repeatedly from the same source, it is as good as certain we are dealing with an attack and need to isolate that.

Linking repetitive failures with an offending source is a quite reliable assumption. Singular failures do occur, for example if a softphone user types in a wrong SIP password an authentication failure event is

reported. Yet if the same event is repeated many times, the more likely explanation is we have encountered a password-cracking attack. Leaving such an attack unattended creates a ticket for troubles. At the pace of 2800 authentication attempts per minute (45 per second) shown in our example, an attacker could crack a trivial password taken from Oxford Advanced Learner's Dictionary (185,000 entries) in less than 70 minutes!

Similarly, when a source continues to exceed traffic limits we are dealing with a Denial of Service attack, and when the ABC SBC is receiving repeatedly 403s from a downstream SIP service we know we are dealing with a scanning attack in which an attacker is trying to find a gap in a dial-out authorization policy.

In such a situation banning the originating source address at the OS layer is the safest way to keep the attack from the infrastructure. Care needs to be applied if in the network topology multiple users exist behind a single IP address: then legitimate users could be banned as well as the actual offender. This could be for example the case with peering traffic from behind a SIP proxy, or multiple users behind a single NAT.

The immediate effect of automated IP Address Blocking can be seen in Figure Number of Events with and without Automatic IP Address Blocking: At the very moment when it is enabled, the storm of authentication attempts calms down. It continues to appear briefly when either the attacker changes his IP address or the maximum "banning time" expires – then the detection mechanism strikes in again and the attacks vanish.

Once a source IP address is detected as a repeating offender, all of its traffic will be silently dropped. The list of all currently banned IP addresses can be found in the menu under **"System ▸ Firewall ▸ Auto FW blacklist ▸ Full List "** together with the remaining time they are supposed to spend on the list.

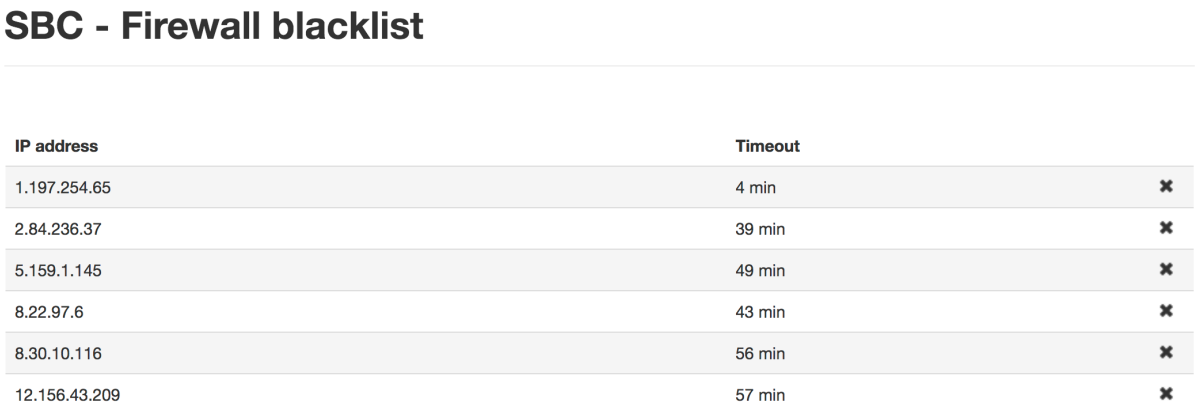

Figure 155: Automated Firewall Blacklist

Note: there is also another list under **"System ▸ Firewall ▸ External FW blacklist "**. This is intended to allow blacklisting IP addresses provided to Sbc from external source and is not used by default.

Scoring system

This effect is achieved by ABC SBC monitoring various occurrences that add to a "score" of a potential offender. To be banned, traffic of an offender must induce several serious events within a pre-configured period of time. Once the score is high enough to identify the originating IP address as "serial offender", the address is put on a blocking list and stays there for a pre-configured time.

The events that add to the score are:

- *limit* for excessive traffic,
- *message-dropped* for messages that the administrator chose to drop using the *drop* action,
- *auth-failed* for failed authentication attempts,

- *log-reply* for transactions which were declined by a downstream SIP entity,
- and significant errors to pass SIP compliance sanity checks.

SIP compliance sanity checks include:
- Request sequence number violation (based on CSeq checking).
- Request parsing errors:
  - malformed first line,
  - missing `Via`, `CSeq`, `From`, `To` or `Call-ID` header field,
  - Unparsable `Via`, `CSeq`, `From`, `To`, `Call-ID` or `RAck` (if included) header field.

Please note: sanity checks errors do not trigger any event.

Each of these events count as **1 offense**, with a **negative score of 1**.

The scoring system is implemented like a leaky bucket into which water is poured regularly. Once the bucket is empty, the offending IP is blacklisted:
- each new IP address starts with a bucket filled with a certain amount of water in it (**start score**).
- each offense decreases that score by 1.
- for every second passed, some water is poured into the bucket (**time bonus**).

If the start score is not considered, a certain IP is allowed `time bonus x time` offenses per `time`. For example, if the time bonus is set to `0.0001`, this means that `0.0001 x 3600 = 0.36` offense are allowed per hour. With `0.005`, this raises to `0.005 x 3600 = 18` offenses per hour, or `0.3` offenses per minutes.

The start score raises the score at the beginning so that the first offense does not cause blacklisting immediately (except if a huge time bonus is setup, which is not recommended), so that in normal cases it should be set to a value greater than 1.

Once an IP has been blacklisted, and the blacklisting expired, the score starts fresh as for a new IP.

If no offense has been registered in a certain amount of time (**time to remove entries**), the IP record is deleted, so that the next offense for that IP will reset the score to its starting value.

Given these settings, different strategies can be implemented:
- **trust strangers**: this strategy starts with a high start score (`> 5`), but won't allow any other offenses after that by using a very low or 0 time bonus.
- **forgiver**: the forgiver will forget about IPs that show a good conduct very fast (`< 300s`).
- **close watch**: the close watch will not allow much from the beginning (start score low; `~ 1`), allows an offense every now and then (time bonus `~ 0.0005 / s = 1.8 / hour`) and takes a long time to forget (time to remove entries `> 3600`).

Please note that these strategies can be combined together to allow for proper functionality without letting bad behavior slip through.

Setting up automatic blacklisting

Automated blacklisting is turned off by default. To enable it perform the following steps:
- Turn it on. Under **"Config ▸ Global Config ▸ Firewall"** turn on **"Blacklist IP addr for repeated signaling failures"**. This will enable the automated blocking process that will process the "score" for IP addresses.
- Fine-tune it if necessary:
  - The options **"Signaling failures blacklist: IP address start score before any offense"** (recommended value: 2.8) and **"Signaling failures blacklist: rate per second used to calculate a time-related bonus between offenses"** (recommended value: 0.0005) in the same global configuration section allow to specify a threshold. When exceeded, the offending IP address will be blacklisted. The first parameter specifies an initial "allowance" that helps to overcome initial

problems like forgotten password. The other parameter sets an error rate which can be tolerated over time.
  - The option **"Signaling failures blacklist: time in seconds to remove entries for which no event has occurred from score calculation:"** states how long an IP address continues to be suspected after it produced its first security events. Recommended value is 600.
  - The option **"Time in seconds to blacklist IP addr for signaling failures:"** determines how long an offending IP address stays on a blacklist. Recommended value is 3600.
- Define what occurrences add to the blacklisting score:
  - To include authentication failures and SIP protocol sanity checks, enable the options **Sanity** and **Auth** under **"Realms ▸ Call Agents ▸ Edit ▸ Firewall Blacklisting"**. If this CA option is not set, the traffic coming from IP addresses within this CA will not be blacklisted.
  - Additionally scripting actions used for constraining undesired traffic may be set up to add to the blacklisting score. To enable the "drop" action to add to the score, check its "Blacklist by firewall if repeated" option as shown in the Figure The Drop Rule Options. To count originators of requests that were rejected by a downstream server, use the action **Log message / Event for replies**, include message codes you are concerned about and turn on the option *Log to firewall blacklist*. The Figure Scoring Rejected Requests shows an example of such a rule intended to decline scanning attacks trying out calls to various telephone numbers. The guesses frequently fail and cause the replies with code 604. If this happens, the action *Log message / Event for replies* increases the blacklisting score.

Figure 156: The Drop Rule Options

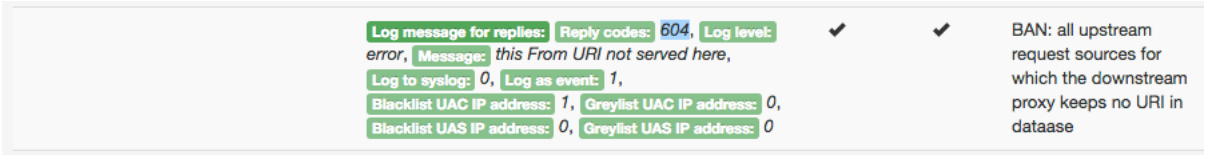

Figure 157: Scoring Rejected Requests

Note again that blacklisting can impair legitimate users who share the same IP address with an offending user. This is often the case with NATs or a trunk Call Agent represented by a single IP address and a single user that is misbehaving. In such a case, it may make sense to turn off auto-blacklisting for such a Call Agent, and deal with the misbehaving URI using SIP-layer filtering as shown in Section Manual SIP Traffic Blocking.

Automatic Proactive Blocking: Greylisting

Sometimes an automated blacklisting policy may be too reactive in that it begins to block traffic sources only when they have been already "caught" misbehaving. An alternative automated and sterner policy, greylisting, may be used instead to block suspicious traffic coming from an interface preemptively.

The basic idea is very simple: Permit signaling traffic from unknown sources for only a temporary "probation period", accept it if some legitimate criteria is established within this period and block (greylist) it otherwise. In this case, all packets coming from the IP address will be blocked at OS layer for maximum performance. This concept is stronger than blacklisting in that it doesn't wait until a misbehavior is spotted. An attacker trying to remain "under the radar" will not be tolerated any more. A single useless probing packet from his IP address to an ABC SBC signaling port will get him greylisted.

To enable grey-listing, you need to establish what makes legitimate traffic. An often used criteria is completion of authenticated SIP registration. To set up greylisting, proceed with the following steps:

- Turn greylisting on for an interface. Go to **System ▸ Interfaces ▸ Edit ▸ Greylist**. At this moment signaling coming over this interface from an IP address will be dropped if the criteria does not establish its legitimacy within a strict time window.
- Define the legitimacy criteria. This is achieved using the actions **Log to grey list** and **Log message / Event for replies**. The former immediately accepts a request source IP address. The latter does so later only when an answer with required status code comes back and can do so for UAC, UAS or both.
- Fine-tune greylisting global parameters if needed:
  - **time delay in seconds to give IP a chance to prove validity**,
  - **time period in seconds when IP can be blacklisted if repeats and did not prove validity**,
  - **time in seconds to keep IP on blacklist**,
  - **time in seconds to keep IP on whitelist**,
  - **additional ports or port ranges (a:b) to check in addition to signaling ports, space separated**.

Source IP addresses of cached registration bindings are implicitly accepted after receiving a successful response from the downstream registrar. This helps with a single administrative domain: an authenticated registration is quite a credible proof of sender's legitimacy.

However in scenarios with peering domains and other scenarios where SIP devices do not register, legitimacy of the senders must be established using some explicit criteria. To asses such a non-registering SIP sender, administrator must choose SIP transactions that demonstrate the sender is not an offender. This requires knowledge of a site's policy. For example, accepting an IP address based on an arbitrary 200-completed SIP transaction may be too relaxed, as any sender of a SIP OPTIONs "PING" packet that is "PONGed" would then qualify. Insisting on 200-completed INVITEs may be too harsh on the other hand, as a canceled call attempt would result in graylisting the caller. Therefore the acceptance policy must be chosen with knowledge of what SIP transactions shall or shall not be accepted by the downstream SIP elements.

The qualifying SIP transactions are tagged using the "Log to grey list" and if dependent on the resulting transaction status the "Log message / Event for replies" actions. When a transaction is processed using either action, and completes with a matching response code, then IP address of the SIP UAC, UAS or both will be accepted and will not be greylisted.

An example of such an A-rule is shown in Figure Greylisting Rule Example: Accept 200 REGISTERs and selected non-REGISTER codes. It accepts IP addresses from which REGISTERs come that complete with the 200 status code, and any other SIP requests that complete using some of the specified status codes. In all other cases, the IP address sending a packet to the ABC SBC will be blacklisted. That includes the cases when it is a non-SIP packet that doesn't even make it to rule processing, a REGISTER which doesn't result in a 200, and for example an INVITE which completes with the 604 code.

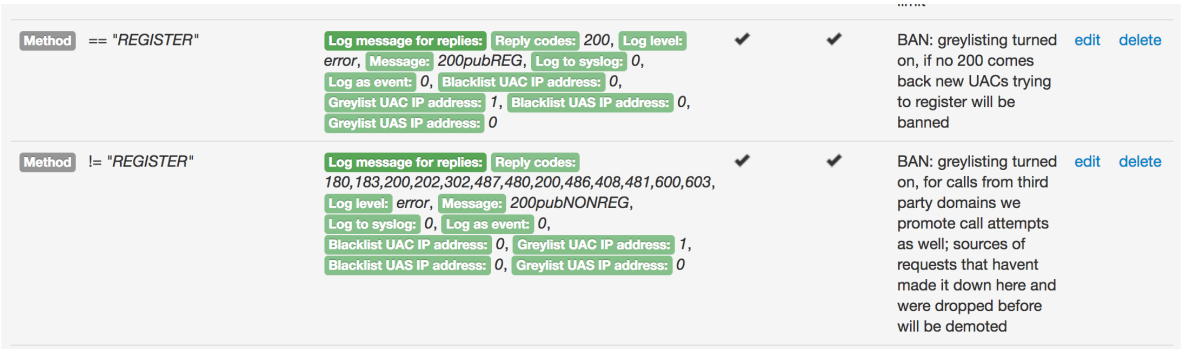

Figure 158: Greylisting Rule Example: Accept 200 REGISTERs and selected non-REGISTER codes

If we wanted to craft a more relaxed policy which does not inspect SIP answers coming back, we could use the action **Log to Grey List** instead (Figure Rule Example: Accepting an INVITE Sender's IP Address). It accepts all IP addresses from which an INVITE comes. Its actual impact depends on where in the rules this action is placed. If it was in beginning of the rules, it would only block offenders sending non-SIP or non-INVITE packets to the signaling ports. Therefore it is typically placed after several rules that drop undesirable traffic, such as request from well-known scanners or unsolicited OPTIONs.

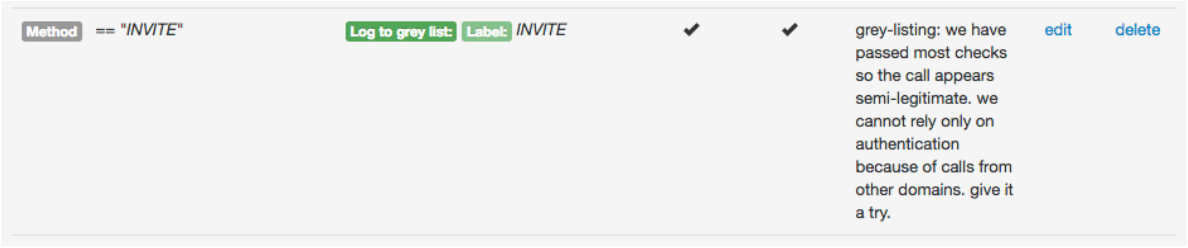

Figure 159: Rule Example: Accepting an INVITE Sender's IP Address

We also have to care about outbound SIP requests. Answer packets coming back trigger the greylisting process and we need to have an acceptance policy as well. Typically it is quite simple under the assumptions that requests sent to outside express consensus to communicate with the outside IP address. Therefore installing a rule in C-rules to accept the destination address regardless of the response coming back will form a reasonable policy. Such a rule is shown in Figure Rule Example: Permit UAS's IP Address for Any Replies. A destination appears on the greylist only if it sends no answer.

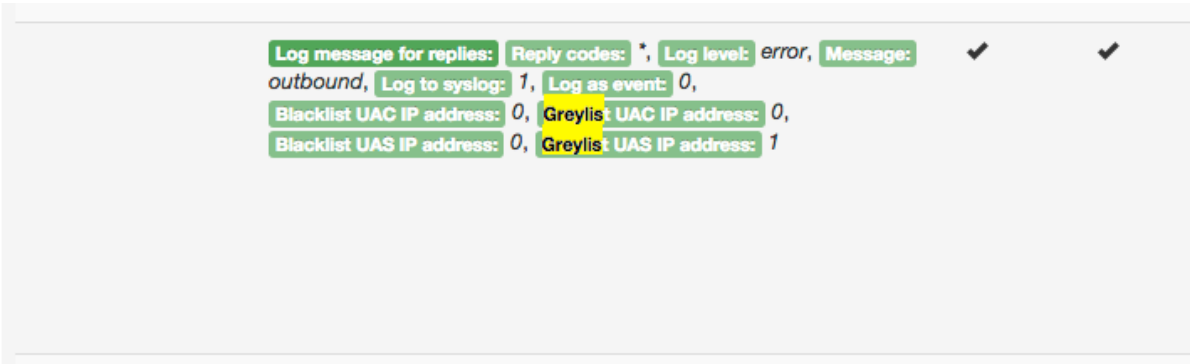

Figure 160: Rule Example: Permit UAS's IP Address for Any Replies

Note that like with blacklisting, greylisting may have side-effects when there are multiple users behind a single IP address. A legitimate user who proves himself and promotes his IP address by the greylisting procedures makes traffic of other users behind the same IP also legitimate.

Blacklisting and greylisting may be used at the same time. In this case the side-effects of blacklisting will prevail as blacklisting goes first in the processing order. Then even if an IP address is accepted by the greylisting criteria, and a misbehaving user will cause the IP address to be blacklisted, all traffic from the IP address will be blocked.

It is also important to know that ABC SBC resets greylists upon every restart and starts re-learning them. This makes re-configuration and/or rapid failovers more robust against grey-listing innocent IP addresses. Otherwise a change of greylisting policies could fail to accept an IP address that has been already spotted under a previous policy. Similarly, a fail-over back and forth may also result in greylisting a legitimate IP address.

Checking the actual status of an IP address can be done on the administrative page **"System ▸ Firewall ▸ Search IP"**, from where one can also retrieve the full current blacklist and greylist.

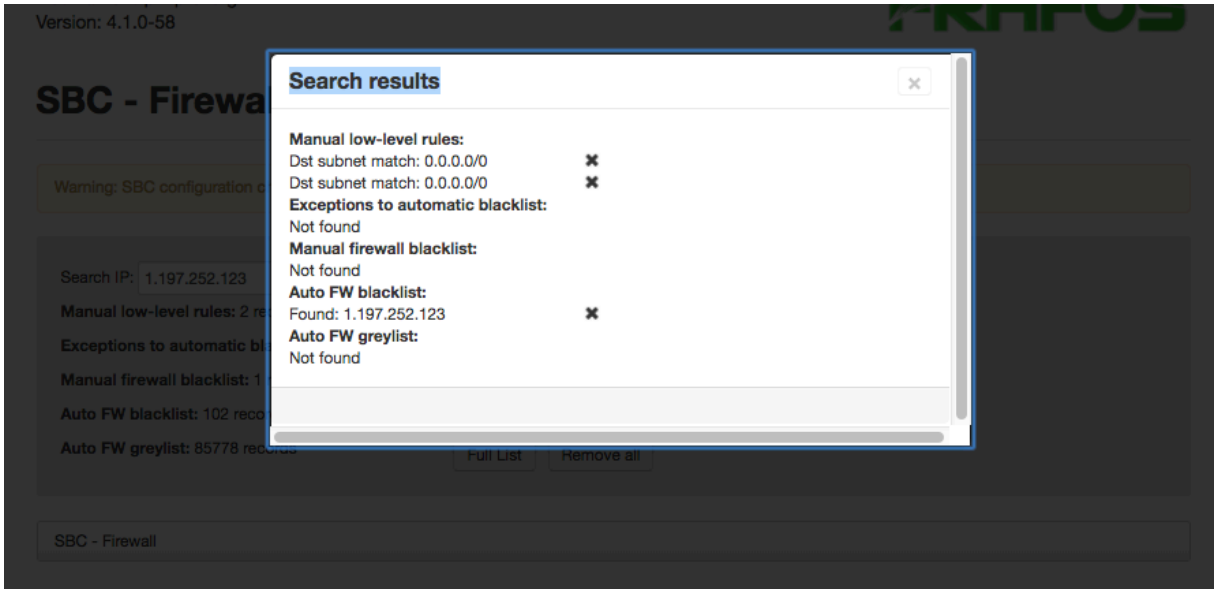

Figure 161: Firewall Search IP Box Results

Manual SIP Traffic Blocking

The manual blocking is used to block well-known offending traffic using SIP-layer criteria. The SIP-layer blocking allows to establish SIP-layer filtering criteria, and it also allows to indicate to the upstream SIP client why a request is being denied using a SIP response code.

The reasons for using this type of blocking can be multifold: declining traffic from unsupported call agent types, refusing to process some unsupported applications like SIP for presence, or banning traffic from SIP users that have become unwelcome and cannot be dealt with using IP-layer blacklisting because they share IP address with other legitimate users.

A call can be refused silently or using a SIP response using either of the following methods:

- *Reply to request with reason and code.* This action declines a SIP request using response code and phrase. Optionally a header field may be attached to the response. Replacement expressions can be used in the response phrase and header field. Multiple header fields can be introduced by putting $\backslash r \backslash n$ between them. An event of type "call-attempt" is generated for declined INVITEs.
- *Drop request.* This action drops a request silently and generates an event of type "message-dropped". Events can be grouped by a key in which case the events repeat within short interval of time (ten seconds) if their keys differ. If there is no key, the event does not repeat until offending messages stop to arrive for ten seconds.

If either action is executed, rule processing stops immediately and no further rules are processed. Neither do the requests count towards limits (see Section Traffic Limiting and Shaping) if the limits are placed behind the reply/drop actions.

The remaining question is how to discriminate between trusted and untrusted traffic. The ABC SBC can use any of its rule conditions described in Section Condition Types. The most often used conditions include:

- SIP header elements (Section Blocking by User-Agent, From and Other SIP Headers Fields)
- Source IP address (Section Blocking by IP Address)
- Registration status (Section Blocking a User by his Registration Status)
- Geographic origin (Section Blocking by Geographic Origin)

The following subsections documents the cases that are commonly used to filter out unwanted traffic based on different message elements. In the simple case, the tested elements are checked against fixed values like in the Figure The Drop Rule Options where the SIP requests are dropped if their Header

Field User-Agent contains “scanner” or “sipcli”. If the list of values to check against is longer, devising many rules may become cumbersome, use of provisioned tables is recommended as shown in the Section Provisioned Table Example: URI Blacklist.

Blocking by User-Agent, From and Other SIP Headers Fields

SIP request elements include many header fields upon which an administrator may make an accept/reject decision. For example, a SIP user may be found problematic and blocking his IP address is impossible because there are other legitimate SIP users behind the same IP address. In such a case it makes perfect sense to block all SIP requests with an offending address in the SIP From header field. Alternatively a whole domain can be blocked the same way. Conditions for this, From URI and From Domain, are available in ABC SBC rules, others are described in the Section Condition Types.

Not all header field names are available in the SBC rule conditions, and some may be even custom-made. Therefore there is also the possibility to refer to a header field by header name. That can be particularly useful when checking for some well known User Agent types that show their signature in the *User-Agent* SIP header field. Cases have been reported when this type of filtering has been used to block traffic from SIP devices with new untested firmware causing registration storms. Other common case is blocking well-known SIP scanners, one of such being known as “friendly-scanner”. Their packets look like this:

```
OPTIONS sip:100@212.79.111.155 SIP/2.0.
Via: SIP/2.0/UDP 37.187.191.144:5064;branch=z9hG4bK-3414626242;rport.
Content-Length: 0.
From: "sipvicious"<sip:100@1.1.1.1>;tag=643434663666396231336334013337313833393333235.
Accept: application/sdp.
User-Agent: friendly-scanner.
To: "sipvicious"<sip:100@1.1.1.1>.
Contact: sip:100@37.187.191.144:5064.
CSeq: 1 OPTIONS.
Call-ID: 383887304209490351968881.
Max-Forwards: 70.
```

A rule to detect, drop and record such requests from inbound (A) rules is shown in Fig. Inbound rule for refusing calls from a certain user agent.

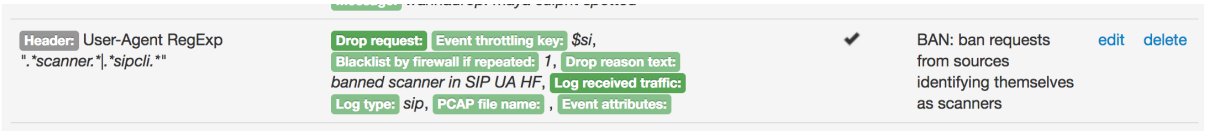

Figure 162: Inbound rule for refusing calls from a certain user agent

If the number of blocked elements become too long to have a separate rule for each of them, one can also utilize the provisioned tables as shown in the Section Provisioned Table Example: URI Blacklist.

Blocking by IP Address

It is possible to block a single IP address or multiple IP addresses matching a text pattern with actions configured with Source IP condition in the inbound (A) rule see Fig. Inbound rule for refusing calls from a certain IP address.

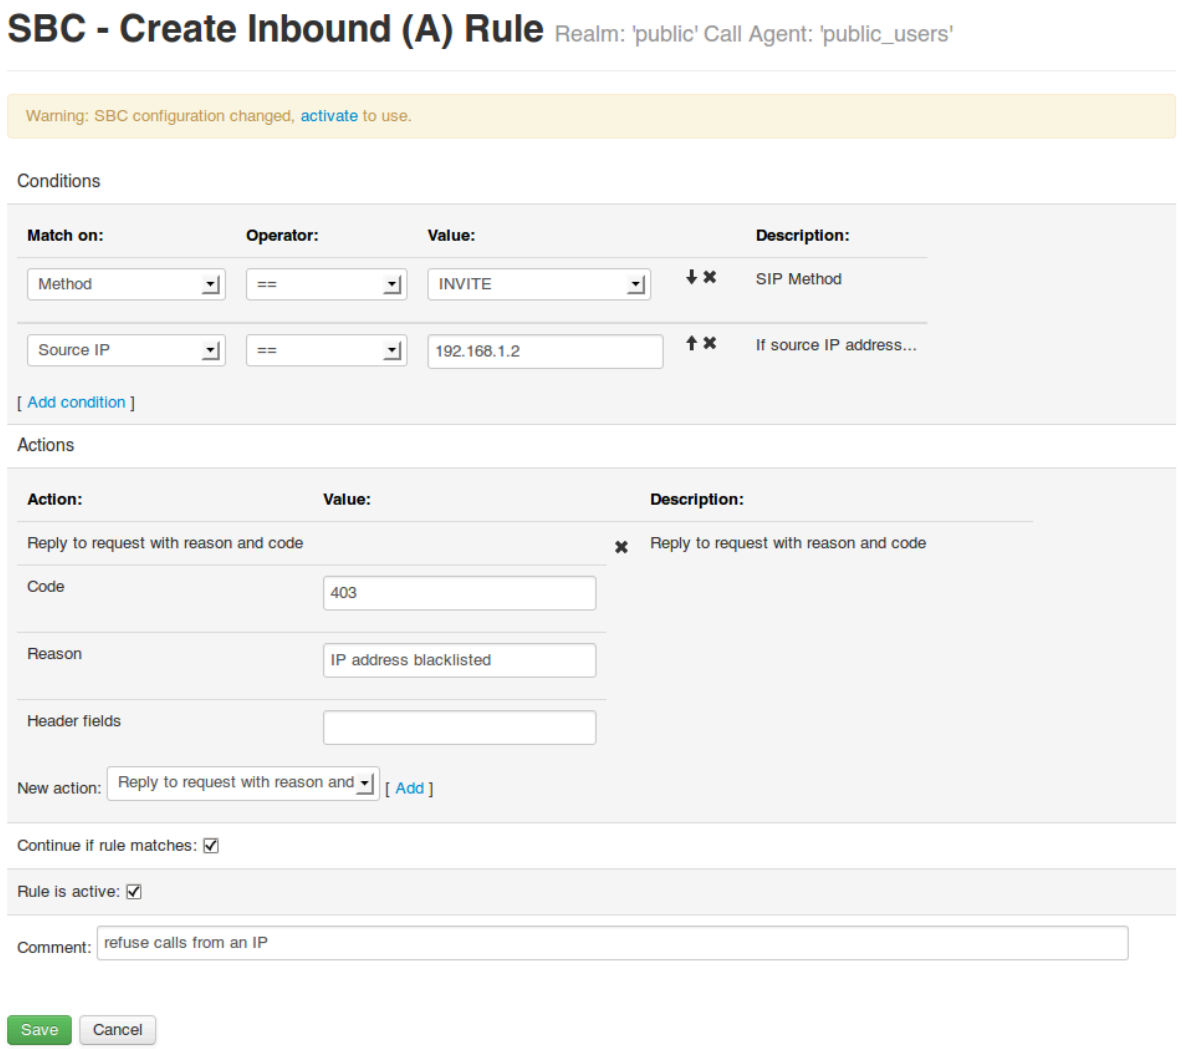

Figure 163: Inbound rule for refusing calls from a certain IP address

Blocking by IP Address Range

The simplest way to block a range of IP addresses is to create a Call Agent for such an IP address range, see Fig. Definition of a Banned Call Agent, and create an inbound (A) rule for this call agent without conditions that will refuse all messages from it see Fig. Rules for a Banned Call Agent,

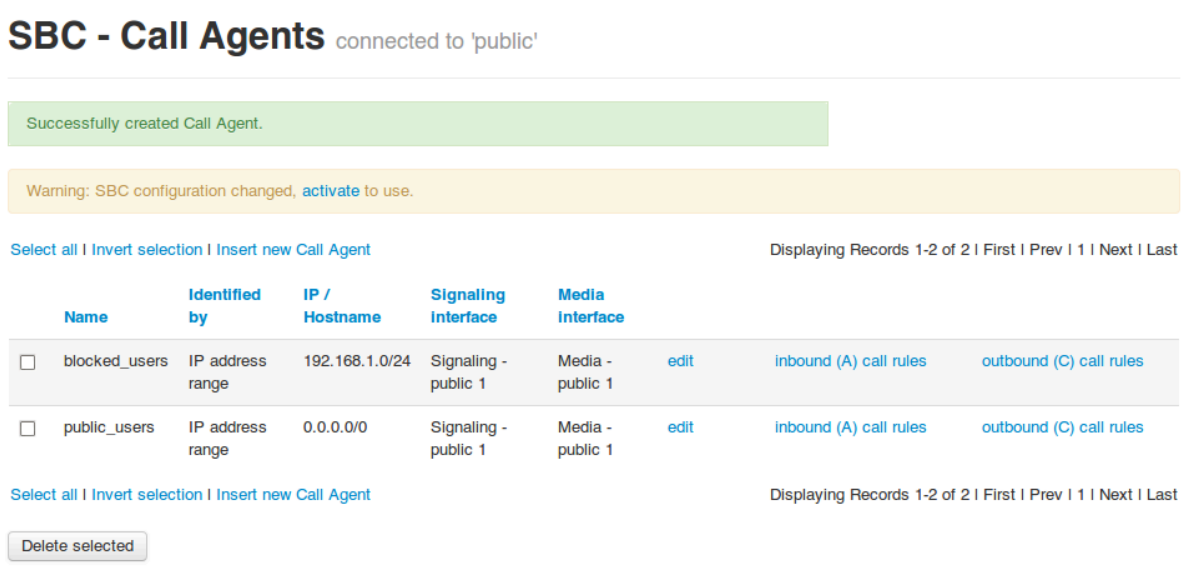

Figure 164: Definition of a Banned Call Agent

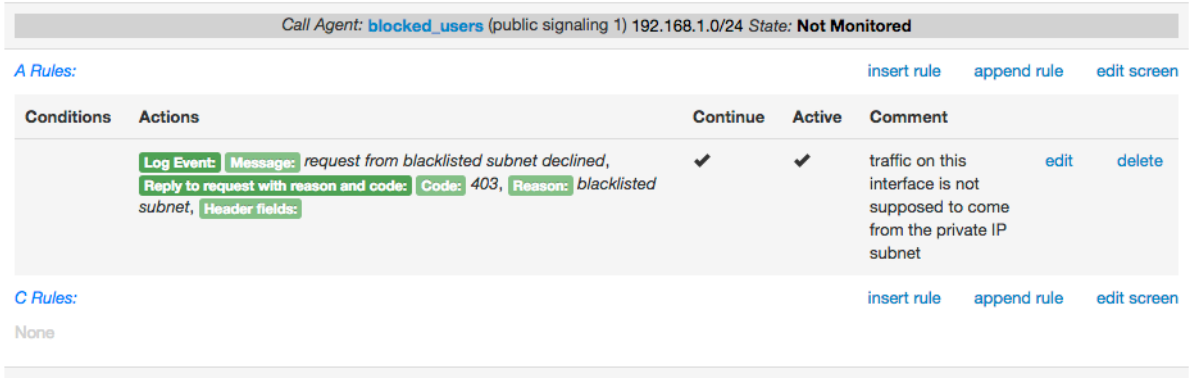

Figure 165: Rules for a Banned Call Agent

Additionally this rule example uses the "Log Event" action to alert administrator of traffic violating his policy.

Blocking a User by his Registration Status

Inbound (A) rules offer a possibility to enforce an administrative policy by blocking the request (usually an INVITE) if its initiator is or is not registered by using condition *Register Cache*. It also can be used as some form or caller prioritization if used together with CAPS limit. The test against the register cache is made using one of the following keys:

- From URI (AoR+Contact+IP/port)
- From URI (AoR+IP/port)
- Contact URI (Contact+IP/port)
- To URI (AoR)
- R-URI (Alias)

Such requests can be refused with **Refuse call with reason and code action**, see Fig. Inbound rule for refusing calls based on registration status.

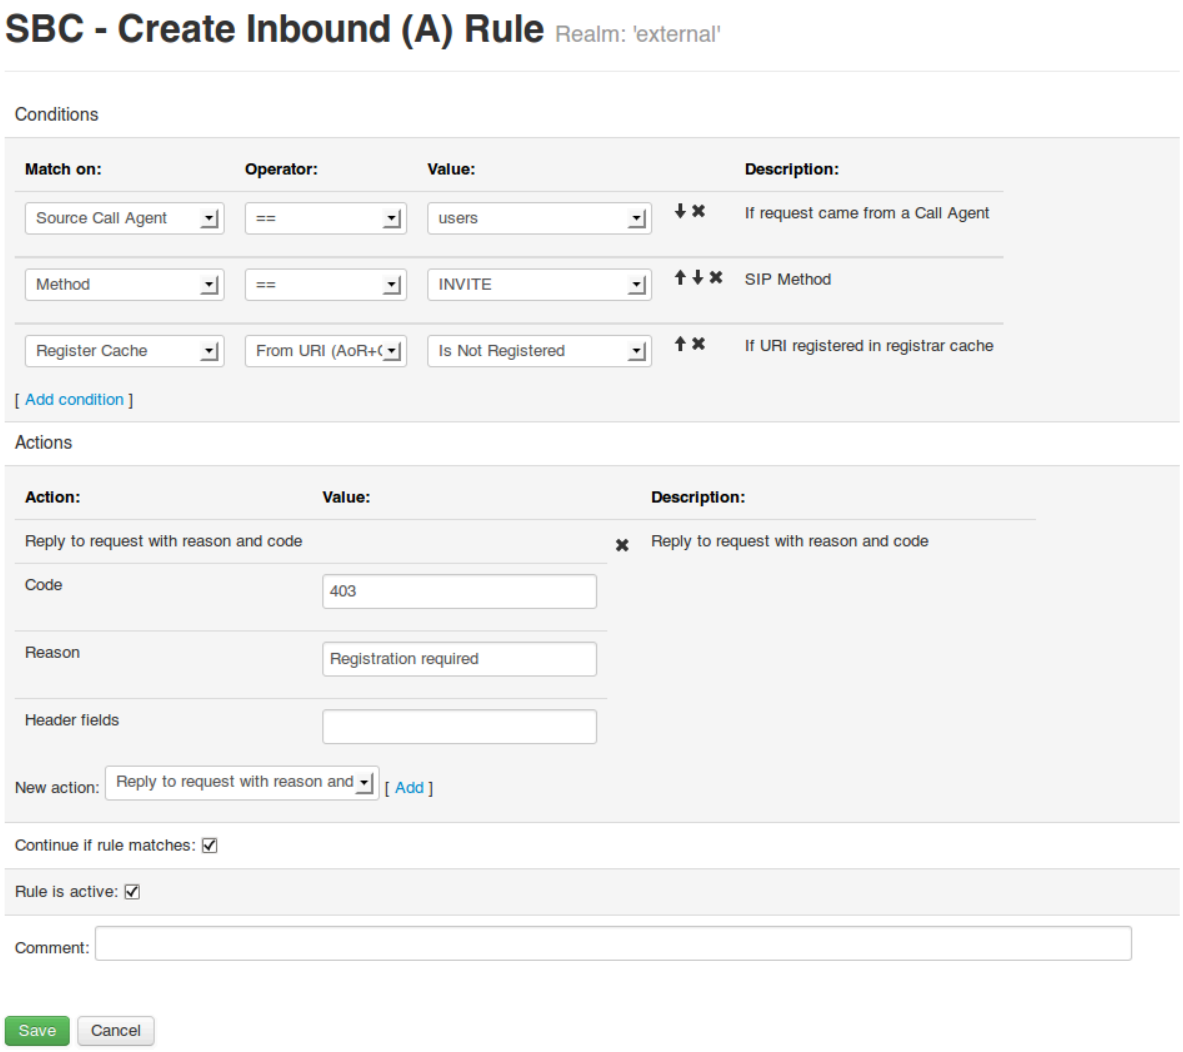

Figure 166: Inbound rule for refusing calls based on registration status

Blocking by Geographic Origin

The ABC SBC can also block or otherwise discriminate incoming requests based on the country code of the region from which they are coming. The region is determined using a Geo-IP database from request's source IP address. The example here generates custom events when a request comes from France.

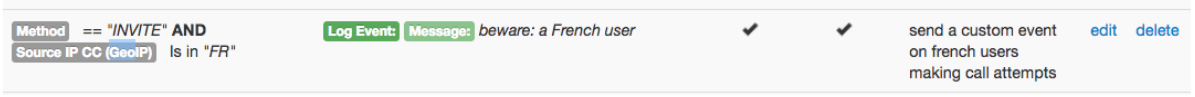

Figure 167: Inbound rule for Reporting on French Request Originators

Traffic Limiting and Shaping

Like any other Internet-based service VoIP servers can be the target of denial of service attacks. By generating a flood of SIP requests a malicious attacker can overload the VoIP infrastructure. Such overload conditions can negatively impair established calls and calls in progress and need to be controlled. Similarly, authorized users of a SIP service can access the service in a way that reaches abusive dimension and needs to be controlled as well. For example, a provider offering a flat-rate service to consumers may find that whole PBXs are connected to the SIP accounts. This would result in losses since the pricing calculation anticipated different usage calculation. Therefore traffic control is also needed in such a situation.

The ABC SBC offers several forms of controlling SIP and RTP traffic which are described in this Section. These are implemented as rules which can be placed in A or C rule-basis. If used in inbound A-rules, the limitations refer to traffic **coming from** a Call Agent or Realm. If used in outbound C-rules, the limitations refer to traffic **sent to** a Call Agent or Realm. In either case the limitations only affect calls that passed the limitation action. Reversely, calls that have not been processed using a limit action are not subject to such a limit. To make sure that all calls within a realm or call-agent are subject to a limit, the action must be placed in beginning of the rules without any condition.

More often, the call limits need to be related to a subset of traffic. For example only one parallel call may be permitted per IP address. The criteria can vary depending on use-case and therefore the limiting actions have an optional variable key parameter. The key specifies which traffic portion the limit applies to, and can use replacement expressions (see Section Using Replacements in Rules). All messages (and no other) that have the same key count towards the limit. Two often used keys are source address and combination of source address with From URI. The former (denoted as **"$si"**) checks all traffic coming from any single IP address against the respective limit. The combination of source address with AoR (denoted as **"sifu"**) allows that requests with distinct From URIs count against their own limits even from behind a single IP address – particularly useful when the IP address belongs to a PBX which serves numerous SIP addresses.

When the *"Is global key"* option is kept unchecked, the indexing key is scoped to the entity the rule belongs to (Realm or Call Agent). This means that the real key used to index the corresponding measurement is a compound of the indexing key and the entity. If, however, the key is declared to be global (by checking the *"Is global key"* option), the index is solely determined by the key entered in the *"Key attribute"* field. This means that if the same indexing key is used in another rule block (for example for another Realm or Call Agent), the limit will be applied jointly for calls on which this other rule block applies.

The traffic limiting actions also generate events when traffic does violate the limits. This is important for administrators to be able to notice such conditions and consider how to deal with such violations further: Whether to recognize these as illegitimate and continue blocking the originators, or to lift the limits if they find the above-limit usage has a legitimate reason. Only one event is produced for a detected excess of traffic limit, regardless of its duration. However, if the excess calms down and emerges again after ten seconds, a new event will be generated.

To make sure that an administrator can be alarmed even before a limit hits and starts to drop traffic, some of the traffic limit actions have the "soft-limit" option that creates diagnostic **notice** alarms but does not drop any traffic. Also, in the case that the traffic violates the "hard" limit repeatedly, the option "Report abuse" allows to block the offending traffic source – see Section Automatic IP Address Blocking for additional information.

The following call limit actions are available for use in A- and C-rules:

- **Limit parallel calls** - Set limit for number of parallel calls. New calls arriving in excess of this limit will be declined using the *403* SIP response. To make it easier to find the cause, the response includes a *Warning* header field with an additional hint: *Warning: Caps limit reached.* For example, to limit the number of parallel calls from the ABC SBC to a Realm or Call Agent, add a **Limit parallel calls** action to its outbound rules. Incomplete call attempt in progress whose context resides in memory also count temporarily towards the limit together with established call. That's an important security aspects: it makes sure new calls in progress are declined and cannot establish calls later that would exceed the limit. To limit the number of parallel calls from a Realm to the FRAFOS ABC SBC, add a **Limit parallel calls** action to the Realm's inbound rules. The action includes the following parameters:
  - *Limit parallel calls* – the actual number of parallel calls permitted.
  - *Key Attribute* and *Is Global Key* optionally define which partition of traffic counts towards the limit.
  - *SIP response code* and *SIP response reason* specify what type of reply is sent in response to a request that violated the limit. Optionally, header fields such as `Warning` may be added to the response using

the *SIP Header* option. This option is intended to provide upstream client and troubleshooters with additional information explaining why a request is
  - *Soft-limit* value allows to specify the "soft" threshold which if exceeded will generate a diagnostic event.
  - *Report abuse* checkbox makes occurrence of a traffic limit violation count against automated IP address blocking score.
  - *SIP response code* and *SIP response reason* specify what type of reply is sent in response to a request that violated the limit. Optionally, header fields such as `Warning` may be added to the response using the *SIP Header* option. This option is intended to provide upstream client and troubleshooters with additional information explaining why a request is
- **Limit CAPS** - Set limit for SIP request rate. If the request rate exceeds this limit, new call attempts will be declined using a *403* SIP response. Note that when a request is authenticated using SIP digest, it results in two transactions, both of which count towards the CAPS limit. New dialog-initiating (e.g. SUBSCRIBE) and out-of-dialog (e.g. unsolicited NOTIFY) requests also count against the CAPS limit and will be dropped if they exceed it. SIP requests belonging to a dialog that has previously passed the limit test will all be accepted. Retransmissions do not count towards the SIP limit. The action includes the following parameters:
  - *Limit CAPS* – the number of permitted SIP requests per unit of time. These two values define the permitted signaling rate.
  - *Time Unit* – length of time unit in second. Even if the number of permitted requests grows proportionally with length of time unit and yields the same signaling rate limit, longer time units are more permissive as they can accommodate more intense bursts.
  - *Key Attribute* and *Is Global Key* optionally define which partition of traffic counts towards the limit.
  - *SIP response code* and *SIP response reason* specify what type of reply is sent in response to a request that violated the limit. Optionally, header fields such as `Warning` may be added to the response using the *SIP Header* option. This option is intended to provide upstream client and troubleshooters with additional information explaining why a request is being dropped.
  - *Soft-limit* value allows to specify the "soft" threshold which if exceeded will generate a diagnostic event.
  - *Report abuse* checkbox makes occurrence of a traffic limit violation count against automated IP address blocking score.
- **Limit Bandwidth (kbps)** - Set bandwidth admission limit for codecs. If current total sum of maximum bandwidth as signaled in SDP exceeds this limit, the signaling request will be rejected using a 403. For example, the limit of 30 kbps will reject any incoming INVITE that, among others, offers G.711 codec (64 kbps) in its SDP using a SIP 403 response. This type of limit only serves as initial admission control and does not guard the actual RTP usage. A sender is not hindered to send more RTP traffic than advertised in SDP unless the **Limit Bandwidth per Call** action is applied. The action includes the following parameters:
  - *Limit Bandwidth (kbps)* – maximum permitted bandwidth
  - *Key Attribute* and *Is Global Key* optionally define which partition of traffic counts towards the limit.
  - *SIP response code* and *SIP response reason* specify what type of reply is sent in response to a request that violated the limit. Optionally, header fields such as `Warning` may be added to the response using the *SIP Header* option. This option is intended to provide upstream client and troubleshooters with additional information explaining why a request is being dropped.
  - *Soft-limit* value allows to specify the "soft" threshold which if exceeded will generate a diagnostic event.
  - *Report abuse* checkbox makes occurrence of a traffic limit violation count against automated IP address blocking score.
- **Limit Bandwidth per Call (kbps)** - Set limit for RTP traffic per call. This action observes all RTP streams, video and audio, of a call, and if the actual traffic rate exceeds the limit, the RTP packets

will be dropped. This action has the only parameter, the threshold value in kbps. RTCP traffic is not counted against the bandwidth limit and this bandwidth limit is only effective if RTP anchoring is enabled for the call. The limit includes RTP packets including RTP headers and excludes lower layer overhead (UDP,IP). For example for g.711 that makes 68.69 kbps (64kbps codec, 4.69 kbps RTP) and excludes 10.91 kbps overhead (3.13 RTP, 7.81 IP). For GSM the audio and RTP bandwidth is 17.69 (13 kbps GSM, 4.69 RTP), IP and UDP overhead is 10.94 kbps.

Note that limits are only applied to SIP requests that encounter the respective limit rule. That means that a newly introduced limit does not affect established calls. It also means that if call processing is stopped due to declining or dropping the call before the limit rule is evaluated, the declined call attempt doesn't towards the limit. Example of such rules where calls are declined before counted against a CAPS limit is shown in Figure Order of Rules Matters: Dropped Calls Don't Count Towards Limits.

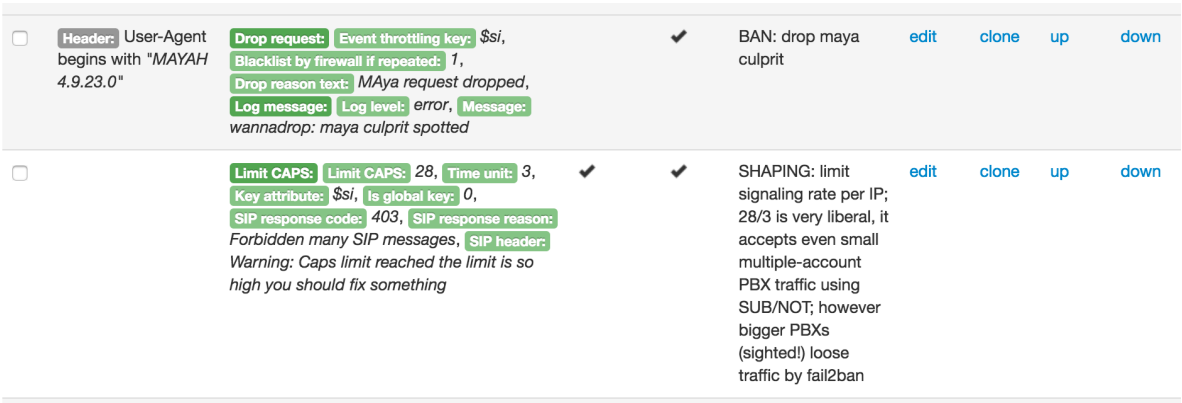

Figure 168: Order of Rules Matters: Dropped Calls Don't Count Towards Limits

The most delicate part when setting the limits is finding the appropriate threshold values. Definition of appropriate values depends on what type of SIP User Agents are being used and how. Specific aspects causing higher traffic rates need to be considered to make sure that legitimate traffic will not be discarded:

- soft-clients often support SIP for presence (RFC3856). The amount of traffic, especially when such a client starts, can be high and grow with the length of the buddy list.
- Registration throttling (Section Registrar off-load) is often used to keep NAT bindings alive. The limit rate needs to be adjusted to the throttling rate.
- PBXs and Integrated Access Devices and most importantly trunking peers send traffic for many users from a single IP address.

Traffic Limiting and Shaping by Example

In the following example we implement a policy to shape incoming traffic for a public SIP service for personal use. The example is intended to be rather liberal and sets the threshold relatively high for the anticipated use to make sure it doesn't break some traffic-intensive use-cases.

We start policing VoIP calls in the first rule. To make sure that even a nervous caller attempting to reach a busy destination doesn't exceed his limits, we permit 10 requests every 30 seconds for every source IP address (the "$i" in the key parameter). Note that the actual number of call attempts may be lower by one half, since SIP authentication attempts preceding the actual call attempts count towards the limit as well and double the number of requests.

The next rule throttles registrations. We know that several popular consumer Integrated Access Devices (IADs) offer several SIP accounts. We want to make sure that the devices don't get locked out when they boot and send REGISTERs for all of their SIP accounts. We also need to account for the authenticating transactions. Therefore we set the limit to 10 requests every thirty seconds and key the limit by combination of IP address and From URI. That means that the limit can only be exceeded by requests coming from the same IP address and bearing the same From URI. In other words, even if many REGISTERs

come from behind a single IP address the limit will only be hit if they use the same URI. If the URI is registered from an IP address at a rate beyond the limit, parallel registrations of the same URI from behind a different IP address will not count towards the same limit.

Further we impose a general limit on all SIP transaction types. Especially soft-phones are known to send a lot of "noise": SIP PUBLISHes, SUBCRIBEs, NOTIFYs, OPTIONS and other request types. We permit 28 requests every three seconds from every single IP address.

Last but not least: we limit the number of parallel calls to 5 per IP address.

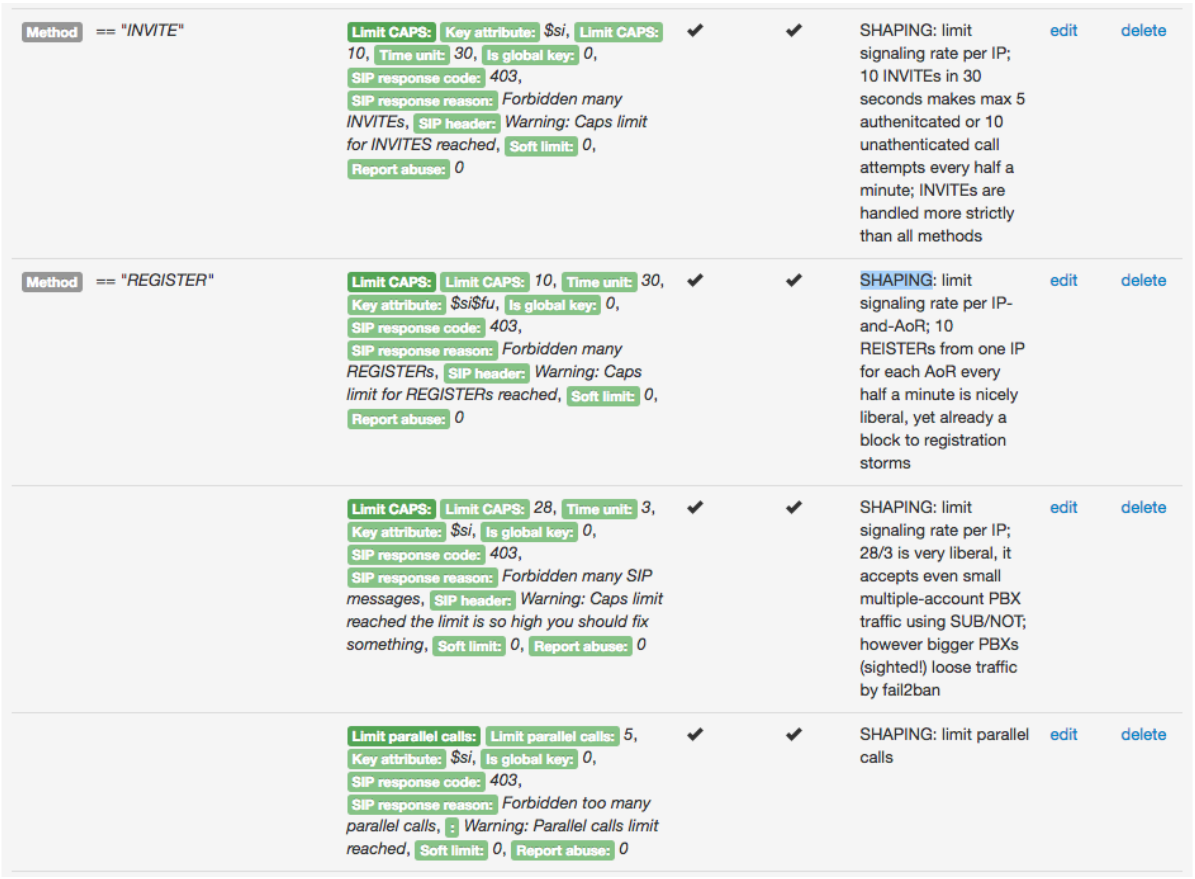

Figure 169: Limit on number of Call Attempts per Second

Bandwidth limits by example

The ABC SBC can limit the bandwidth admitted for a calls' media streams. The action **Limit Bandwidth (kbps)** has as a parameter the bandwidth in kilobits per second to which the call should be limited, see Fig. Example of Shaping Rules. Attempts to set up calls exceeding this bandwidth will be rejected using a 403 response.

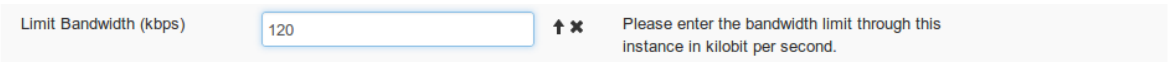

Figure 170: Example of Shaping Rules

Call Duration Control

By limiting the maximum duration of calls one can on the one hand prevent "bill shocks" when some customer fails to terminate a call in a proper manner. Additionally, attackers might try to deplete the resources of the SBC by generating calls with long durations causing a saturation of the call processing capacity of the SBC.

Setting Call Length Limits

The **Set call timer** action sets the maximum duration of the call, in seconds. If a call exceeds this limit, the ABC SBC sends a BYE to both call participants and generates an event of the call-end type.

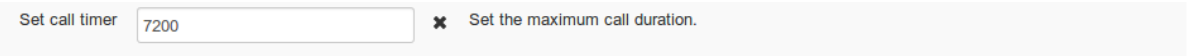

Figure 171: Set call length

If this action is executed several times, the call duration will be the lowest call timer set, regardless of the order in which the actions are executed.

Controlling SIP Session Timers (SST)

SIP Session Timers (SST) is a mechanism defined in **RFC 4028** that can be used to make sure that calls are ended after a period of time even if one endpoint disappeared without properly terminating the call. This is especially important if calls are billed using data derived from the signaling messages, but also makes sure that unused resources are released properly. In order to achieve this, periodically a refresh of the SIP dialog is done by using a re-INVITE or an UPDATE. If the refresh request fails, e.g. if it times out, per the standard SIP mechanisms the dialog is torn down and related resources are released.

Note that the session refresh, i.e. the re-INVITE or UPDATE that is done here, is a normal re-INVITE or a normal UPDATE, so all the normal rules regarding the re-negotiating of the session apply. For example, a re-INVITE which is triggered by Session Timers may modify the session by selecting another codec or other codec parameters.

SIP Session Timers is a mechanism to negotiate which of the endpoints in the SIP dialog does this refresh. After the negotiation that happen with some specific headers (Session-Expires/x,Min-SE) in the INVITE or UPDATE and the responses to it, it is clear who has the refresher role.

Unless explicitly configured otherwise, the negotiation is conducted directly between the caller and the callee, with the ABC SBC merely passing SST-related information between the two call legs. However, if necessary, the ABC SBC can take control of the negotiation.

The ABC SBC supports SIP Session Timers even if one or both endpoints do not have support for Session Timers.

There are separate actions for enabling SST on the caller and the callee leg

- Enable SIP Session Timers (SST) - caller leg
- Enable SIP Session Timers (SST) - callee leg

The remote UA may leave it open who should be the refresher (by not including a refresher=uac or refresher=uas parameter in the Session-Expires/x header). In this case the ABC SBC has the possibility to select whether the ABC SBC should take on the refresher role or the remote UA should do it. This can be selected by the 'let remote refresh' option. It may be safer to do the refresh from the ABC SBC, as some UAs do not perform the refresh properly, even if they have said they would do it. On the other hand, if NATs are involved, there is no keepalive and the refresh interval chosen is too long, only the remote UA which is behind the NAT may be able to do the refresh thus reopening the NAT, in which case it may be safer to let the remote UA do it.

The SST mechanism also negotiates the time after which the refresh is done. The timer parameters - proposed Session Expiration and Minimum Expiration, in seconds - can be set individually for each leg.

Setting RTP Inactivity Timer and Keepalive Timer

These two timers help to detect situations in which due to some network failure a phone call has already stopped. It requires media anchoring to be enabled.

It often occurs that a call party becomes suddenly unavailable and a call remains "hanging". This may happen for example due to a software error in a softclient or a disconnected IP network. To make sure such a call doesn't continue, an RTP inactivity timer may be configured. If configured and either party stops sending RTP, a call is discontinued by the ABC SBC after the preconfigured period of inactivity. Eventually an event of type "call-end" is reported with originator field set to "rtp-timer-terminated".

The timer can be configured under "Config ▸ Global Config ▸ RTP handling ▸ RTP timeout". If set to zero, the timer is deactivated.

To make sure that a peering SIP device using a similar kind of timer doesn't disconnect a call which just occurs to produce no media (voice inactivity detection, on-hold), the ABC SBC may also be configured to generate RTP keep-alive packets. If set to a non-zero value (in seconds), the ABC SBC sends keep-alive RTP packets periodically.

This timer can be configured under "Config ▸ Global Config ▸ RTP handling ▸ RTP keep-alive frequency".

Both timers can be also set on a call-by-call basis under parameters of media anchoring (see Section Media Anchoring (RTP Relay)).

4.13.3 Collect Events: Gathering Usage Data in the ABC Monitor

Knowledge is never too dear. Sir Francis Walsingham, Queen Elizabeth's Principal Secretary

An administrator can only craft reasonable security policies if he knows what is actually going on. He must have access to detailed history of SIP user behavior, security incidents and unusual activities. This is indeed the purpose of "events". Events are detailed reports on user activity that encompass registration, call attempts, and security incidents.

Many individual events can identify need for administrator's attention. For example if a packet is dropped because it is coming from a SIP scanning software, an administrator may want to act and ban the source IP address.

Some events in isolation may alone not indicate a threat and need to be monitored by their quantity and trend. For example, an isolated authentication failure can be caused by a password typo during SIP authentication process. However if many such occur in series, chances stand high it is some kind of password-guessing attack. Being able to recognize such repetition allows the ABC SBC to act automatically and even ban offending traffic without the human administrator's intervention. (see Section Automatic IP Address Blocking).

Most of the events are always produced by the ABC SBC, and administrators don't need any extra action to enable them. They just need to be able to understand and analyze them.

The rest of this Section is concerned with the cases when reporting events is optional and needs to be turned on explicitly to alert on possible departures from a SIP site's policy.

Reporting Security Events

As security events failures are reported when a particular administrator-defined policy is being enforced. The only exception is an authentication failure which is always considered a security threat.

The events are reported *only if* corresponding actions are executed and proper parameters are set. Therefore it is the rule conditions that primarily determine when to trigger an event.

The following table summarizes how rules must be set up in order to generate proper events. All shaping actions report limit violations if event reporting is enabled. So does the `drop` action when executed on an incoming SIP request, and `log-reply` when a request is rejected downstream.

Particularly the `log-reply` action is important as in some cases the downstream SIP elements may know better than the ABC SBC that a request is illegitimate. This is for example the case in a scanning

attack when an attacker attempts to probe all possible SIP addresses. The ABC SBC is unaware of individual users and is not in position to repel such an attack straight off. However the downstream server knows subscriber details and can reveal by proper SIP response codes that the requests are for invalid destinations. It may respond back with 604 for non-existing users or 403 for forbidden addresses. This way the ABC SBC can infer from the response codes received from downstream that the upstream request originator should be better blocked.

Event Type	Required Action	Required Parameter	Additional Information
limit	Limit parallel calls	Report Abuse	Traffic Limiting and Shaping
limit	Limit CAPS	Report Abuse	Traffic Limiting and Shaping
limit	Limit Bandwidth	Report Abuse	Traffic Limiting and Shaping
limit	Limit Bandwidth per Call	none	Traffic Limiting and Shaping
message-dropped	drop	Blacklist by firewall if repeated	Manual SIP Traffic Blocking
log- reply	Log message for replies	Log to Firewall blacklist	

Setting up Diagnostic Events

Diagnostic events are also of great importance to the process of continuous refinement of security policies and bridging the gap between liberal and strict policies. A too liberal policy may lead to exposure of a security gap. On the other hand a too strict policy that filters all unknown SIP elements is likely to break some SIP features. Diagnostic events allow to strike a compromise, in which a policy remains open and diagnostic events report on suspicious traffic patterns. An administrator can then inspect these in details and choose whether they are legitimate and can be preserved, or whether they shall be better banned.

An example of such policy is reporting on call from unregistered users (see Figure Rule Example: Report Calls of Unregistered Users). If an administrator feels uncertain whether such calls are legitimate or not, he may initially just observe them. To do so, he places `log-action` in an appropriate condition and then watches the reported events. These include detailed information about the calls in question and provide the administrator with insights needed for further refinement of his policies. He may for example find out that the call attempts are coming from a peering domain and are perfectly legitimate. Or he may find that they have no traceable originator and should be better blocked.

The following table lists actions that can be used to provide customized reports on observed activities. The shaping actions can include an additional lower limit to report on high traffic before the “hard limit” is hit and traffic begins to be declined. The `action-log` can report on any conditions identified in the ABC rules: unexpected URIs, traffic from unregistered users, and anything else that can be captured by conditions specified in Section Condition Types. The `message-log` event is used analogously, in addition to the event details it also collects the actual traffic that triggered the event.

Event Type	Required Action	Required Parameter	Additional Information
limit	Limit parallel calls	Soft Limit	Traffic Limiting and Shaping
limit	Limit CAPS	Soft Limit	Traffic Limiting and Shaping
limit	Limit Bandwidth	Soft Limit	Traffic Limiting and Shaping
limit	Limit Bandwidth per Call	none	Traffic Limiting and Shaping
action- log	Log Event	none	

Event Type	Required Action	Required Parameter	Additional Information
log- reply	**Log message for** replies	Log as Event	
message-log	**Log received** traffic	none	

4.13.4 Practices for Devising Secure Rule-basis

While we have shown in previous sections how to police traffic, collect diagnostics information and analyze it there is still a remaining question: how to put all of this together in a consistent configuration using the ABC SBC rules. The way the rules are compiled can have significant impact on the logic of the service.

When devising the rule-base, the following important choices must be made:

- Whether to use **Media Control** or not. Relaying media (Section Media Anchoring (RTP Relay)) provides the ABC SBC with more control and insight into calls at the price of performance and media latency. Also it is a necessity when NAT traversal (NAT Traversal) needs to be implemented, IP addresses of infrastructural elements behind an ABC SBC need to be hidden, transcoding (Section Transcoding) or RTP-to-SRTP conversion (Section RTP and SRTP Interworking) is needed. If used, which is nowadays the default choice, the latency impact can be mitigated by geographic dispersion (see Section Introducing Geographic Dispersion for more information on what difference geographic distribution makes in a cloud environment).
- Whether to use **Topology Hiding** or not. Topology Hiding obfuscates signaling so that it is hard for an external party to find IP addresses of the infrastructural elements behind the ABC SBC . We are describing the rules to be used for topology hiding in the Section Topology Hiding. Note that obfuscation of SIP traffic may make its analysis quite difficult. If used tracing of traffic using ABC Monitor is recommended.
- How to organize policing rules. A reasonable practice is to start with rules that identify and instantly drop undesired signaling traffic before “heavier-processing” rules, such as media control or database queries begin. We show our recommendations in the Section Devising a secure rule-base.

Topology Hiding

Some service providers are worried about disclosing IP addresses of their infrastructure to third parties, attackers and competitors. Unfortunately the SIP protocol does such a disclosure in a grand style: SDP payload shows IP addresses from/which media is sent, Contact header-field shows the IP address of an end-point, and so does pre-RFC3261 Call-ID header-field. Via, Route and Record-Route header-field disclose the path of a SIP message exchange. Other standardized and / or proprietary header fields can also carry IP addresses.

Therefore service providers concerned about such disclosures prefer obfuscation of the respective SIP message elements. It needs to be pointed out though, that what makes life harder for attackers makes it similar hard or even harder for service operators. Correlating messages with each other for sake of troubleshooting is much harder when they are modified.

In the following subsections, we will review the default topology hiding behavior and how to make it more transparent or more obfuscated.

Default Address Hiding

The default configuration of the ABC SBC tries to strike a good balance between the two extremes, full disclosure and full obfuscation. Already the back-to-back user-agent (B2BUA) design of the ABC SBC contributes significantly. The whole SIP path, as disclosed in *Via*, *Route*, and *Record-Route* header

fields is split in two call-legs, each of them terminated by the ABC SBC. As result, each SIP dialog party sees the SBC as its peer in these Header fields. Additionally the ABC SBC by default rewrites dialog information (the triple Call-ID, From-tag, To-tag) so that IP address present in pre-RFC3261 implementations Call-ID is obfuscated.

If more signaling transparency is need than this default behavior implements, transparent dialog ID can be enabled by an action as shown in the next Section. Also in some rare scenarios, the downstream elements in the SIP path may need to inspect the Via stack for the upstream leg. The ABC SBC reintroduces it when the following action is enabled: **Show Via**.

Transparent and Non-Transparent Dialog ID

The other concern is Call-ID – old-fashion SIP implementations pre-dating RFC3261 followed the RFC2543 specification and disclosed its address in Call-ID header-field. To make sure that the addresses do not get disclosed through this header-field when out-of-dialog or dialog-initiating requests are created by an elderly SIP User Agent, the ABC SBC can replace the Call-ID values with obfuscated values.

The choice whether to do is is administrator's. By default the ABC SBC does change the Call-ID Header Field value.

We recommend that administrators consider preserving the Call-ID for sake of troubleshooting. Leaving it unchanged makes correlation of incoming and outgoing SIP messages significantly easier. Enabling it is easy, what needs to be done is to place the following action in a Call Agent's or Realm's rules: **Enable transparent dialog IDs**.

Another advantage is that some non-standardized SIP extensions may want to take reference to a Call-ID value which becomes invalid once the ABC SBC changes it.

Hiding Addresses in Well-known SIP header-fields

When an operator is indeed concerned about disclosing internals of a Call Agent, the very step is to make sure that occurrences of the Call Agent's address in well-known header-fields are replaced with ABC SBC 's. Doing that is as simple as turning the Topology Hiding checkbox on under Call Agent's attributes. Once enabled all the following header fields will be rewritten accordingly:

- `P-Asserted-Identity` (**RFC 3325**)
- `P-Preferred-Identity` (**RFC 3325**)
- `Diversion` (**RFC 5806**)
- `History-Info` (**RFC 4244**)
- `Remote-Party-ID` (proprietary pre-3325)
- `Call-Info` (**RFC 3261**)
- `Warning` (**RFC 3261**)

Note that if the `Warning` header field is obfuscated, it is denoted using an additional `;topoh` parameter. This makes it clear that the address in the header-field is not genuine – otherwise a troubleshooter may be misled.

Hiding Contact Header in REGISTER

The *Contact* header is by default obfuscated by the SBC in all dialog-initiation transactions. Contact header field in REGISTER requests remains however untouched. If obfuscation is desirable, ABC SBC's register cache must be used that replaces the original Contacts with aliases.

The Section Registration Caching and Handling provides details about configuring registrar cache. This may be a reasonable option to be turned on alone for its "shock absorbing" and "NAT keep-alive" capabilities.

Hiding All Other Header Fields

Additional header-fields, standardized (Service Route **RFC 3608**) or proprietary, may appear and convey IP addresses. The ABC SBC only obfuscates the documented header-fields and doesn't interfere with others. If other header-fields are present and disclosure of IP address is a concern, the administrator can remove them at the risk of affecting the purpose they are serving. He can either remove the specific header-fields or use header field whitelist, i.e. remove all but well-known header-fields. SIP manipulation is described in detail in the section SIP Mediation, of particular interest is the action **set header whitelist**.

Concealing Media

Similarly like with SIP, the ABC SBC can put itself in the middle of the path and present itself to each call as its peer while hiding the other party. If the ABC SBC doesn't do it, IP address used for sending/ receiving media will be seen in SDP and in the actual RTP packets.

To conceal the media sender/receiver, enable **Enable RTP anchoring**.

The downside is that all media visits the ABC SBC, while increasing media latency and bandwidth imposed on the server. A detailed discussion can be found in the Section Media Handling .

Preventing SIP Digest Leak:

In order to effectively prevent malicious UAs from requesting a SIP Basic Authentication Digest from another UA with which a call has been established, it is necessary to take some measures to prevent authentication requests to be forwarded to UAs from other UAs that should not send any authentication requests.

For Call-Agents facing single end-user UAs, a simple method can be used to effectively block these authentication requests: 2 **UAC auth** actions are configured in the **A-rules** of the Call-Agent facing the end-user UAs (one toward "caller", and one toward "callee"). These actions shall be configured with a bogus username and password which will be used to reply the authentication requests from the malicious UA.

For Call-Agents representing a trunk line, a simple **header blacklist / whitelist** can be used to effectively filter out the following header fields:

- `Proxy-Authenticate`
- `WWW-Authenticate`

Preventing Resource Exhaustion Attacks:

To effectively prevent the SBC from being subject to resource exhaustion attacks (flooding based) but also from high traffic peaks, it is necessary to configure so called **Server Transaction limits** (see Server Transaction limits).

Devising a secure rule-base

Developing a reasonable security policy may be a delicate task for a SIP service administrator. A too strict policy may too easily "throw the baby out with the bathwater" and impair legitimate traffic. The other extreme, a too liberal policy, may be too inviting for an attacker. Finding the right balance between serving users and protecting them from attackers requires understanding of the service goals and risks and drawing a balance between them.

The policy represented by the ABC rules also has performance implications associated with it. Some rules, such as database lookups, have higher latency and lower throughput than others.

We therefore suggest that policies are crafted in order of increasing complexity, starting with rules that instantly reply certain requests and continue to more complex rules. Basically, all undesirable SIP messages should be eliminated by rules in the initial rule-base part before processing for the accepted messages starts in the other part. The following subsections show examples of such rules in such order.

Shaping the Signaling Rate

It makes sense to begin processing with a check against SIP rate limits. Placing the check in the very beginning makes sure that all incoming SIP requests are checked against these limits including requests that are dropped by rules.

In Figure Rule Example: CAPS Shaping we are showing a simple rule example for sake of this Section. The rule checks all incoming SIP messages against a request rate and declines messages in excess of the limit.

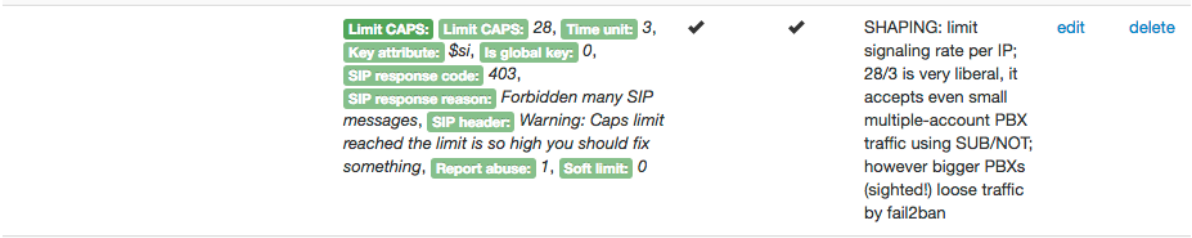

Figure 172: Rule Example: CAPS Shaping

More sophisticated examples for shaping rules have been given in Figure Limit on number of Call Attempts per Second in Section Traffic Limiting and Shaping by Example.

Instant Responses

Many requests come that do not require any sophisticated decision making: the right action is just to send a reply instantly. The reply can be positive or negative. Positive replies are typically sent in answer to some SIP User Agents' SIP-layer NAT keep-alives. Negative answers are sent when requests request some unsupported service, do not comply to some local URI conventions, or come from a User Agent type known to malfunction.

The following rule examples show both cases: positive reply (Figure Rule Example: Confirming SIP Keep-alives) for keep-alive messages and negative replies to decline a request for unsupported Message Waiting Indication server (Figure Rule Example: Declining an MWI Request).

Figure 173: Rule Example: Confirming SIP Keep-alives

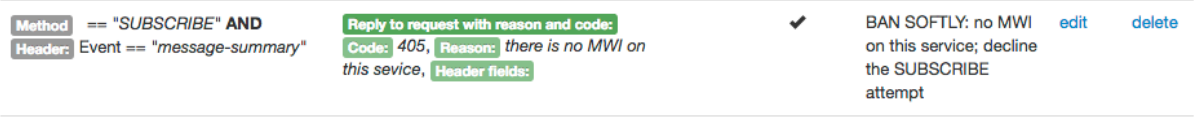

Figure 174: Rule Example: Declining an MWI Request

Dropping

With SIP requests who appear a security threat, dropping them silently is a safer choice than declining them. The less information an attacker gets, the harder it is for him to find a security gap in a SIP service. If an IP address is sending clearly offending traffic, it may even make sense to ban it entirely.

A typical reason for deploying such a restrictive rule is elimination of SIP scanner traffic. SIP scanners are automated tools that probe Internet address space to see if there are some SIP services running.

Such tools are even publicly available[7]. When such a tool finds a responsive SIP service, it continues looking for legitimate SIP addresses and it may even proceed to mounting a password-guessing attack. Such attacks are real: Once you start a SIP service on the public Internet, it takes no longer than few hours until the first scanning packets arrive. Note however that filtering such traffic is only eliminating naively crafted attacks that advertise themselves as such. More sophisticated attacks will certainly not do it and must be detected and repelled using other methods such as traffic shaping.

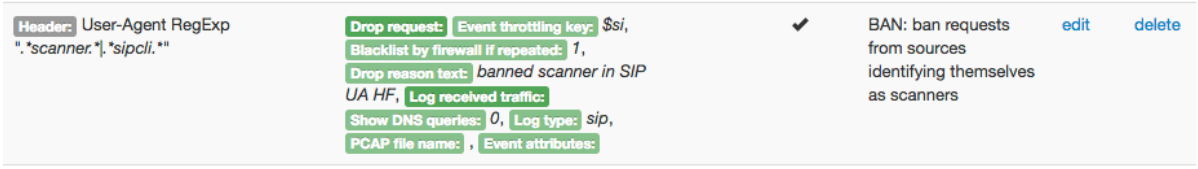

Figure 175: Rule Example: Eliminating Traffic from SIP Scanners

The rule has an important option turned on: "Blacklist by firewall if repeated". That means if the offending traffic appears repeatedly, the originator's IP address will be blacklisted.

Database Checks

By now, we have eliminated most of the unwanted traffic: we have declined excessive traffic, gracefully handled SIP-layer keep-alives, declined politely messages for unavailable services and dropped obvious security threats. The remaining traffic has been reduced to a level where we can deploy more expansive policy checks and dig in database. Often there are offending users identified by their URIs. A straightforward way to eliminate their traffic is to provision a list of such users and block SIP traffic if it comes from such users. Figure Rule Example: Prohibited URIs shows a rule that looks up SIP requests From URI in such a table and if the URI is found, drops the request silently.

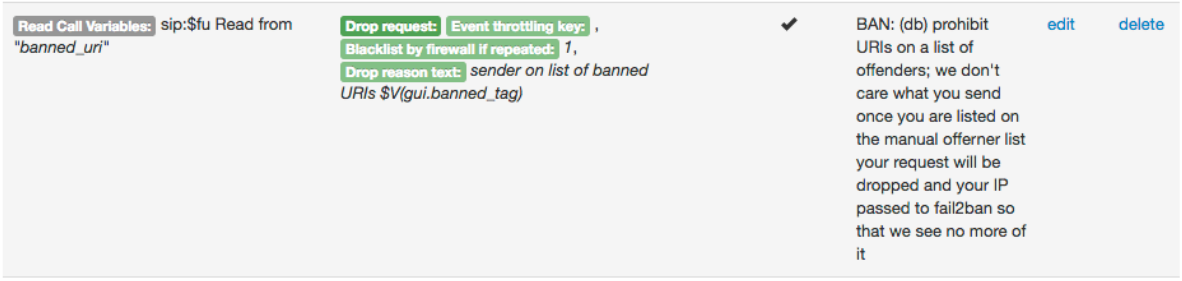

Figure 176: Rule Example: Prohibited URIs

More Limits

We may also want to constrain the number of parallel calls. We didn't place such a limit in beginning of our rule-set. The reason is that too many call attempts are rejected in the initial part of rule-set and count towards the limit too. When we place the parallel call limit in the rule-base after all unwanted traffic is rejected, the call attempts we chose to decline won't count towards the limit.

Figure Rule Example: Limit Parallel Calls shows rule for enforcement of maximum five parallel calls per single IP address. Also note that in this rule, no soft-limit warning is enabled and limit violations add to the security score computed by automated blocking (Section Automatic IP Address Blocking).

[7]SIP Vicious: https://github.com/sandrogauci/sipvicious

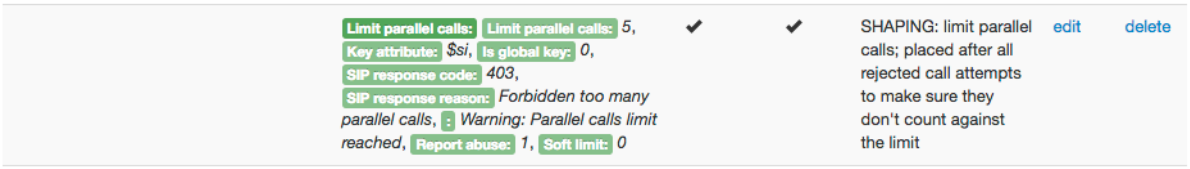

Figure 177: Rule Example: Limit Parallel Calls

Diagnostic Events

Often SIP messages do appear whose purpose is not entirely clear. Devising a policy that drops them may be premature – they may have some legitimate use which the administrator doesn't understand. It is therefore a good practice to observe them before considering a policy adjustment. This moment of rule processing is perfectly right for this purpose: all traffic that shall be dropped is dropped already.

Example of such is shown in Figure Rule Example: Report Calls of Unregistered Users. This rule reports on all non-REGISTER requests for users who have not registered previously. This may be perfectly reasonable for a peering trunk and quite suspicious for a residential user. Gathering these diagnostic events puts an administrator in position to analyze the traffic and create well-targeted policies.

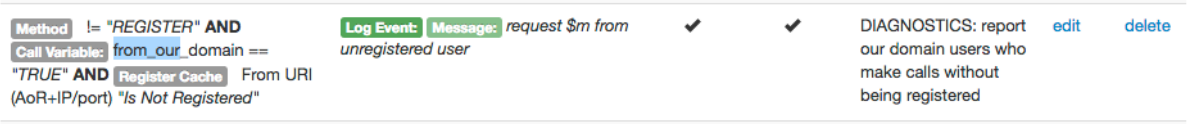

Figure 178: Rule Example: Report Calls of Unregistered Users

Processing Legitimate Traffic

At this stage of rule processing we have eliminated well-known offending traffic and reported on suspicious traffic. It is time to devise rules that process the traffic considered legitimate: mediation rules, media processing rules, topology hiding, etc. The most important fact for sake of this Section is placement of these rules: they are placed in the very end of a rule-base after all other checks have eliminated unwanted traffic.

Figure Rule Example: Processing Legitimate Traffic shows such rules: they implement registration caching and media anchoring to facilitate NAT traversal and off-load the infrastructure behind the ABC SBC . These two rules also contribute to topology hiding: use of media relay hides the actual RTP receivers and registration caching hides the registered contacts.

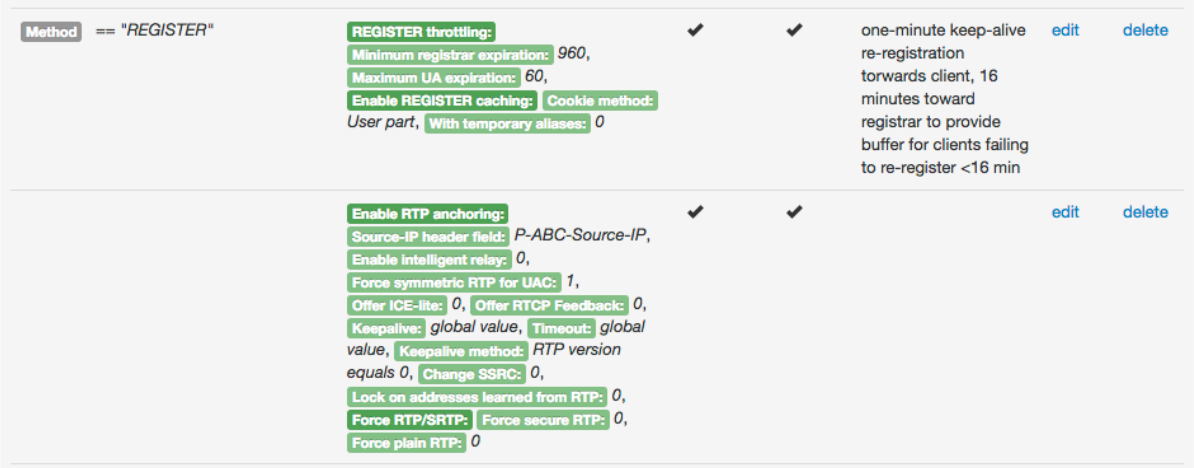

Figure 179: Rule Example: Processing Legitimate Traffic

4.14 Preview of Experimental Features

This chapter summarizes features that are scheduled to appear in future releases and may, under circumstances, become available earlier in experimental releases. The availability, maturity and scope of the features is subject to change without prior notice. Consult FRAFOS professional services if you wish to inquiry about use of these features.

4.14.1 Using Two-Factor Authentication for Users

Two-factor authentication (2FA) is a new experimental feature that helps to preserve security of a whole VoIP system even when security of a component is compromised. What sometimes happens is that PBX passwords leak in various ways and stolen passwords are used to make fraudulent calls that appear legitimate. The 2FA system allows to manage "shadow passwords" for SIP users. If a user's account begins to show irregular patterns, identity of the user can be verified using this shadow password. The shadow password is a short string of digits (PIN) which is stored internally at the SBC in parallel to user's SIP credentials.

The system works as follows. On his or her first attempt to make a call, a user is challenged to enroll by submitting a PIN code using DTMF. The user must remember the PIN code for future verification. Subsequent calls work as normally as long as the status of the user doesn't change. The status can be changed manually from "ok" to "soft-limit" by the administrator or an automated tool such as the FRAFOS ABC Monitor. When a user attempts to initiate a call in the "soft-limit" status, he will be challenged to prove his identity using his PIN code. If the user fails to submit the proper PIN code, his status will change to "hard-limit" and further calls will be blocked using an announcement. Otherwise the verification timestamp will be stored and the user will not be prompted anymore for some convenience period of time.

This basic scenario documented below is programmed using ABC rules and can be adjusted to the needs of a specific scenario.

Prerequisites

For the system to work, the following preparations must be made:

- the ABC SBC must be up and running in cloud configuration with a designated configuration master. This allows changes of user status to propagate to all attached SBCs.
- An administrative account and password must be known for use by SBC to manipulate the user status. Ideally a special user is created for this purpose using **"System ▸ Users ▸ Create User"** . In our examples below, we are assuming a user **rpcuser** with password **rpcrpc**.
- the PIN database must be created. To create the database, use **"Tables ▸ Add New Table"** on the configuration master. Make sure you choose **2FA** as the table type.

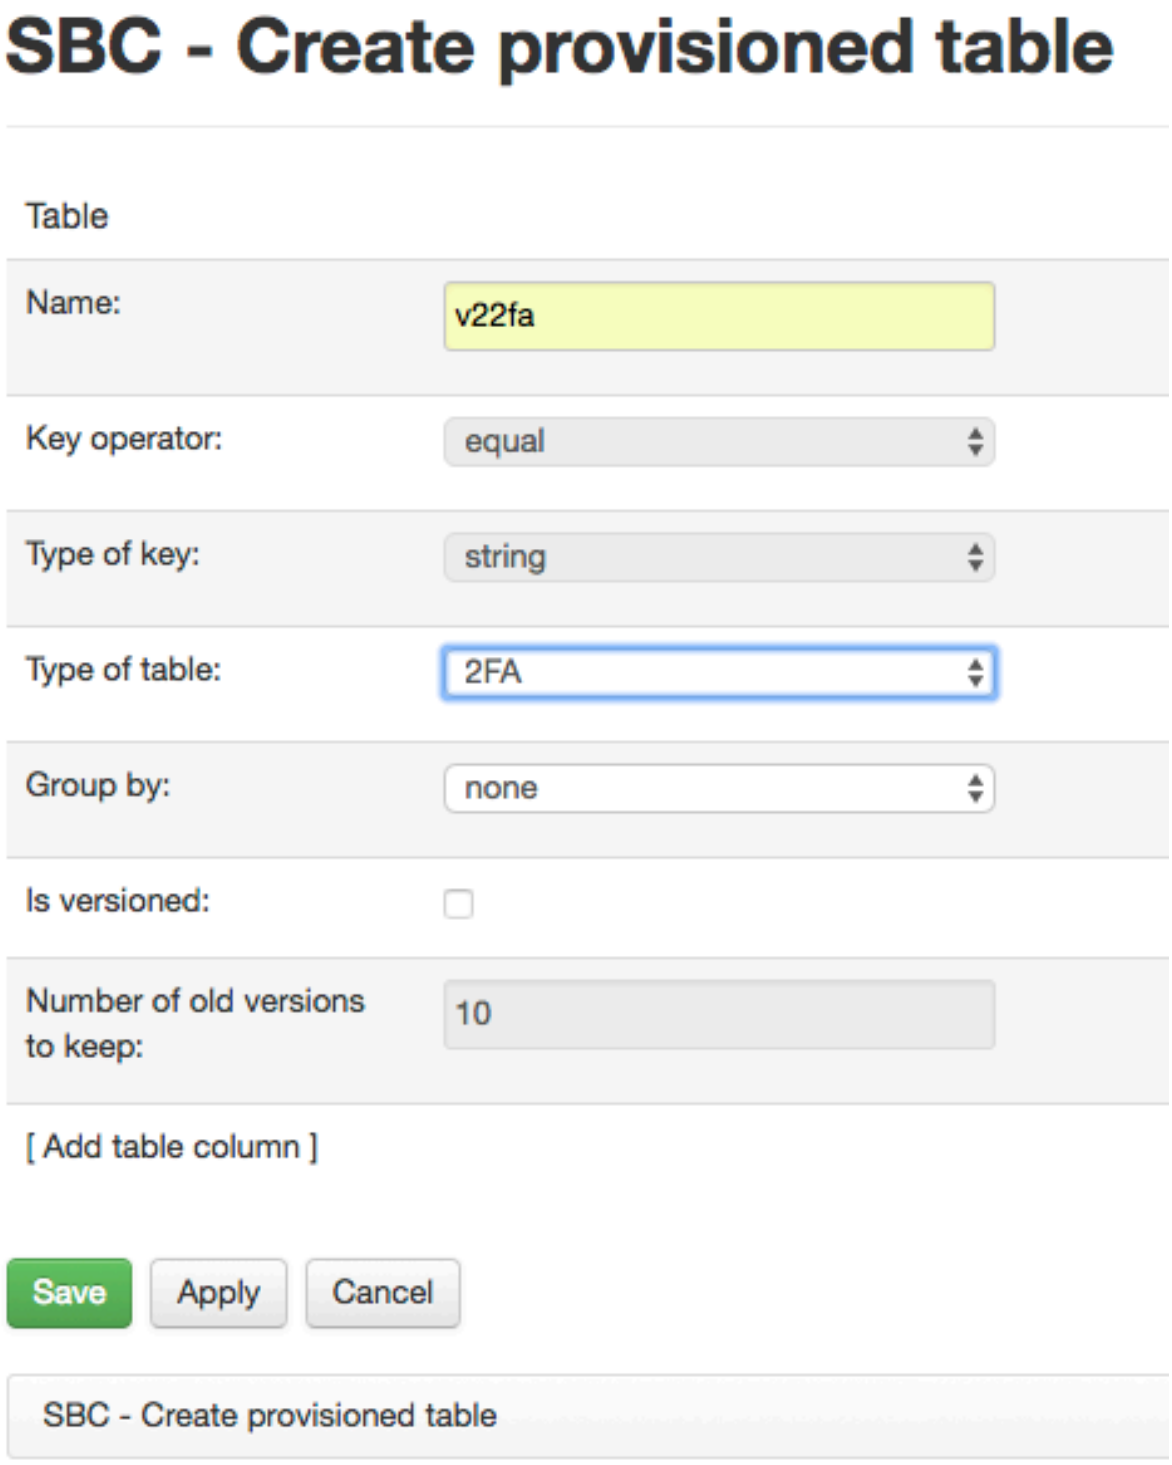

Figure 180: Two factor authentication: Add a New PIN Table

- a system is required to manipulate user status. This can be done manually by editing the provisioned table, or by this-party tools using XML-RPC, or by deploying an extension to FRAFOS ABC Monitor. An example of XML-RPC use is shown below.

```
#!/usr/bin/python
import xmlrpclib
import ssl
if hasattr(ssl, '_create_unverified_context'):
  ssl._create_default_https_context = ssl._create_unverified_context
# find IP address of config master: grep MASTER /data/sbc/etc/sbc-pullconf.conf
servernew = xmlrpclib.Server('https://rpcuser:rpcrpc@192.168.0.83:1443/rpc.php')
data = {"key_value":"sip:3@abc.com","status":"soft-limit","pin":"1111"};
print servernew.tables.insert_update_rule('twoFA',data);
```

- a loopback CA bound to a loopback interface must be setup. The 2FA is running on a loopback interface so that multiple realms can use the same logic by forwarding traffic to loopback, and the 2FA logic doesn't interfere with the actual realms rules. The following screenshots show how to set up the signaling interface, media interface and the actual CA. The interfaces must use systems physical "lo" interface. There must be also a realm to which the CA belongs.

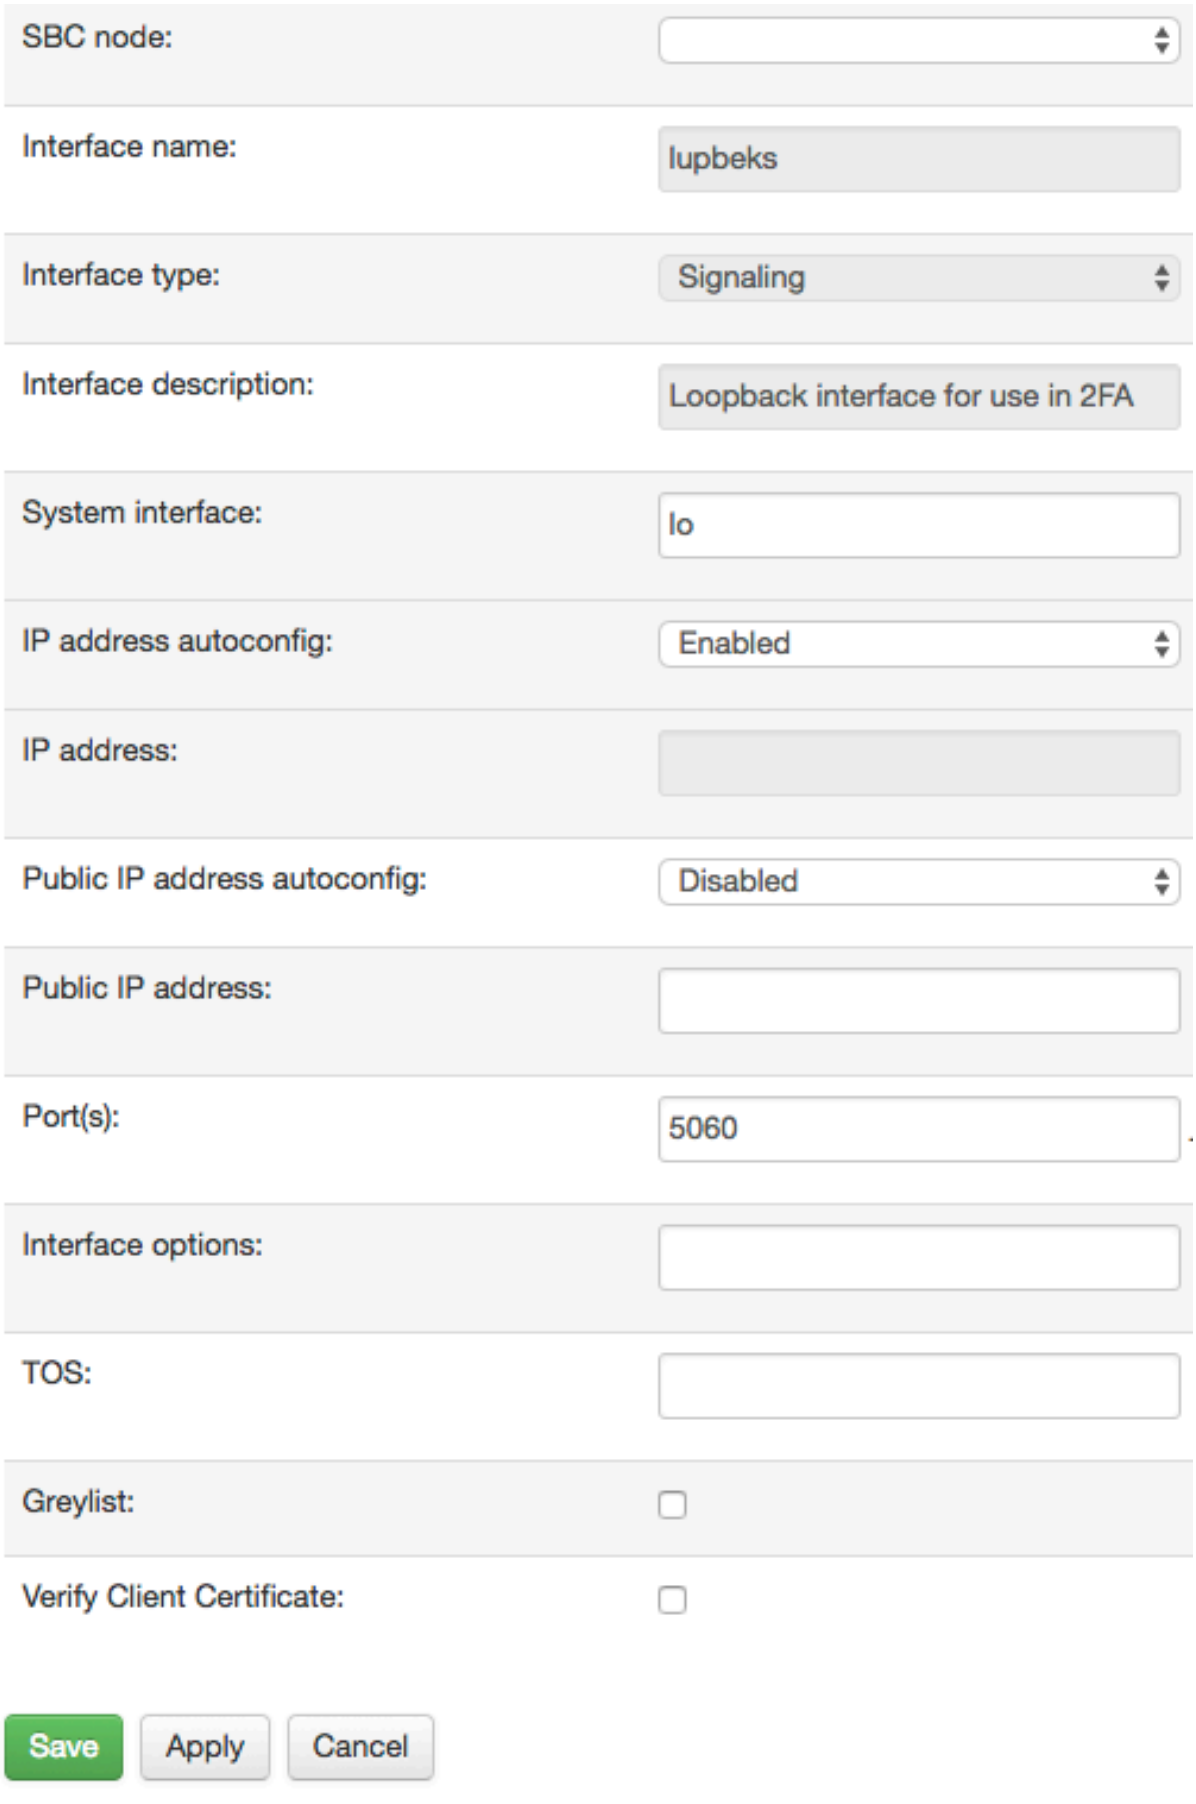

Figure 181: Loopback signaling interface for two factor authentication

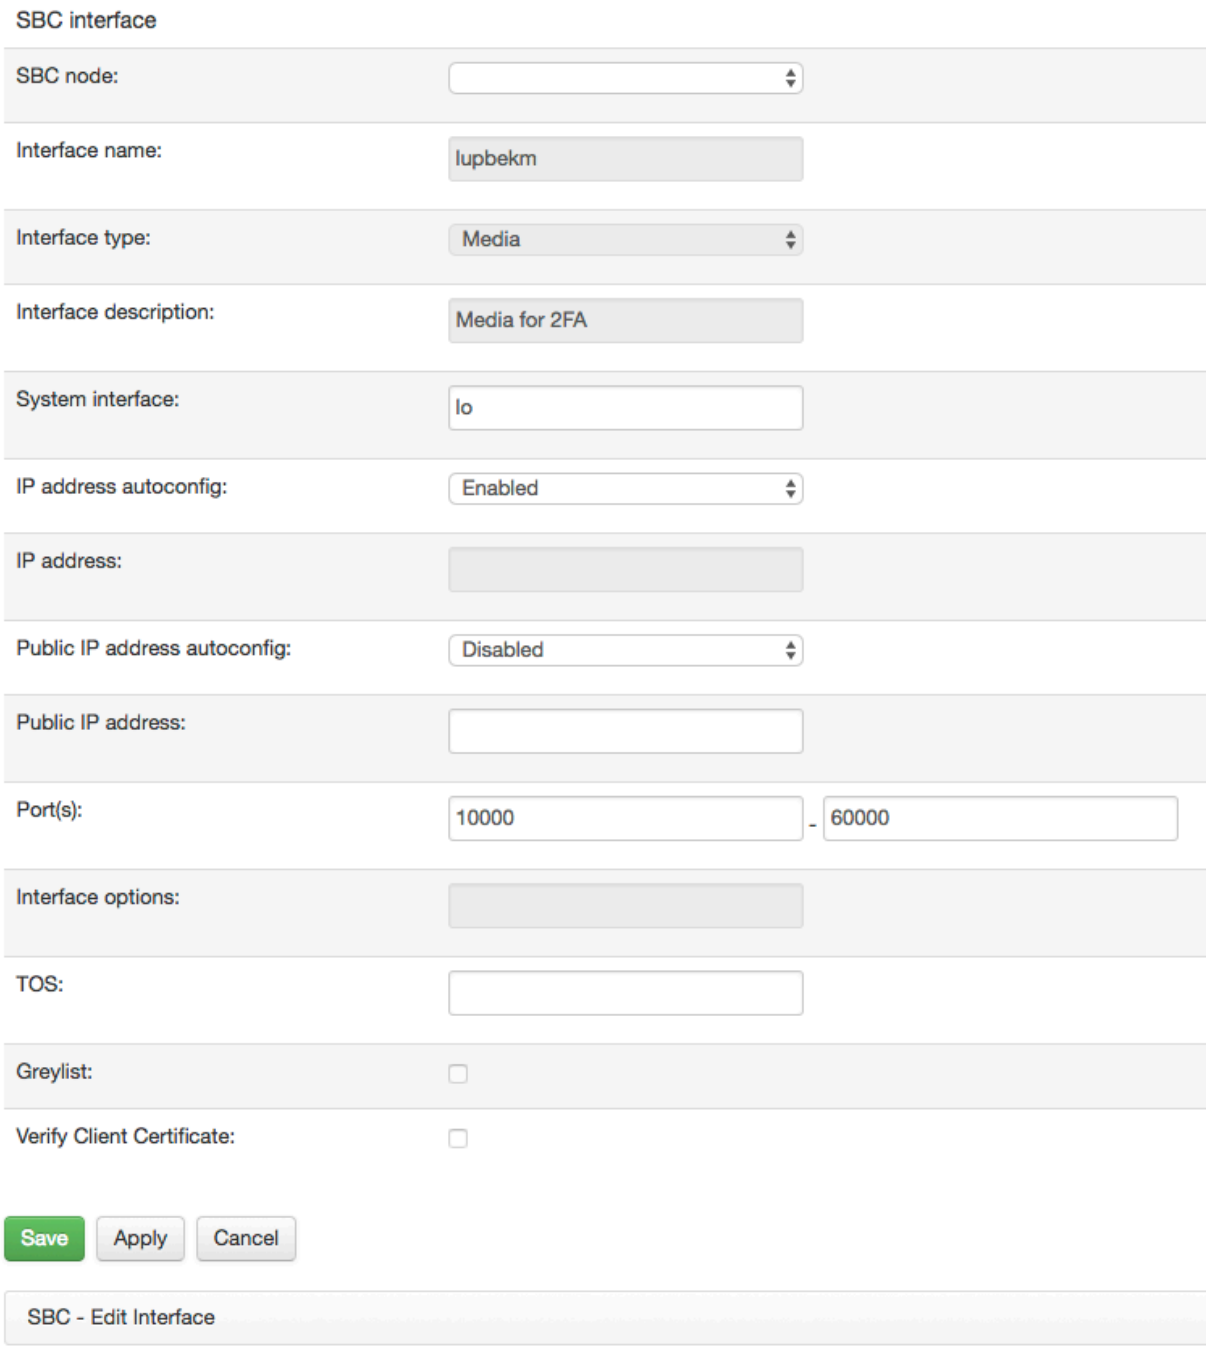

Figure 182: Loopback media interface for two factor authentication

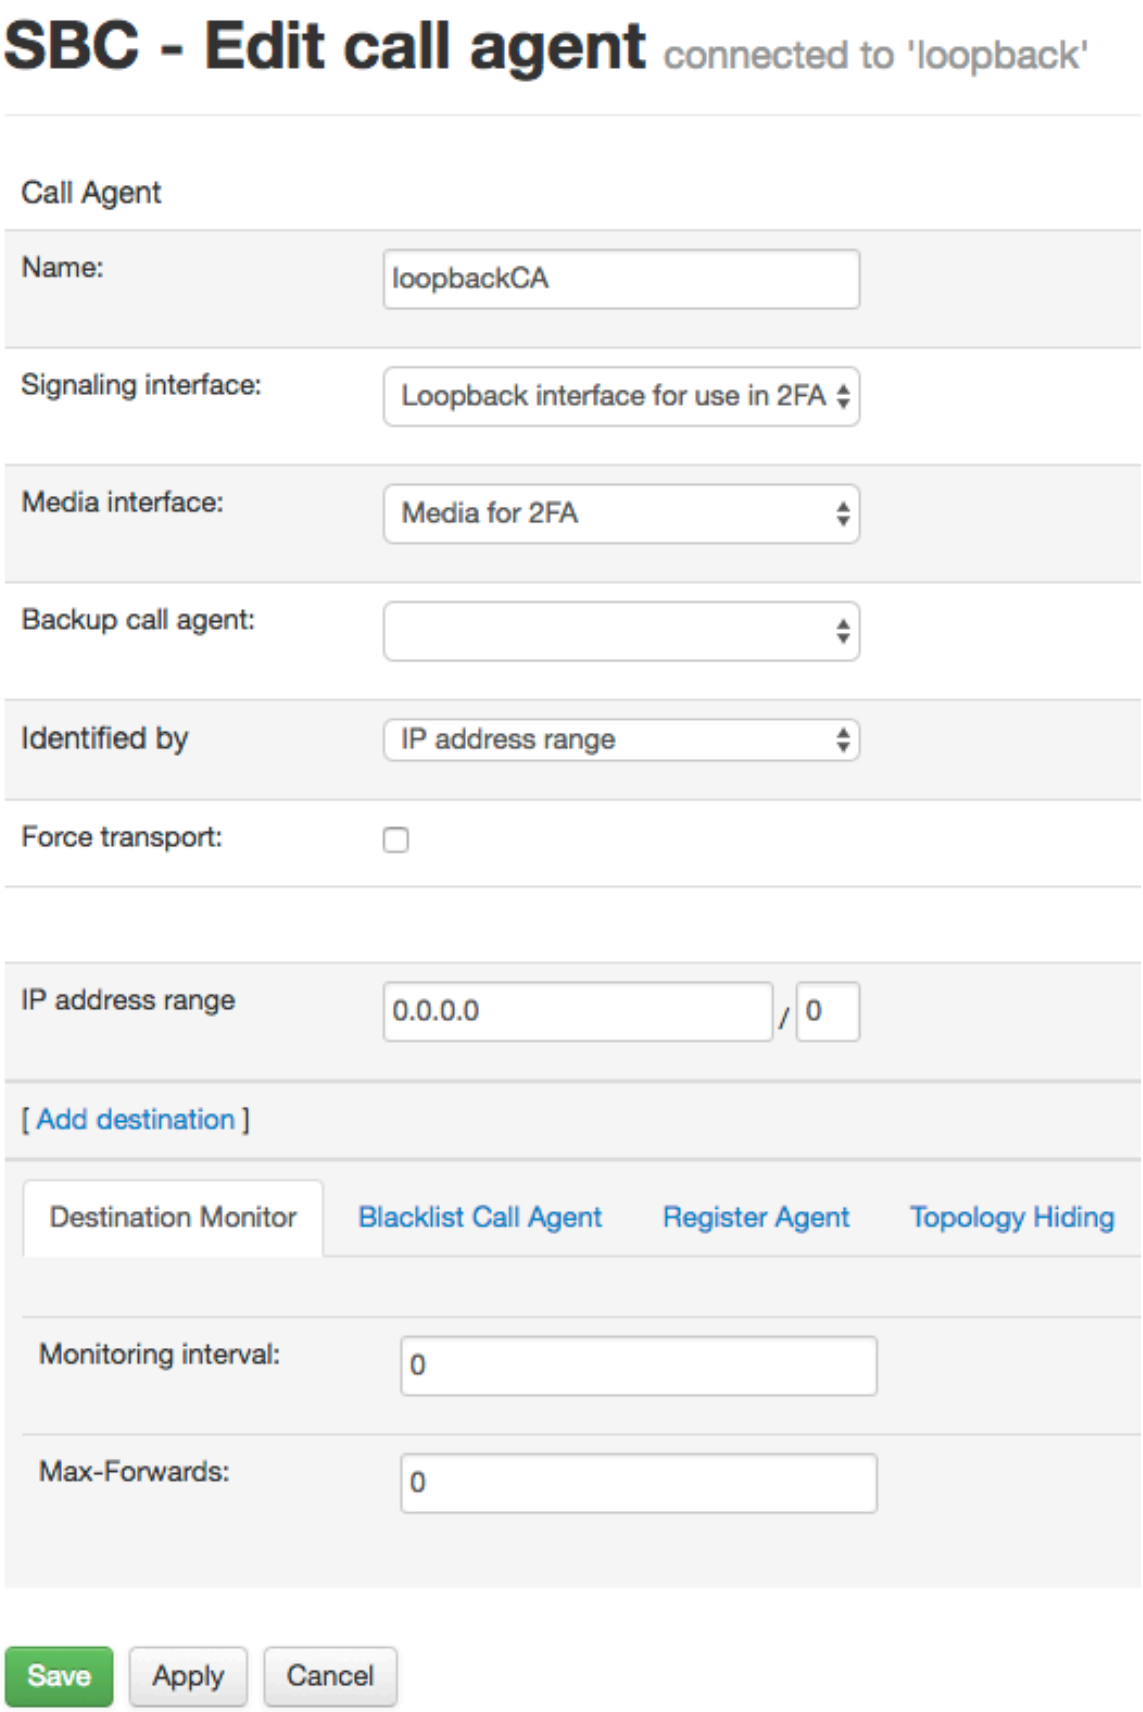

Figure 183: Loopback Call Agent

Rules for Two Factor Authentication Processing

As mentioned bellow, the actual rules for two factor authentication processing will be tied to a loopback interface. This allows to share the rules for multiple Call Agents, in that the Call Agents forward relevant requests to the loopback interface. This is achieved by rewriting userpart of request URI to indicate the desired action and rewriting the hostpart to the loopback address 127.0.0.1.

The following screenshot shows how the loopback rules are formed. They assume that the user part of the request URI indicates what shall be done with INVITEs for the caller. Whether this is a request from a new user who needs to be enrolled, or a user whose status shall be verified, or a user whose request shall be rejected using a voice announcement.

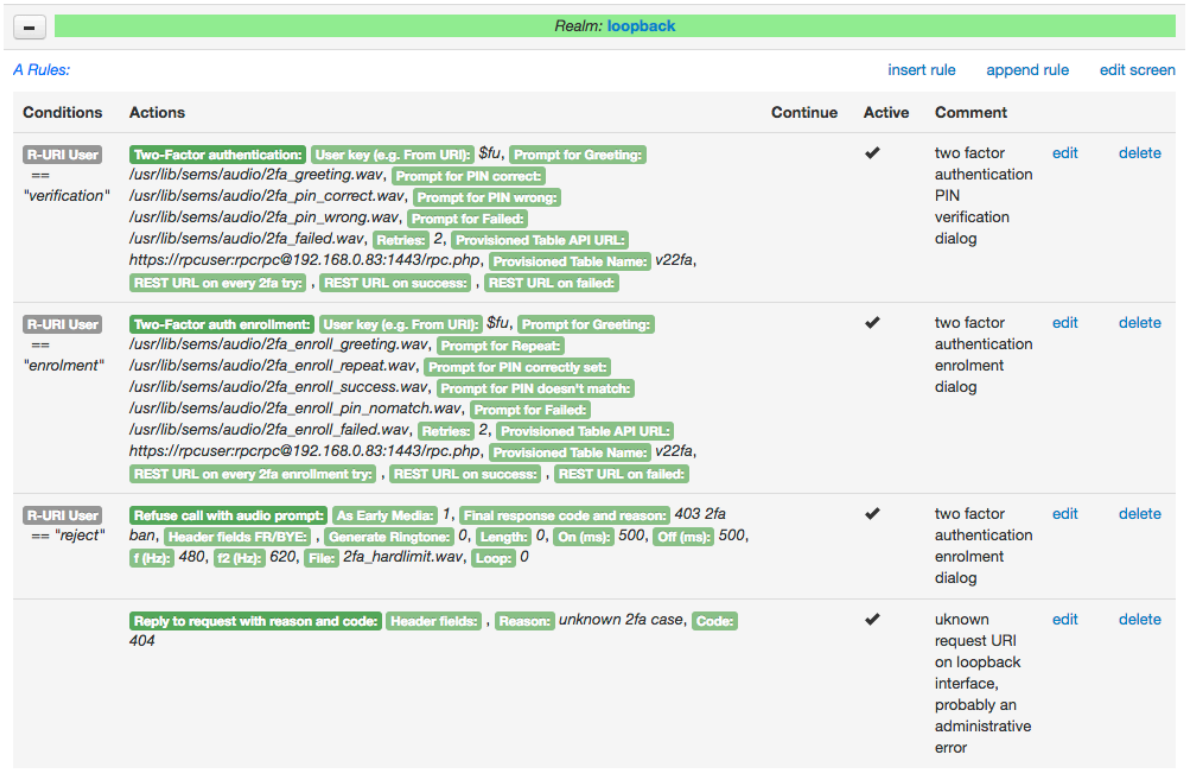

Figure 184: Loopback Rules

There are two actions: **Two factor authentication** which guides an existing user through a verification dialog. If the user types his proper PIN using DTMF, the action updates user's verification timestamp and the user can continue using the service without further disruptions.

The other action, **Two factor auth-enrollment**, prompts a user to submit his PIN. The PIN is then later used to verify identity of the user.

The rule actions have a couple of parameters. The most important parameter is that of the RPC URI – this is the address of the configuration master and legitimate username and password. Both actions update the PIN database based on their completion status. The audio filenames can be changed for a different user experience. RESTful URIs can be optionally used to notify external applications when a rule action finishes with success or failure.

The parameter "Source IP" can be used to set the remote IP address which is recorded with the two factor authentication record in the database. As the processing is executed after being looped on a loopback interface, the real remote IP may be passed on, e.g. by adding a header 'P-ABC-Source-IP' with the correct remote IP. That can be used here with the value '$H(P-ABC-Source-IP)'.

Rules for determining User Status and discriminating by it

The more sophisticated part of the rules discriminates how to treat respective users. An example of such is shown in the following screenshot.

The very first rule retrieves the user status from the current database. The table attributes, such as PIN and status, are then available for processing as call variables.

The first condition selects users whose status is set to “hard-limit”. In that case, incoming INVITEs are forwarded to the loopback interface with “reject” in userpart of request URI, and rejected from there.

The next rule targets users for whom no records are available. Their INVITEs are forwarded to loopback for enrollment.

Subsequent conditions try to determine whether the user is in the **soft-limit** status, and how recently he has verified his identity. If the last verification is too long ago (24 hours in this example), the INVITE is forwarded to the loopback interface for PIN verification.

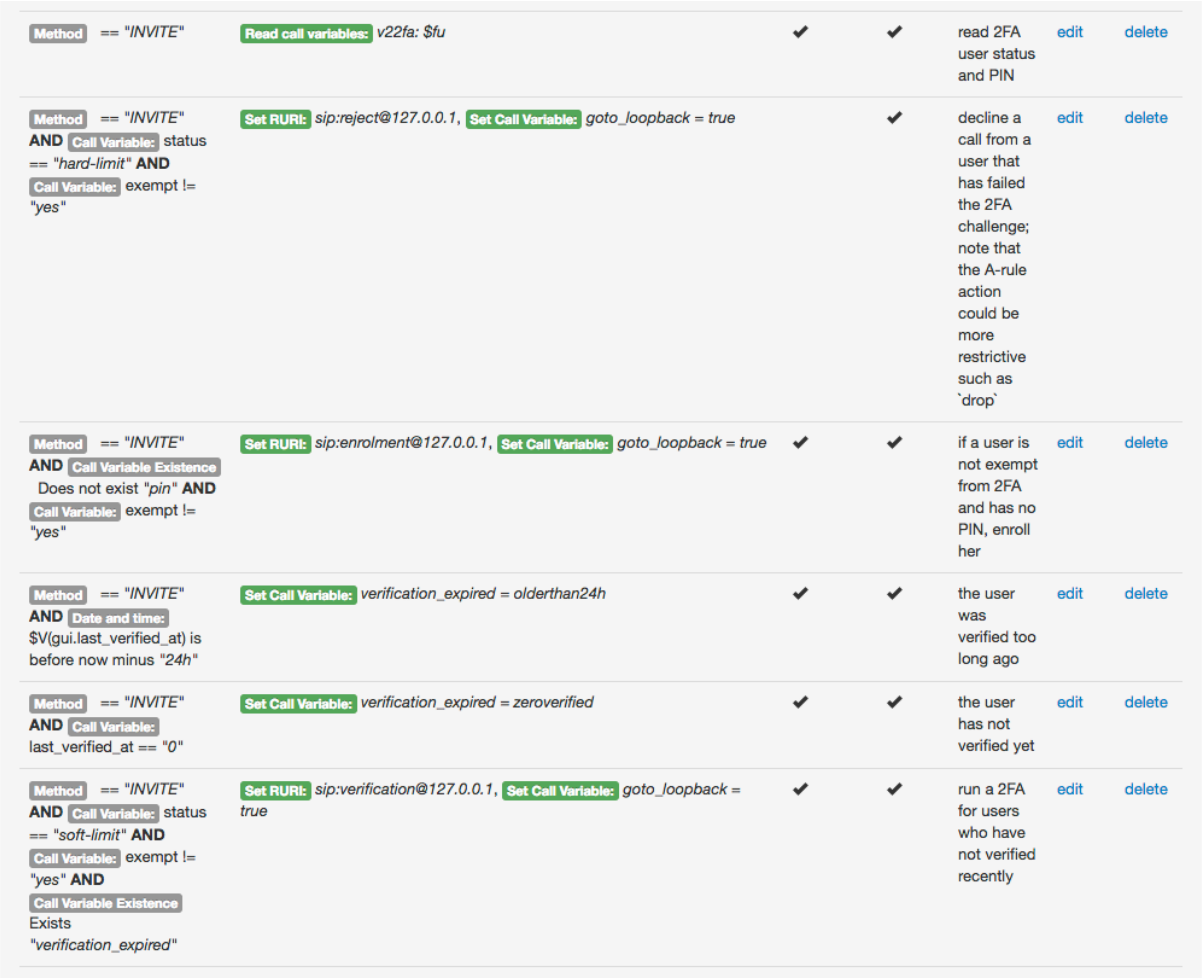

Figure 185: User Discrimination in Two factor authentication Ruleset

If none of these conditions matched, the rule processing continues “as usual”. That’s the case when user status is “ok” or if the user’s identity was verified recently.

Routing Rule to Connect Two Factor Authentication Processing and User Discrimination

There is one remaining piece to connect the user discrimination and 2FA processing rules: a routing rule. The user discrimination rules have already set the loopback address in request URI and defined a variable "goto_loopback". We still need to act upon these. This is fortunately easy to set up:

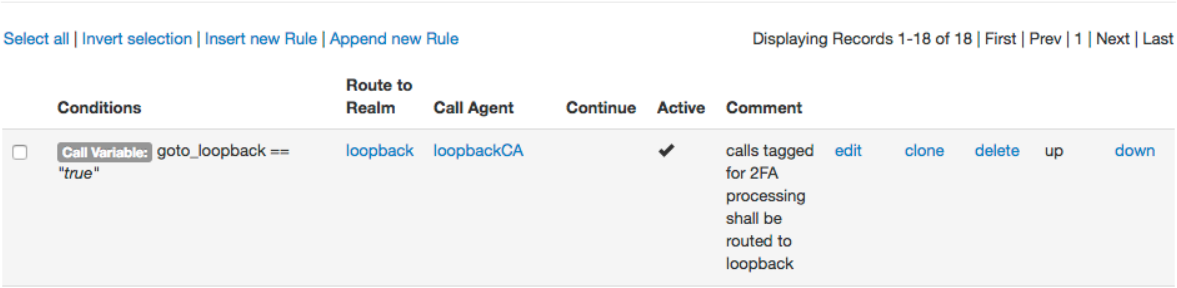

Figure 186: Two Factor Authentication Routing Rule

Scenario Modifications

The two-factor authentication scenario can be modified in many different ways using the ABC rules. The period for which a user doesn't have to be re-validated may be extended or shortened. The IP address of a request can be checked against the IP address from which the identify was verified the last time (last_verified_from_ip) – then a successful verification only validates calls from the same IP address. The way calls from a user in *hard-limit* status are declined can be changed. Note however, that the 2FA application is still in experimental status and untested scenarios may or may not work as expected.

4.14.2 AWS: Reputation Lists

Monitor-steered firewalling allows to combine Monitor intelligence about misbehaving users and ABC SBC's capability to filter out their traffic before it causes harm. It is based on the notion of reputation list and works as follows: an ABC Monitor collects events from all associated SBCs as usual. It uses its data-streaming logic to identify misbehaving traffic sources and posts such to a reputation list. SBCs are subscribed to the list, receive a notification when a new IP address is published by the ABC Monitor, and block the source then.

Use of the ABC Monitor to decide which IP addresses to block is particularly advantageous for several reasons:

- the ABC Monitor collects big data about users and their behavior and is therefore in a very good position to make sophisticated decisions which IP addresses should be blocked.
- the centralized nature of the ABC Monitor allows to convey problematic IP address to all managed SBCs as soon as any of them detects them
- the ABC Monitor has a global view of multiple ABC SBC and can identify misbehaving traffic sources even when they spread their traffic to remain low profile on each managed SBC but their total traffic is beyond a critical mass.

Currently, the reputation list is facilitated using AWS Simple Notification Service (SNS). It is not necessary for the ABC Monitor and ABC SBC to run on AWS but both must have access to AWS SNS service.

Setting Up ABC SBC for Use of Reputation List on AWS

Before you begin, the following prerequisites must be set up in AWS and in the ABC SBC configuration:

- In AWS, there must be an SNS topic, to which the Monitor is allowed to write and from which the ABC SBC is allowed to read.

- AWS Identity must be properly configured on the ABC SBC under **"Global Config ▸ AWS"**. Set AWS region for the SNS, and key id and secret key for identity that can subscribe to the SNS topic. We recommend that you set up a special IAM user with privileges limited to receiving SNSs for this purpose.
- On the ABC SBC, an XMI interface must be properly set up and run on an IP address reachable by SNS. Private IP address not connected to AWS via a VPN will be unreachable for the notifications and the subscription will fail. If the ABC SBC is running on AWS, the option "Public IP address autoconfig" must be therefore set to "Amazon Method".
- RESTful interface for processing notifications must be enabled on the XMI interface. To do so, choose the XMI interface name under **"Global Config ▸ Misc ▸ RESTful interface XMI name "**. Make sure that the port number under **"Global Config ▸ Misc ▸ RESTful interface XMI port"** is open in the SBC's security group.

To subscribe to the SNS, find **"System ▸ Firewall ▸ Subscription to AWS Notification Service"**, click **"Subscribe"** and include the SNS topic's ARN. After the subscription is successfully completed, the IP addresses learned from the reputation list appear under "External FW blacklist".

Setting Up ABC Monitor for Use of Reputation List on AWS

- In AWS, there must be an SNS topic, to which the Monitor is allowed to write and from which the ABC SBC is allowed to read.
- ABC Monitor instance must be assigned a proper IAM role to publish SNS messages.
- ABC Monitor instance must be configured to post to the topic identified by its region and ARN. The settings are under **"Settings"**. In this configuration section, it is also possible to set threshold for Exceeded Limits that may eventually cause a source address to be published on a reputation list.
- The administrator must choose which type of Exceeded Limits will place a source address on the reputation list. To do so turn on/off checkboxes under **"SNS Settings"**.

4.14.3 Server Transaction limits

The server transaction limits allows for limiting the number of running SIP server transactions (UAS transactions). This mechanism, when properly configured, offers a very effective protection against burst of new transactions typical for denial of service attacks as well as against resource exhaustion (mostly RAM) on sustained high SIP traffic or flooding attacks.

These limits will allow the administrator to be warned (events are generated toward the ABC Monitor) when certain limits are passed (soft limits) and limits to be enforced by rejecting new transactions (hard limits) with `503 Overloaded`. In case requests are actively rejected, ABC Monitor events are sent as well.

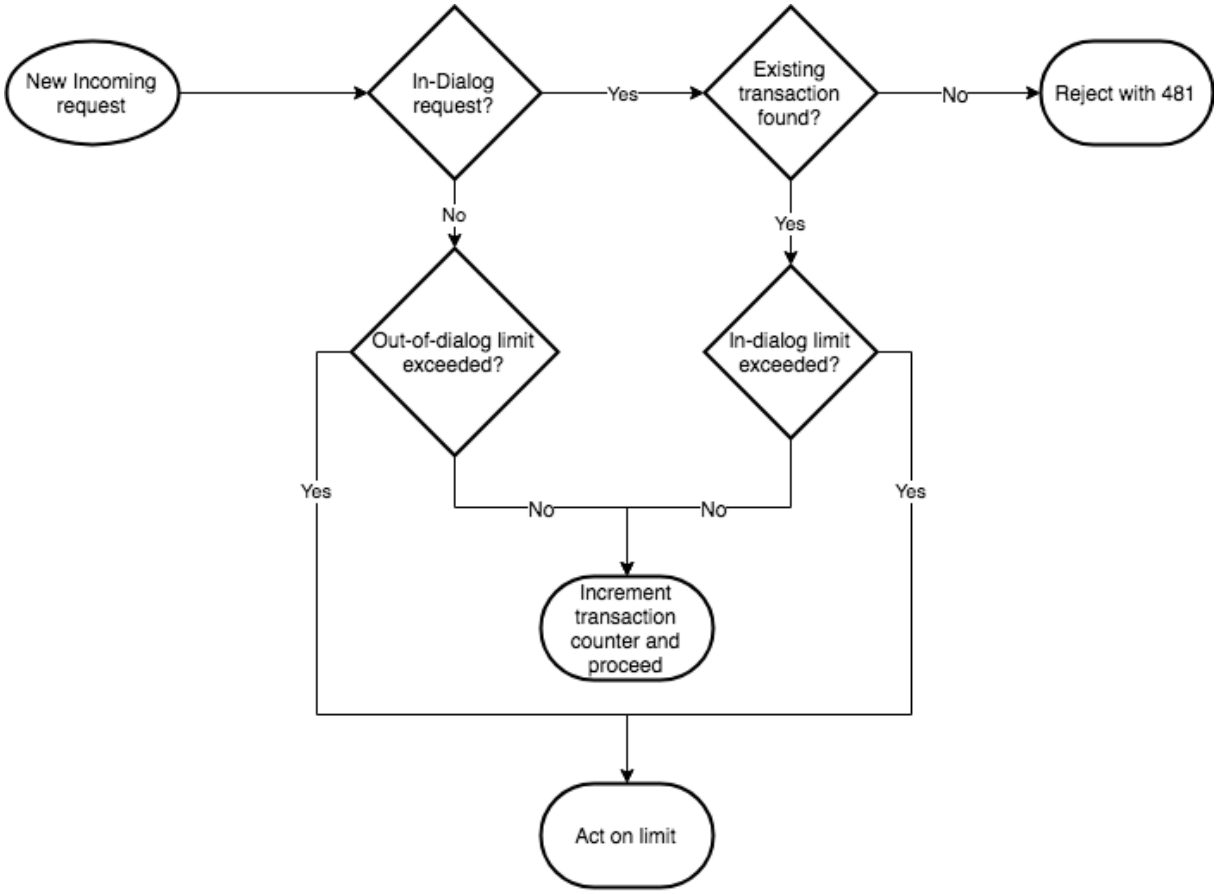

Figure 187: Transaction limit flow diagram

The diagram above shows the behavior related to transaction limits when a new request (not a retransmission) is received. The action taken on a limit breach depends on the type of limit (soft vs. hard limit), as described above.

It is important to note that the very same transaction counter is used to check both types of limits (in-dialog and out-of-dialog), so that the in-dialog limits must necessarily be higher than the out-of-dialog limits. The difference between both is the guaranteed number of in-dialog transactions that can be held.

Soft limit for out-of-dialog transactions (event logging only, use 0 to disable):	0
Hard limit for out-of-dialog transactions (event + reply 503, use 0 to disable):	0
Event throttling for soft/hard OOD limit (in seconds, use 0 to disable):	0
Soft limit for in-dialog transactions (event logging only, use 0 to disable):	0
Hard limit for in-dialog transactions (event + reply 503, use 0 to disable):	0
Event throttling for soft/hard DLG limit (in seconds, use 0 to disable):	0

Figure 188: Transaction limit settings

The transaction limit settings can be found in the global config parameters, in the category “SEMS”.

Setting proper limits

The easiest method for setting proper limits is to monitor the number of UAS transactions while the SBC is operating under normal conditions with the ABC Monitor and to apply a factor of at least 2 to these numbers before setting limits.

For any of the transaction limits fields, a value of 0 means that the limit is deactivated.

In order to start using the limits without impairing production traffic, the soft limits should be set first, together with the event throttling to avoid generating too many events.

Once the soft limit give satisfactory results, meaning that events are generated only on significant load peak, the hard limits can be with a safe margin (at least +30%).

The in-dialog limits should be set very carefully, as it impacts greatly the stability of the system. In particular, BYE requests could be lost in case the in-dialog transaction limit is set improperly.

4.14.4 New restify CDR process

For information about the legacy behavior, please refer to Call Data Records (CDRs).

Since ABC SBC 4.5, the new CDR-ng feature may be enabled in Cluster Config Manager under `Global config` > `CDRs` > `Enable new version of CDRs (CDR-NG)`. The new process monitors event from a target redis list. If a event type is matched with one from the watch list (`call-end` or `conf-leave` for example), a CDR is generated in a CSV format. The CSV content is based on the event fields, in a format specific to the event type. The CDR is then forwarded to a specific syslog facility.

CDRs Location

Please note that by default, `syslog-ng` is configured to redirect a process's messages from a facility to a target file. ABC SBC default configuration for CDR is the following :

Process	Syslog facility	Target file
restify-cdr	LOCAL1	/data/cdr/cdrNG.log
restify-cdr	LOCAL2	/data/cdr/cdrNGconf.log

CDRs configuration

The configuration file is located in `/etc/frafos/restify-cdr.conf`. By default :

- `call-end`, `call-attempt` and `conf-leave` event are watched
- `call-end` and `call-attempt` CDR are forwarded to the `LOCAL1` facility
- `conf-leave` CDR are forwarded to the `LOCAL2` facility

The following formats are defined by default :

- `classic`: 1-1 call CDRs (`call-end`, `call-attempt`)
  - webconf: `web conference based CDRs` (conf-leave)

CDR Format

classic

- Source Realm (event field: `src_rlm_name`)
- Source Call Agent (event field: `src_ca_name`)
- Destination Realm (event field: `dst_rlm_name`)
- Destination Call Agent (event field: `dst_ca_name`)
- From user part (event field: `caller_user`)
- From host part (event field: `caller_host`)

- From name part (event field: `caller_name`)
- To user part (event field: `callee_user`)
- To host part (event field: `callee_host`)
- To name part (event field: `callee_name`)
- Local tag (ID for call) (event field: `id`)
- Timestamp when the call was initiated (format - 2012-05-04 02:22:01) (event field: `start_tm`)
- Timestamp when the call was connected (format as above) (event field: `connect_tm`)
- End Timestamp of the call (format as above) (event field: `end_tm`)
- Duration from start to end (sec.ms) (event field: `duration`)
- Duration from start to connect/end (for established/failed call; sec.ms) (event field: `setup_duration`)
- Duration from connect to end (for established call; sec.ms) (event field: `bill_duration`)
- SIP R-URI (event field: `sip_req_uri`) **note**: the field differ from the original CDR
- SIP From URI (event field: `sip_from_uri`)
- SIP To URI (event field: `sip_to_uri`)

webconf

- Conference identifier (event field: `conf_id`)
- Participant identifier (event field: `participant_id`)
- Call identifier (event field: `call-id`)
- Timestamp when user joined the conference (event field: `ts-join`)
- Timestamp when user leaved the conference (event field: `ts-leave`)
- Duration from start to end (sec.ms) (event field: `duration`)
- From (event field: `from`)
- Call local tag (if one) (event field: `local_tag`)
- Remote URI (event field: `r-uri`)
- Caller source IP (event field: `source`)
- Caller source port (event field: `src-port`)
- Callee (event field: `to`)

Tweak

CDR format may be tweaked as needed, by adding / removing fields from the configuration file entry (`/etc/frafos/restify-cdr.conf`). All fields from the linked events are available via config.

Chapter 5

Reference book

5.1 Reference of Actions

The actions are grouped as follows:

SIP Mediation

Request URI manipulation

- Set RURI
- Prefix RURI user
- Set RURI user
- Append to RURI user
- Strip RURI User
- Set RURI Host
- Set RURI Parameter

To/From manipulation

- Set From
- Set From display name
- Set From User
- Set From Host
- Set To
- Set To Display Name
- Set To User
- Set To Host

Contact HF manipulation

- Set Contact-URI user
- Set Contact-URI host
- Set Contact-HF parameter whitelist/blacklist
- Set Contact-URI parameter whitelist/blacklist
- Forward Contact-HF parameters
- Forward Contact-URI parameters
- Keep Contact user
- Add Dialog Contact Parameter
- Add Contact header field parameter

Authorization

- UAC auth
- UAS auth

Common header manipulation

- Remove Header
- Add Header
- Replace header value
- Replace header value (on leg)
- Insert or Replace header (on leg)
- Set header whitelist
- Set header blacklist
- Update Supported header
- Update Require header
- Update Allow header
- Replace URI header user
- Replace URI header host
- Replace headers of URI header
- Insert or replace headers of URI header
- Diversion to History-Info
- Set Max Forwards
- Map Replaces header
- Forward Via-HFs
- Add X-Org-ConnID header

Session timers

- Enable SIP Session Timers (SST) - caller leg
- Enable SIP Session Timers (SST) - callee leg

Others

- Absorb Re-INVITEs (on leg)
- Absorb UPDATEs (on leg)
- Set SIP Timers
- Handle INVITE with Replaces header

- Relay 503 Reply (on leg)
- Reply In-Dialog Request (on leg)
- Translate Reply Code
- Enable transparent dialog IDs
- Call transfer handling
- Pin TLS Certificate To Dialog (on leg)
- Set Content Type whitelist/blacklist
- Insert or Replace SIP Message Body (on leg)
- Replace SIP Message Body (on leg)

SDP Mediation

- Set CODEC Whitelist
- Set CODEC Blacklist
- Set CODEC Preferences
- Set SDP attribute whitelist
- Set SDP attribute blacklist
- Set SDP bandwidth limit
- Set Media whitelist
- Set Media blacklist
- Drop early media
- Drop SDP from 1xx replies
- Insert or Replace SDP Session Attribute (on leg)
- Replace SDP Session Attribute (on leg)
- Insert or Replace SDP Media Attribute (on leg)
- Replace SDP Media Attribute (on leg)
- Disable SDP Media
- Remove SDP Media Attribute (on leg)
- Insert or Replace SDP Payload Attribute (on leg)
- Replace SDP Payload Attribute (on leg)
- Limit telephony event list (on leg)
- DTLS Setup Preference (on leg)

Monitoring and Logging

- Increment custom counter
- Log received traffic
- Log Event
- Set log level
- Log Message
- Log Message for Replies
- Log to grey list
- Disable privacy monitor mode

Traffic Shaping

- Limit parallel calls
- Limit CAPS
- Limit Bandwidth per Call
- Limit Bandwidth
- Set call Timer

Media Processing

- Enable RTP anchoring
- Restrict media IP to signaling IP (on leg)
- Force RTP/SRTP
- SRTP Fallback to RTP (on leg)
- Activate audio recording
- Activate transcoding
- Process RTP Header Extension
- Join meet-me conference
- Meet-me conference set PIN
- Refuse call with audio prompt
- Play prompt on final response
- Generate Ring-Back Tone
- Activate Music On Hold

DTMF handling

- Convert DTMF to AVT RTP
- Convert DTMF to SIP INFO
- Activate Inband DTMF Detection
- DTMF Termination Same SSRC (on leg)
- DTMF Termination Stable Duration Increments (on leg)

SIP Dropping

- Reply to request with reason and code
- Drop request
- Allow unsolicited NOTIFYs

Scripting

- Set Call Variable

Register Processing

- Enable REGISTER caching
- Retarget R-URI from cache
- REGISTER throttling
- Save REGISTER contact in registrar
- Restore contract from registrar

External Interaction

- ENUM query
- Read call variables over REST
- Read call variables from table

NAT Handling

- Enable dialog NAT handling

Other

- Support serial forking proxy
- Fork

5.1.1 SIP Mediation

Set RURI

Set request URI of the outgoing request to a new value.

Can be used in A and C rules.

> **Info**
>
> This action affects outgoing, dialog initiating request only. In-dialog requests in both directions follow SIP protocol and use content of remote peer's Contact header for building the request URI.

> **Warning**
>
> Using an invalid value will lead to processing error and outbound request wouldn't be sent. "Parser failed on generated request" error will be logged in SEMS log in such case.

Parameters

new URI

New value of request URI.

Accepts replacement expressions.

Prefix RURI user

Prefix user part of request URI.

The values are cumulated thus using this action twice will lead to adding two prefixes.

Adding a prefix (for example `AA`) to an URI without username part (`sip:domain.com`) will create the user part (`sip:AA@domain.com`).

Can be used in A and C rules.

Parameters

prefix string

Prefix that should be prepended to the user part of request URI.

Accepts replacement expressions.

Info

This action affects outgoing, dialog initiating request only. In-dialog requests in both directions follow SIP protocol and use content of remote peer's Contact header for building the request URI.

Warning

Using value that will break R-URI syntax will lead to processing error and outbound request wouldn't be sent. "Parser failed on generated request" error will be logged in SEMS log in such case.

Set RURI user

Replace user part of request URI.

Parameters

new user part

Append to RURI user

Add a suffix to user part of request URI. The result is accumulated if actions is used multiple times.

Parameters

suffix

Strip RURI User

Remove leading characters of user part of request URI.

Parameters

number of leading characters

Set RURI Host

Replace hostport part of request URI.

Parameters

new host part

Set RURI Parameter

Set request URI parameter.

Parameters

parameter name

parameter value

Set From

Replace From Header Field Value.

Parameters

From HF value

Set From display name

Replace From Display name.

Parameters

new From Display name

Set From User

Replace user part of From URI.

Parameters

new From user part

Set From Host

Replace hostport part of From URI.

Parameters

new From hostname

Set To

Replace To Header Field Value.

Parameters

To HF value

Set To Display Name

Replace To Display name.

Parameters

new To Display name

Set To User

Replace user part of To URI.

Parameters

new To user part

Set To Host

Replace hostport part of To URI.

Parameters

new To hostname

Set Contact-URI user

Set the Contact-HF URI user part used for the dialog.

Available since: 4.2

Set Contact-URI host

Override host part of Contact URI used by ABC SBC in appropriate direction.

If this action is not used, ABC SBC uses IP address of appropriate signaling interface (or "Public IP address" if configured) to compose its Contact header. With this action, the host part of generated Contact URI is overridden with the configured value.

Can be used in A and C rules. If used in A rules, it overrides SBC's Contact header in requests or replies (even in-dialog ones) being sent towards caller. If used in C rules, it overrides SBC's Contact header in messages sent towards callee.

Available since: 4.5

> **Warning**
>
> The Contact header field is used by peers to send in-dialog messages to the ABC SBC. If the syntax is broken or if it doesn't point to the appropriate signaling interface, in-dialog messages couldn't be sent by peers (i.e. for example BYE won't be properly delivered and thus calls couldn't be properly terminated).

Parameters

Host

New value of Contact header host.

Replacement expressions and back-references are allowed.

Apply on

Can be used to control on which message type (request, reply or both) the modification is to be applied to.

Available since 5.1.

Only on reply codes

Can be used to control on which reply codes the modification is to be applied to.

If empty, replies with code less than 300 (i.e. provisional and success class responses) are affected.

This is only effective if 'Apply on' is set to a value that will affect replies.

Available since 5.1.

UAC auth

Authenticate on behalf of UAC against an UAS. Any request passing this action and challenged to authenticate by a downstream server will be resent with credentials passed in the action's parameters.

> **Info**
>
> Note that the input fields support replacement expressions. If i.e. password contains special characters such as `$`, they need to be escaped with a backslash.

Parameters

username

password

realm

UAS auth

Authenticate a UAC against the SBC. Either HA1 or password can be provisioned on the SBC; HA1 is safer as the plaintext password does not need to be saved on the SBC. The HA1 can be calculated as MD5(username:realm:password) on the command line. Can be used together with provisioned tables and the Save REGISTER contact in registrar action to create a full registrar.

Parameters

username

realm

H(A1) or password

Remove Header

Removes all occurrences of a header field. The action is applied to initial message, newly added header fields are not removed.

Parameters

header field name

Add Header

Add a new Header Field to a request.

> **Info**
>
> '100 Trying' replies are generated by the SBC. So an action on C-rules with `direction = A leg` will not work on 100-replies because they are not coming from the B-leg. Action on A-rules will work as fine with respect to 100-replies.

> **Info**
>
> Replacement expressions are evaluated once at the beginning of the call (initial request) and the result is re-used throughout the call.

Parameters

HF Name

HF Value

Request or reply

Can be used to control on which type of messages the header will be added on.

Available since 5.1.

Initial or in-dialog

Can be used to choose to only add the header on initial or in-dialog requests, or both.

Available since 5.1.

Direction

Can be used to choose whether to add the header on messages going towards a-leg, b-leg or both.

Available since 5.1.

Replace header value

Replaces matching header field values based on regular expression search and replace.

Parameters

header name

search

replace with

Replacement expressions are allowed, so for example a call variable value may be used here (for example: `$V(gui.fullname)`).

Also, with "replace with", one can use regular expression back-references to use parts of the expression in "match" parameter.

I.e. to replace host part in a header containing a URI, search for `^<sip:([^@]*)@[^?;]*(.*)>` and replace with `<sip:\$1@a.b.c.d\$2>` can be used.

Note that you can only back-reference from 1 to 9 sub-matches, meaning that `\$123` will replace as `<sub-match-1>23`.

Replace header value (on leg)

Same as Replace header value but acts on messages on call leg only.

E.g. putting a rule on A rules of CA1:

```
[CA1] INVITE -> [SBC] -> [CA2] 200 OK -> [SBC] rule-applied -> [CA1]
```

E.g. putting a rule on C rules of CA2:

```
[CA1] INVITE -> [SBC] rule-applied -> [CA2] 200 OK -> [SBC] -> [CA1]
```

Available since: 4.6

Parameters

header name

search

replace with

Insert or Replace header (on leg)

Tries to insert a header field to messages. Unless "replace existing" is enabled, a new header will be added even if a header with the same name exists. If "replace existing" is enabled, the header is replaced with the given value.

Available since: 4.6

Parameters

header name

header value

Replacement expressions and regular expression back-references are allowed.

replace existing

Absorb Re-INVITEs (on leg)

Absorb re-INVITEs coming from the leg if they are considered identical to the previous (re-)INVITE. The decision is done based on:

- All headers match except the following ignored headers: Call-Id, Contact, Content-Length, Content-Type, From, To, Via, RAck, CSeq, Route and Record-Route.
- If the request has a body, the body type is SDP, and the SDP is considered identical (see below).

When SDP is being checked, the SDP of the session is considered. I.e. SDP negotiated via late-oa, or an UPDATE affects this.

For comparing body, if there's application/sdp, only the SDP is taken into account. Within the SDP, only s=, c=, other session-level a=, m= and everything media-level is taken into account. If there is no body on an incoming request, then its body is considered equal to the previous one. If there is a body but the body does not contain application/sdp, then it is considered not-equal to the previous one.

Available since: 4.6.

Parameters

Session-Expires Percentile

If Session-Expires Percentile is set, the INVITE will not be absorbed if the time elapsed has exceeded the set value since the last relayed INVITE. I.e. if percentile is set to 10 and last (re-)INVITE has Session-Expires: 90, then a re-INVITE will be relayed if more than 9 seconds has passed since the last relayed (re-)INVITE even if it is considered identical.

If Session-Expires percentile is empty, then the SBC will absorb re-INVITEs even if they were supposed to refresh the session.

If there is no Session-Expires header in a received UAC Request, then Session-Expires percentile is not checked. This ends up with the same effect as not setting the Session-Expires percentile.

If there was no Session-Expires header in the last received UAC Request and it was not absorbed, then the SBC will not check for the validity of Session-Expires percentile on the following requests. This ends up with the same effect as not setting the Session-Expires percentile.

If Absorb UPDATEs action is also used, then Session-Expires calculations are done in a common way on both INVITEs and UPDATEs.

Ignore Headers

If Ignore Headers is set, then request headers do not affect the decision on absorbing the INVITE or not. This effectively means that only the SDP is compared to previously sent SDPs for equality. Note that the Session-Expires parameter is still honoured if set.

Ignore Body

If Ignore Body is set, then request body does not affect the decision on absorbing the INVITE or not.

Absorb UPDATEs (on leg)

Absorb UPDATEs coming from the leg if they are considered identical to the previous UPDATEs. The decision is done based on:

Parameters and behavior is the same as Absorb Re-INVITEs (on leg).

Note that the first UPDATE will never be absorbed, unless Ignore Headers parameter is enabled. Headers of the UPDATE request are compared separately and the first UPDATE will mark the initial state for the previous headers.

Available since: 5.4.

Relay 503 Reply (on leg)

Normally, per **RFC 3261 Section 16.7**, 503 replies are converted to 500 before sending the reply out to the CA. With this action, 503 replies are relayed to the call leg it is on.

Available since: 5.1.

Reply In-Dialog Request (on leg)

Reply In-Dialog requests matching "Method" (case-insensitive) with a reply with the code "Code".

Parameters

Method

Code

Set header whitelist

Removes all but mandatory and white-listed header-fields.

The list is applied to the final appearance of the INVITE request after all A and C rules have been processed.

Parameters

header-field names

Comma-separated, case-insensitive list of header field names.

> ⚠ **Warning**
>
> compact form needs to be mentioned explicitly!

Set header blacklist

Removes all blacklisted header-fields.

The list is applied to the final appearance of the INVITE request after all A and C rules have been processed.

Parameters

header-field names

Comma-separated, case-insensitive list of header field names.

> ⚠ **Warning**
>
> compact form needs to be mentioned explicitly!

Insert or Replace SIP Message Body (on leg)

Allows inserting or modifying SIP message body based on mime type.

Available since: 5.4

Parameters

Mime-type

Mime type to match. Replacement expressions and back-references are supported. The mime-type `application/sdp` cannot be used here and the action will not be applied if a replacement results in that.

Pattern

RegExp pattern to match. If `Replace with` is enabled, matched part will be replaced with `Value`. If `Replace with` is not enabled, this is ignored. Replacement expressions and back-references are supported.

Value

When `Replace with` is enabled and given mime-type exists, matched part is replaced with the given value. When `Replace with` is enabled and given mime-type does not exist, `Pattern` is ignored and sets the content-type (or makes the message multipart and inserts a new part if the message already has a body) and content to the given value. When `Replace with` is not enabled, `Pattern` is ignored and sets the content-type (or makes the message multipart and inserts a new part if the message already has a body) and content to the given value. Replacement expressions and back-references are supported.

Replace with

When checked, if given mime type already exists, runs a replacement on it instead of inserting. Replacement expressions and back-references are supported.

Replace SIP Message Body (on leg)

Allows modifying SIP message body based on mime type.

Available since: 5.4

Parameters

Mime-type

Mime type to match. Replacement expressions and back-references are supported. The mime-type `application/sdp` cannot be used here and the action will not be applied if a replacement results in that.

Pattern

RegExp pattern to match. Replacement expressions and back-references are supported.

Value

Value to replace the matched part with. Replacement expressions and back-references are supported.

Update Supported header

Allows simplified manipulation with Supported header field content.

Available since: 4.5.

Parameters

operator

Specifies how to use given list of tags.

Add tags

Add the listed tags to the current list of supported tags.

Remove tags

Remove listed tags from the current list of supported tags.

Set tags

Overwrite current list of supported tags with the listed ones.

Whitelist tags

Remove tags that are not listed. Available since: 5.5

comma-separated list of option tags

Update Require header

Allows simplified manipulation with Require header field content.

Available since: 4.5

Parameters

operator

Specifies how to use given list of tags.

Add tags

Add the listed tags to the current list of required tags.

Remove tags

Remove listed tags from the current list of required tags.

Set tags

Overwrite current list of required tags with the listed ones.

comma-separated list of option tags

Update Allow header

Allows simplified manipulation with Allow header field content.

> **Info**
>
> "Add" operator will not add unless Allow header already exists, set via "Set" operator or "Default tags" are specified.

Available since: 4.6.

Parameters

operator (Add / Remove / Set tags)

comma-separated list of option tags

Direction

Apply on

Default tags

Replace URI header user

Allows modifying "user" part on headers containing an URI. I.e. `Refer-to: sip:USER@host`.

Available since: 5.0.

Parameters

Header name

Search

Replace with

Replace URI header host

Allows modifying "host:port" part on headers containing an URI. I.e. `Refer-to: sip:user@HOST:PORT`.

Available since: 5.0.

Parameters

Header name

Search

Replace with

Replace headers of URI header

Allows modifying headers in headers containing URIs.

I.e. Call-ID in `Refer-to: <sip:user@host?Call-ID=55432%40alicepc.atlanta.example.com>` can be manipulated with "header name = refer-to", "name of the header in URI = call-id", "Search = 432@alice", "replace with = 433@bob".

Available since: 5.0.

Parameters

Header name

Name of the header in URI

Search

Replace with

Insert or replace headers of URI header

Allows modifying headers in URI of headers containing a URI.

I.e. NEW-hdr in `Refer-to: <sip:user@host?Call-ID=55432%40alicepc.atlanta.example.com&NEW-hdr=value>` can be added with this.

Parameters

Header to modify

Header name

Header value

Replace if exists

Add Dialog Contact Parameter

Add parameters to the Contact URI generated by the SBC. Acts on all out-of-dialog, dialog-initiating or in-dialog messages.

I.e. `new-param` in `Contact: <sip:user@host;new-param=value>` can be added with this.

Parameters

Leg: A or B parameter name parameter value

Add Contact header field parameter

Add Contact header field parameters to the Contact header generated by the SBC. Acts on all out-of-dialog, dialog-initiating or in-dialog messages.

The given value parameter needs to be already escaped.

Note that this action will not work together with REGISTER messages when `Save REGISTER contact in registrar` or `Enable REGISTER caching` actions are used.

Note that this action will not work on provisional replies when it is in C-rules and the `Leg` parameter is `A`.

I.e. `new-param` in `Contact: <sip:user@host>;new-param=value` can be added with this.

Parameters

Leg: A or B parameter name parameter value

Set Contact-HF parameter whitelist/blacklist

Specify which Contact header field parameters in incoming request to forward downstream.

Parameters

comma-separated list of parameter names

Set Contact-URI parameter whitelist/blacklist

Specify which Contact URI parameters in incoming request to forward downstream.

Available since: 4.6.

Parameters

comma-separated list of parameter names

Forward Contact-HF parameters

Forward all Contact header field parameters “as is” downstream.

Forward Contact-URI parameters

Forward all Contact URI parameters “as is” downstream.

Available since: 4.6.

Keep Contact user

Keep Contact URI user part as received from the other peer in Contact header generated by ABC SBC.

Without this action, ABC SBC generates its Contact URI with username part representing the dialog identifier. If this action is used, the username part from incoming Contact URI is preserved and used in SBC’s Contact URI towards the other peer and new Contact URI parameter `dlg-id` is added and used to identify the dialog instead of the URI username.

Can be used in A and C rules and affects the appropriate call leg only.

If this action is used in A rules, the callee’s username in Contact URI is preserved and sent in Contact header in messages towards caller. For example:

Caller sends INVITE with its Contact header:

```
INVITE sip:104@vku-test.com SIP/2.0
...
Contact: <sip:101@192.168.13.221:6010;ob>
...
```

ABC SBC forwards the INVITE with usual Contact header ("Keep Contact user" is not used in C rules):

```
INVITE sip:104@192.168.13.221:6040;ob SIP/2.0
...
Contact: <sip:21F67A8F-64DF20AB0005698E-923FF6C0@192.168.13.51;transport=udp>
...
```

Callee replies with its Contact:

```
SIP/2.0 200 OK
...
Contact: <sip:104@192.168.13.221:6040;ob>
...
```

ABC SBC forwards the Contact URI username to caller ("Keep Contact user" is used in A rules) and adds `dlg-id` parameter:

```
SIP/2.0 200 OK
...
Contact: <sip:104@192.168.13.51;dlg-id=4D1B3203-64DF20AB00055FCD-B833D6C0;transport=tcp>
...
```

If it is used in C rules, the caller's username is used in Contact header in messages towards callee. For example:

Caller sends INVITE with its Contact header:

```
INVITE sip:104@vku-test.com SIP/2.0
...
Contact: <sip:101@192.168.13.221:6010;ob>
...
```

ABC SBC forwards the Contact URI username to callee ("Keep Contact user" is used in C rules) and adds `dlg-id` parameter:

```
INVITE sip:104@192.168.13.221:6040;ob SIP/2.0
...
Contact: <sip:101@192.168.13.51;dlg-id=386CF1E5-64DF2A70000DDF70-921FD6C0;transport=udp>
...
```

Callee replies with its Contact:

```
SIP/2.0 200 OK
...
Contact: <sip:104@192.168.13.221:6040;ob>
...
```

ABC SBC forwards the reply with usual Contact header ("Keep Contact user" is not used in A rules):

```
SIP/2.0 200 OK
...
Contact: <sip:01E3A1EE-64DF2A70000DD992-B833D6C0@192.168.13.51;transport=tcp>
...
```

Translate Reply Code

Translate SIP reply codes to other value.

Parameters

matching reply code

new reply code

new reason phrase

Set Max Forwards

Reset the number of hops a request can be forwarded to specified value.

Parameters

the new value of Max-Forwards header field

Enable transparent dialog IDs

Enforce use of the same dialog IDs on both sides of a call.

Parameters

To-tag

Controls To-tag handling. Can have following values:

Stick to first received to-tag

Keeps the first seen to-tag in the early responses throughout the rest of the dialog, even if it changes in the final reply.

Re-set to-tag with final reply

Will switch the to-tag from early to established dialog (on first final reply sent to caller).

Forward Via-HFs

Force the SBC to keep the Via header fields while forwarding the request.

Diversion to History-Info

Converts SIP Diversion header-field into History-Info.

Call transfer handling

Defines the mode in which REFERs are handled: rejection, local processing or forwarding.

Parameters

Mode

REFER processing mode. Can be one of

REFER pass-through

Handle REFER internally

Reject REFER

Reconnect on all failures during unattended transfer

Reconnect if transfer ends in 4xx during unattended transfer.

Do not terminate after unattended transfer

Do not terminate referrer leg when the unattended transfer completes.

Only NOTIFY 100 & final sip replies

Disables relaying of provisional replies of transferee to referrer as NOTIFY messages. It can come useful in scenarios where backup CA agent is tried and provisional replies of latter CA might confuse the referrer.

Set SIP Timers

Allows setting SIP timers per call.

Parameters

SIP Timers

Failover reduce factor

This parameter is used to divide B, F & M timers when destination call agent has a backup CA. This allows for a faster failover. Leaving it empty uses the default value of 4.

Handle INVITE with Replaces header

Activates internal processing of INVITE with Replaces header.

Map Replaces header

Activates mapping of dialog identifiers in INVITE with Replaces.

Pin TLS Certificate To Dialog (on leg)

This action causes remembering the initial client certificate that's used while initiating the dialog and rejects any in-dialog request that do not use the same certificate.

This action requires "Verify peer certificate" to be enabled on the TLS Profile of the signaling interface.

Note that non-TLS messages, messages with no associated TLS client certificates or messages with different different certificates compared to the pinned one will be:

- Rejected with 403 if it is an initial request.
- Rejected with 481 if it is an in-dialog request.
- Dropped if it is a reply or an ACK.

When used in A rules:

- If SHA256 fingerprint is empty, then the fingerprint of the certificate used in the initial request is pinned.
- If SHA256 fingerprint is given, then it is pinned for the dialog and the certificate used in the initial request will also be compared against it.

When used in C rules, SHA256 fingerprint must be given.

In order to get the SHA256 fingerprint of a certificate, the following command may be used: `openssl x509 -noout -fingerprint -sha256 -inform pem -in <CERT>`

Available since: 5.2.

Parameters

SHA256 fingerprint

Set Content Type whitelist/blacklist

Specifies which SIP payload types (such as SDP) will be permitted.

Parameters

comma-separated list of content types

Enable SIP Session Timers (SST) - caller leg

Enforce the use of session timers for the caller. Support for session timers is not advertised to the callee (the `timer` extension is removed from the `Supported` header if present) unless the Enable SIP Session Timers (SST) - callee leg action is also used.

Even if the caller does not support session timers, ABC SBC will periodically refresh the session by sending UPDATE or re-INVITE requests to the caller.

If the session timer negotiation results in the caller being responsible for session refreshes, the appropriate session refresh requests will be propagated to the callee unless the Absorb Re-INVITEs (on leg) or Absorb UPDATEs (on leg) actions are used in the caller's call leg.

Parameters

session expiration (sec)

minimum expiration (sec)

let remote refresh

Enable SIP Session Timers (SST) - callee leg

Enforce the use of session timers for the callee. Support for session timers is not advertised to the caller (the `timer` extension is removed from the `Supported` header if present) unless the Enable SIP Session Timers (SST) - caller leg action is also used.

Even if the callee does not support session timers, ABC SBC will periodically refresh the session by sending UPDATE or re-INVITE requests to the callee.

If the session timer negotiation results in the callee being responsible for session refreshes, the appropriate session refresh requests will be propagated to the caller unless the Absorb Re-INVITEs (on leg) or Absorb UPDATEs (on leg) actions are used in the callee's call leg.

Parameters

session expiration (sec)

minimum expiration (sec)

let remote refresh

Add X-Org-ConnID header

The X-Org-ConnID header field contains a unique value that remains constant for the duration of the transaction and any dialog created from this request.

By enabling this action, a X-Org-ConnID header is added to every outgoing initial SIP INVITE request product of this dialog.

The header helps to correlate calls that have been internally redirected (due to a 302 SIP response) or blindly transferred (due to a REFER SIP request).

The value can be retrieved in the CDR by specifying the keyword "$x_org_connid" in the cdr_format (see cc_syslog_cdr.conf).

5.1.2 SDP Mediation

Set CODEC Whitelist

Remove all but listed codecs from SDP.

Parameters

codec list

Comma-separated, case insensitive, list of allowed codecs.

Set CODEC Blacklist

Remove all listed codecs from SDP.

Parameters

codec list

Comma-separated, case insensitive, list of disallowed codecs.

Set CODEC Preferences

Define the order in which available codecs are chosen.

Parameters

comma-separated codec-list

Set SDP attribute whitelist

Removes all but listed SDP attributes from SDP payload.

Parameters

comma-separated list of attribute names

Set SDP attribute blacklist

Removes specified SDP attributes from SDP payload.

Parameters

comma-separated list of attribute names

Set SDP bandwidth limit

Set session bandwidth limit in SDP. Sets (if nonexistent) or limits (if existent) the b= attribute. If Media type is set (e.g.: 'audio' or 'video'), the media-line b= attributes are limited, if Media type is not set, the session level b= attribute is limited. Common limit types: TIAS, AS, CT.

Parameters

Limit Type

Limit type, e.g. TIAS (RFC 3890), or AS, or CT (RFC 4566)

Limit (kilobit per second)

Set media limit for this media type. Only one value is allowed.

Media type

Specify media type, e.g. 'audio' or 'video'

Set Media whitelist

Permit only listed media types.

Parameters

media list

Comma-separated list of enabled media types. For example “audio,video”.

Set Media blacklist

Remove listed media types.

Parameters

media list

Comma-separated list of media types to blacklist. For example “video,image”.

Drop early media

Drop early media (audio only).

Drop SDP from 1xx replies

Drop SDP from listed 1xx replies.

Parameters

list of affected reply codes

Insert or Replace SDP Session Attribute (on leg)

Try to insert a session-level attribute to all requests/replies on call leg. Unless “replace with” is enabled, the insertion will take place even if an attribute with the same name exists. If it’s enabled the value of the attribute with the same name is changed to “Attribute value”.

If the attribute is “known” to the SBC this action can remove other forms of the attribute. I.e. inserting “sendonly” will remove the previous indicator such as “inactive”, regardless of the value of the “Replace with” parameter.

Available since: 4.6.

Parameters

Attribute name

The name to replace.

Supports replacement expressions.

Attribute value

The Attribute value.

Supports replacement expressions and back-references.

Replace with

Replaces if already exists.

Replace SDP Session Attribute (on leg)

Replace an SDP session attribute on all requests/replies on a call leg.

Available since: 4.6.

Parameters

Attribute name

The name to replace, supports replacements.

Search

Regexp to match the part to be replaced.

Replace with

Holds the value to be replaced with. Supports replacement expressions and back-references.

Insert or Replace SDP Media Attribute (on leg)

Try to insert a media-level attribute to all requests/replies on call leg. Unless "replace with" is enabled, the insertion will take place even if an attribute with the same name exists. If it's enabled the value of the attribute with the same name is changed to "Attribute value".

If the attribute is "known" to the SBC this action can remove other forms of the attribute. I.e. inserting "sendonly" will remove the previous indicator such as "inactive", regardless of the value of the "Replace with" parameter.

Available since: 4.6.

Parameters

Attribute name

Name of the attribute to be replaced. Supports replacement expressions.

Media

Regexp matched against the `m=` media lines to select specific ones. Supports replacement expressions and back-references.

Attribute value

The attribute value to be used.

Supports replacement expressions and back-references.

Replace with

Replaces if already exists.

Replace SDP Media Attribute (on leg)

Replace an SDP media attribute on all requests/replies on a call leg.

This action can be used for payload id re-mapping if used with RTP anchor. E.g. attr. name, media, search, replace with values `rtpmap`, `.*`, `^98 XYZ`, `105 XYZ` respectively will replace payload id 98 with 105 in relayed RTP packets.

Available since: 4.6.

Parameters

Attribute name

Name of the attribute to be replaced. Supports replacement expressions.

Media

Regexp matched against the `m=` media lines to select specific ones. Supports replacement expressions and back-references.

Search

Search is a regexp to match the part to be replaced.

Replace with

Holds the value to be replaced with, supporting replacement expressions and back-references.

Disable SDP Media

Disable an SDP media on all requests/replies.

This action can also remove the media line based on the global config option "Remove filtered m-lines".

I.e. in removal of media with payload:

```
m=audio 8012 RTP/AVP 102
a=rtpmap:102 telephone-event/48000
a=content:special
```

"Media" would be compared against `audio 8012 RTP/AVP 102`, "Attribute name" would be compared to `rtpmap` or `content` under that media line, "Attribute value" would be compared against `102 ...` or `special` values.

Available since: 5.1.

Parameters

Media

Regexp matched against the `m=` media lines to select specific ones. Supports replacement expressions and back-references.

Attribute name

Regexp to match an attribute under the `m=` line to be removed. Supports replacement expressions and back-references.

Attribute value

Regexp to match an attribute under the `m=` line to be removed. Supports replacement expressions and back-references.

Remove SDP Media Attribute (on leg)

Remove an SDP media attribute on all requests/replies on a call leg.

I.e. in removal of payload with id 102:

```
m=audio 8012 RTP/AVP 102 103
a=rtpmap:102 telephone-event/48000
a=rtpmap:103 telephone-event/8000
```

"Attribute name" would be `rtpmap`, "Media" would be compared against `audio 8012 RTP/AVP 102`, "Search" would be compared to `102 telephone-event/48000`, and would result in:

```
m=audio 8012 RTP/AVP 103
a=rtpmap:103 telephone-event/8000
```

Available since: 4.6.

Parameters

Attribute name

The name of attribute to remove. Supports replacement expressions.

Media

Regexp matched against the `m=` media lines to select specific ones. Supports replacement expressions and back-references.

Search

Search is a regexp to match the line to be removed.

Insert or Replace SDP Payload Attribute (on leg)

Try to insert a payload-level attribute to all requests/replies on call leg. Unless "replace with" is enabled, the insertion will take place even if an attribute with the same name exists. If it's enabled the value of the attribute with the same name is changed to "Attribute value".

Available since: 4.6.

Parameters

Attribute name

The name of attribute to insert/replace. Supports replacement expressions.

Media

Regexp matched against the `m=` media lines to select specific ones. Supports replacement expressions and back-references.

Codec

Regexp matched against the respective `rtpmap=xyz <CODEC>`. Supports replacement expressions and back-references.

Attribute value

Supports replacement expressions and back-references. I.e. for `fmtp`, it is placed as `fmtp: xyz <VALUE>`.

Replace with

Replaces if already exists.

Replace SDP Payload Attribute (on leg)

Replace an SDP payload attribute on all requests/replies on a call leg.

Available since: 4.6.

Parameters

Attribute name

Name of the attribute to be replaced. Supports replacement expressions.

Media

Regexp matched against the `m=` media lines to select specific ones. Supports replacement expressions and back-references.

Codec

Regexp matched against the respective `rtpmap=xyz <CODEC>`. Supports replacement expressions and back-references.

Search

Regexp to match the part of attribute value to be replaced. I.e. for `fmtp` it is compared against `fmtp:xyz <SEARCH>`. Supports back-references.

Replace with

Replacement value. Supports replacement expressions and back-references.

Limit telephony event list (on leg)

Limit telephony events attribute on all requests/replies on a call leg.

Available since: 4.6.

Parameters

Media

Regexp matched against the `m=` media lines to select specific ones. Supports replacement expressions and back-references.

Telephony events

Comma-separated list such as `0-16,66` that will filter out anything that is not in it.

DTLS Setup Preference (on leg)

This controls whether SBC prefers to be *active* or *passive* for DTLS setup. I.e. when used in A-rules, if the caller signals `actpass` setup, this controls whether the SBC prefers to respond with `active` or `passive`. When used in C-rules, this can be used to configure the SBC to send `active` or `passive` instead of `actpass`.

This action is only meaningful when the RTP anchoring is in use.

Available since: 4.6.

Parameters

Preference

Can have one of the values “active”, “passive”.

5.1.3 Monitoring and Logging

Increment custom counter

Increment a custom counter. Custom counters are available via Prometheus and the legacy SNMP interface.

Parameters

counter name

increment

Log received traffic

Log SIP/RTP traffic concealed with logging into PCAP file.

The general log level is used if none is set for that call.

Parameters

log type

Allows choosing what to log:

- `SIP only` logs SIP messages only
- `SIP and RTP+DTMF` logs SIP messages and the whole RTP traffic.
- `SIP and DTMF only` logs SIP and only the RTP packets that are identified to be using the one of the telephone-event payload types seen in the SDP. This option is available since 5.5.
- `SIP and RTP only` logs SIP and only the RTP packets that do not have a telephone-event payload type as seen in the SDP. This option is available since 5.5.

Note that “RTP” options log all packets that are received on the RTP socket, i.e. RTCP and DTLS as well.

PCAP file name

Use filename with .pcap extension.

Log Event

Generate custom event

Parameters

event text

Set log level

Set a specific log level for this traffic.

Note: The global log level will be applied until this Action is processed.

Parameters

log level

see Section Reference of Log Level Parameters

Log Message

Use syslog facility.

Parameters

log level

message text

Log Message for Replies

Report on a transaction that completed with a specific response code. Depending on parameters, such a report can lead to blacklisting or promoting a whitelisted IP address.

Typically used to alarm on requests that were declined because of a possible security risk. The action can report via events, syslog or suggest that the request originator is put on blacklist or promoted on a greylist.

Parameters

reply codes

Comma-separated list of reply codes that trigger the reports or asterisk for any response code.

syslog level

use syslog

send an event

Blacklist UAC IP Address

Blacklist UAS IP Address

Greylist UAC IP Address

Greylist UAS IP Address

Log to grey list

Promote a source IP address from greylist to whitelist.

Parameters

label

Token that differentiates internally the promotion reason; choose some short descriptive string.

Disable privacy monitor mode

Override global configuration for privacy monitor mode to disable it for certain calls.

Note that when used in C rules, call-attempt event will still not be generated in case B-leg refuses.

Available since: 5.1.

5.1.4 Traffic Shaping

Limit parallel calls

Put a quota on number of parallel calls for some specific part of traffic identified by a key. The limit applies separately to inbound and outbound traffic in A and C rules respectively and realm or CA to which the action's rule is linked unless "global key" is turned on. Exceeding calls attempts are rejected using 403.

Parameters

max number of calls

key (optional) that identifies a subset traffic

global key

SIP header

soft limit

report abuse

SIP response code and phrase

Limit CAPS

Put a quota on number of call attempts per second for a traffic subset identified by a key. The limit applies separately to inbound and outbound traffic in A and C rules respectively and realm or CA to which the action's rule is linked unless "global key" is turned on. Authentication counts towards the limit as well. Exceeding calls attempts are rejected using 403.

Parameters

limit CAPS

Maximum number of request per unit of time.

time unit

length in seconds

key attribute

is global key

SIP response code

SIP response reason

SIP header

soft limit

report abuse

Limit Bandwidth per Call

Put a quota on RTP traffic in kbps. A rules steer bandwidth for inbound calls, C rules for outbound. Exceeding RTP traffic is dropped.

Parameters

limit (kbps)

key and global key

SIP response code and phrase

soft limit

report abuse

Limit Bandwidth

Don't admit signaling if its codecs in SDP exceed a limit.

Parameters

limit (kbps)

Set call Timer

Terminate a call if it exceeds a limit length.

Parameters

max call length

Maximum call length in seconds.

5.1.5 Media Processing

Enable RTP anchoring

Anchors RTP media to the ABC SBC.

Allows to centralize media forwarding. Anchoring is a prerequisite for other media processing such as recording.

Additionally, ICE connectivity checks and RTP keep-alive can be introduced for anchored calls. If RTP timeout is introduced and no RTP packet appears, the call is terminated.

RTCP report generation can also be configured to happen on certain conditions described in "RTCP Gen.". RTP Gen. "Always" disables RTCP relay and sends the generated RTCP (available since 4.6).

Parameters

Media far end NAT traversal

"If RFC1918 is in SDP or signaling" option for "Media far end NAT traversal" enables remote address learning only when an RFC1918 IP is seen on SDP `c=` lines or is the signaling IP for the remote endpoint in the dialog (available since 5.0).

Lock on addresses learned from RTP

Address locking affects the socket pair

"Address locking affects the socket pair" will lock both RTP and RTCP socket addresses if one of them locks before the other receives any traffic. For the socket that is locked this way, without seeing any traffic, the source port is allowed to be changed with the first packet received on that socket.

Don't send to RFC1918 addresses

Using this option will prevent the ABC SBC sending any RTP/RTCP/Other data to RFC1918 addresses on the leg.

Available since 5.0.

Enable intelligent relay (IR)

Source IP Header field for IR

Offer ICE-lite

Offer RTCP feedback

Keepalive (sec)

Timeout (sec)

Ignore ICE Offer

RTCP Generation

RTCP Interval

Change SSRC

If used ABC SBC will change the SSRC in RTP and RTCP packets with a locally generated one. Note that Convert DTMF to AVT RTP action will force-enable this behavior even if it is disabled here. For RTCP packets and SSRC replacement, only SSRC that is advertised in the SDP will get be replaced.

Restrict media IP to signaling IP (on leg)

Restricts the incoming and outgoing media packets to a network which is derived by applying a mask on the signaling IP address.

Packets coming from/going to a non-conforming addresses will be dropped.

Applies to RTP, RTCP and other packets.

> ⚠ **Warning**
>
> This action requires RTP anchoring to be enabled as well.

Available since: 5.0.

Parameters

IPv4 Mask

"IPv4 Mask" expects a CIDR value.

"-1" means everything is allowed for IPv4 RTP.

"0" means IPv4 RTP packets will only be accepted if signaling is also IPv4 (and not v6).

"32" means packets should come from and go to the same address seen in signaling.

IPv6 Mask

"IPv6 Mask" is the IPv6 counterpart of the "IPv4 Mask" parameter.

Allow SDP IP

This option will additionally allow communication with the IP specified in respective `c=` line of the SDP.

Available since 5.2.

Force RTP/SRTP

Enforces conversion to the requested protocol in C-rules.

In A-rules it only admits specified protocol and declines requests otherwise. Requires RTP anchoring to be enabled.

Parameters

Key exchange mechanism (DTLS/SDES)

SRTP Fallback to RTP (on leg)

On the leg using this action, if a request is sent with SRTP and the remote endpoint responds with 488, the request is retried with RTP. This works for both initial INVITE and re-INVITEs / UPDATEs.

If "temporary" is false, once the leg switches to RTP, further SDP offers to it will use RTP. If it is true, then further O/A exchange will still try SRTP if it normally would (i.e. through force-srtp action or the other leg sending SRTP).

Note that if the action is on A-rules and SRTP is converted to RTP with Force RTP action on C-rules, then once a RTP-fallback occurs on A-leg, SRTP will not be retried on re-INVITEs going to a-leg even when "temporary" is set.

> ⚠ **Warning**
>
> This action is only meaningful when the RTP anchoring is in use.

This action will override forcing SRTP via Force RTP/SRTP action.

Available since: 5.1.

Parameters

Temporary

Activate audio recording

Record audio into stereo WAV file or using a SIPREC recording server.

Recording type-specific parameters will be available based on the value of the "destination" parameter.

When WAV file recording is used, the call will be recorded as a stereo WAV file where left & right channels contain audio from A & B legs. "call-end" events will contain a link to the file holding the recording. The link will be indexed by the "audio_file" field.

When "destination" starts with `sip:`, SIPREC recording mode will be used. SIPREC-specific parameters will be available to configure options specific to the SIPREC recording mode.

Parameters

destination

Either WAV file name or SIP URI pointing to SIPREC recording server.

WAV-specific:

Discard non-established

Will discard the recording if the call ends before it is established.

SIPREC-specific:

Start announcement

ABC SBC will play an audio announcement before recording starts.

Beep tone and Beep tone interval

If set, ABC SBC will play a tone at the specified interval during the recording.

Stop announcement

ABC SBC will play an announcement before the recording stops.

Caller URI, Caller display name, Callee URI, Callee display name

These parameters are used to fill the participant fields in SIPREC metadata XML (**RFC 7865**) sent in the INVITE message to the SIPREC server.

Do not start yet

Changes the behavior to not start the recording immediately. When this option is enabled, the recording can be started when SIPREC server sends an in-dialog INFO requests with `x-ASC-Recording` header set to `started` and stopped by sending the same header with a value of `stopped`.

Stop call on SIPREC error

Stop the call when SIPREC session can not continue for any reason.

Available since 5.4.

Additional header fields

This parameter can be used to add extra headers to the messages sent to the SIPREC server.

SIP Body

This parameter can take two values. The default one is `Standard` which adds an application/rs-metadata XML in the request body. In this mode further SIPREC Extension fields can be provided. The second mode is `Custom`, which allows configuring up to 3 custom body parts with custom mime-types, headers and contents via template files.

Available since 5.4.

SIPREC Extension Data Enhancements

Adds the `<extensiondata>` section to the SIPREC metadata XML. Fields in the extension data section can be set using the respective parameters. Available in `Standard` SIP Body type.

SIPREC Extension Data | RURI

will set `<apkt:request-uri>`. Available in `Standard` SIP Body type.

SIPREC Extension Data | Realm

will set `<apkt:realm>` and `<apkt:in-realm>`. Available in `Standard` SIP Body type.

SIPREC Extension Data | Additional header fields

will be added as `<apkt:header>`. Available in `Standard` SIP Body type.

Extra Body Part | Mime Type

will add a new body part with the given mime type. Available in `Custom` body type.

Available since 5.4.

Extra Body Part | Headers

will set the new headers to the respective body part. Available in `Custom` body type.

Available since 5.4.

Extra Body Part | Body Template

will set the content of the new body part. The template engine syntax is described below. Available in `Custom` body type.

Available since 5.4.

> **Info**
>
> All header inputs can take multiple headers by separating them with `\r\n`.

Template Engine Syntax

Comments

`comments {# won't #} render` will result in `comments render`.

Loops

`loop will replaced with {% for i in range(4) %}{{ loop.index1 }}{{ i }} {% endfor %}` will result in `loop will replaced with 10 21 32 43`.

Loops can also be written using:

```
alternative
## for i in range(4)
  {{ i }}
## endfor
```

This will result in `alternative\n1\n2\n3\n4\n` where `\n` are line-feeds. In this syntax, the `##` must be at the start of the line.

Conditions

`{% for i in range(4) %}{% if loop.index1 %% 2 %}odd{% else %}even{% endif %}{% endfor %}` will result in `oddevenoddeven`.

Sorting

`sorted list is {{ sort([3,2,1]) }}` will result in `sorted list is [1,2,3]`.

List join

`hello {{ join([1,2,3], " + ") }}` will result in `hello 1 + 2 + 3`.

String manipulation

`hello {{ upper("there") }}` will result in `hello THERE`.

`hello {{ lower("THERE") }}"` will result in *hello there*.

Escaping

`{{ "{% hello %}" }}` will result in `{% hello %}`.

The SBC adds the following function extensions to the template engine:

`{{ abc_replace("<replacement-expression>") }}`

Wraps the SBC's replacement-expressions. I.e. `abc_replace("$ci")` will be replaced with the Call-ID.

`{{ abc_strftime("<format>") }}`

Converts the current system time to (UTC) to a string according to the given format. Format syntax is the same as C language's `strftime`. The final string must not exceed 127 bytes.

`{{ abc_generate_uuid() }}`

Replaced with a base64-encoded UUID.

The SBC adds the following variable extensions to the template engine:

{{ a_leg_stream_label }}

Replaced with the value that the SBC will put in the SDP (`a=label:<label>`) for A leg's stream.

{{ b_leg_stream_label }}

Replaced with the value that the SBC will put in the SDP (`a=label:<label>`) for B leg's stream.

Activate transcoding

Activate transcoding for list of codecs. Listed codecs are added to SDP and transcoded if selected.

When "strict SDP answer" is enabled, while sending SDP answer, SBC will only add the transcoding codecs that were in the offer. Otherwise, all the codecs in the codec list are added to the answer so that we may avoid transcoding if the UA is able to send them.

Parameters

comma-separated codec list

strict SDP answer

Process RTP Header Extension

Enables relaying of RTP header extension in media processor (i.e. transcoded media).

Only supports ED137A.

Available since 5.4.

Convert DTMF to AVT RTP

Convert detected DTMF to RTP/AVT packets (**RFC 4733**/ **RFC 2833**).

Note that this action will make the SBC replace the SSRC and sequence number in relayed RTP/RTCP packets with locally generated ones. For RTCP packets and SSRC replacement, only SSRC that is advertised in the SDP will get be replaced.

This action can be used to convert DTMF received via SIP INFO messages or inband DTMF when used together with Activate Inband DTMF Detection action.

Parameters

Direction

Direction parameter sets on which direction to apply the conversion on.

E.g. setting it to "To B leg" on C rules would apply the conversion on DTMF generated by A leg (caller). Direction defaults to "To B leg" and "To A leg" in A and C rules respectively.

Available since 4.6.

Default volume

This parameter sets the volume when if the SBC can not figure out the volume by other means. Defaults to 20.

Available since 5.0.

Force volume

Forces the volume parameter to always be effective.

Available since 5.0.

Default duration

Sets the duration of the generated DTMF if the SBC cannot figure it out in any other means.

Available since 5.0.

Force duration

Forces the duration parameter to always be effective.

Available since 5.0.

Convert DTMF to SIP INFO

Same as Convert DTMF to AVT RTP except the end result is DTMF in SIP INFO messages.

When used in A rules, DTMF coming from A leg is sent as SIP INFO to B leg. When used in C rules, DTMF coming from B leg is sent as SIP INFO to A leg.

Parameters

Relay AVT RTP

This parameter can be used to control whether to drop RTP AVT packets or to also relay them.

Join meet-me conference

Make a call join a conference.

Note: it is strongly advised to set the configuration synchronization mode to 'pull' for nodes where 'System-generate rooms/PINs' options is enabled. A large amount of notifications about 'outdated provisioned tables' are to be expected otherwise.

Parameters

Enter room via keypad

Room

System-generated rooms/PINs

Room PINs provisioned table

Provisioned Table API user

Provisioned Table API password

Minimal room length

Unacceptable rooms

Room prefix

Split Room number and participant ID

Position to split room

Room is PIN protected

PIN

Use room's PIN as admin PIN

Record participant name

Participant recording filename

Play the number of participants in the room

Play announcements to all participants of the room

Multi-Language support (MLS)

MLS prompt directories

Meet-me conference set PIN

Set and persist the security PIN of a meet-me conference room into a typed provisioned table.

See Default Audio Files for more information about the defaults prompt files.

Available since: 4.6.

Parameters

Room

PIN

Source IP

Path to WAV directory

Provisioned Table API user

Provisioned Table API user password

PINs Provisioned Table

Refuse call with audio prompt

Play an audio announcement and decline an incoming call.

Parameters

file

The filename, relative to the global config option “Prompts/Base Directory”.

As Early Media

Loop

SIP Reply and HF

Play prompt on final response

Play an audio announcement on receipt of a negative final response from downstream.

Parameters

SIP response codes to trigger the announcement

As Early Media

New response code if “as early media”

Optional header fields

announcement WAV filename OR characteristics of a generated ringtone

Generate Ring-Back Tone

Play an audio file or a dual-frequency tone instead of default ringing tone.

Parameters

On downstream 180

Start playing when a 180 response arrives.

On Timer

Start playing if a number of seconds elapses. Turned off if zero.

Generate Ringtone

If turned on, a dual-tone with specified frequencies and durations will be played; otherwise a specified audio file will be used.

File

Audio file to be played.

Loop

When audio file is chosen this option chooses whether to play it once or in a loop.

Activate Music On Hold

Use this action on a call to play an audio file when a call participant puts the call on hold. It is possible to specify how to signal the on-hold status in SDP.

Parameters

music file name

playback in loop

Hold indication

The method of hold signalling. Either preserve incoming or via SDP attribute (`sendonly`, `sendrecv`, `inactive`) or using connection IP set to 0.0.0.0 (**RFC 2543**).

Activate Inband DTMF Detection

Use this action together with the "Convert DTMF to RTP/AVT" or similar actions to detect and convert inband DTMF.

Note that this action:

- Does not filter the inband DTMF,
- will increase the CPU usage on the RTP traffic processing.

Available since: 4.6

Parameters

Direction

Can be used to set in which direction the detection will be enabled.

Mode

Can be used to i.e. not enable the detection if telephone-event is in the SDP.

DTMF Termination Same SSRC (on leg)

Actions that result in DTMF termination/generation (i.e. transcoding, Convert DTMF to AVT RTP) would generate the DTMF RTP (RFC4733/RFC2833) using a new SSRC. Using this action changes it to injecting DTMF RTP into ongoing RTP stream.

Note that this has the drawback of not being able to generate DTMF RTP if no other RTP packets are being relayed. This is because we can not reliably estimate RTP timestamp unless we see the live RTP traffic.

This action only acts on the packets going towards the leg it is set on. If it is in A-rules, then DTMF packets going towards the A-leg will have the same SSRC. If it is in C-rules, then DTMF packets going towards the B-leg will have the same SSRC.

Available since: 5.0

DTMF Termination Stable Duration Increments (on leg)

Actions that result in DTMF termination/generation (i.e. transcoding, Convert DTMF to AVT RTP) would generate the DTMF RTP (**RFC 4733**/ **RFC 2833**) using variable increments in 'duration', according to the wallclock during the relay of the other RTP packets. Using this action changes it to increment the duration in fixed steps. The step interval is determined using ptime attribute of the SDP, calculated from timestamp increments of the RTP packets or default to 20ms, in that order.

Available since: 5.2

Sticky Stream SSRC (on leg)

When used, the RTPs sent to the respective call agent (A-rules: Caller, C-rules: Callee) use the same SSRC value per stream. The SSRC is generated randomly for each call and is derived to be unique for each stream by using the SDP media index of the stream.

This action is only meaningful when the RTP anchoring is in use. Similar to the `Convert DTMF to AVT RTP` action, using this action will automatically change SSRC regardless of the "Change SSRC" option in RTP anchoring action, forcibly enabling it for both legs.

Available since: 5.4

5.1.6 SIP Dropping

Reply to request with reason and code

Send a response to a SIP request.

Parameters

Code

Reason

Reason phrase

Header fields

Additional header fields (optional).

Blacklist by firewall if repeated

Drop request

Drop request without replying.

Parameters

Event throttling key

Allow unsolicited NOTIFYs

Allow forwarding NOTIFY requests without a prior subscription (either implicit with REFER, or explicit with SUBSCRIBE).

5.1.7 Scripting

Set Call Variable

Stores a computing result in an variable. The variable can be tested using the Call Variable condition and/or referred to from actions using the `$V(gui.varname)` replacement.

Parameters

variable name

variable value

5.1.8 Register Processing

Enable REGISTER caching

Stores a cached copy of REGISTER contacts before forwarding.

Retarget R-URI from cache

Rewrites AoR in request URI with contacts cached using Enable REGISTER caching.

Parameters

enable NAT handling

enable sticky transport

REGISTER throttling

Force UAs to refresh registrations within a time window. Particularly useful to trigger REGISTER-based keep-alives to facilitate NAT traversal.

Parameters

minimum registrar expiration

maximum UA expiration

Save REGISTER contact in registrar

Act as local registrar and store registers locally.

Restore contract from registrar

Restore contact from registrar.

5.1.9 External Interaction

ENUM query

Make an ENUM dip. The queried value may contain replacement expression, suffix is appended to the query.

Parameters

queried value

domain suffix

ENUM services

Read call variables over REST

Do REST query to given URL and set call variables received in reply.

Since ABC SBC 5.3 if content-type is application/json then a json content is parsed.

Please note, neither arrays nor nested objects are supported. Only simple objects similar to the one in example are supported:

```
{
  "attribute_name": "value",
  "foo": "bar"
}
```

Parameters

REST URI

Read call variables from table

Read variables from a provisioned table

Parameters

table name

query key

5.1.10 NAT Handling

Enable dialog NAT handling

Remember during dialog lifetime where the initial dialog-initiating request came from and sends all subsequent SIP traffic there.

5.1.11 Other

Support serial forking proxy

Permit to reset early media upon 181-indicated serial forking.

Fork

Fork a new parallel branch to a URI.

Action is supported in inbound rules only.

Parameters

SIP URI

5.2 Reference of Global Configuration Parameters

This reference lists all global configuration parameters used in SBC. Note that they have default values which are designated to accommodate most use-cases and can have massive impact on operation if

changed: modify them only after careful consideration. The GUI screen is showing recommended default values. When the actual value is changed, the default value is highlighted as bold text.

> **Info**
>
> When the global configuration parameters are updated, a warning message with a link to activate the new SBC configuration is shown in the GUI. No changes are applied until the "activate" link is used. When the configuration changes are applied, appropriate services might be restarted (e.g. SIP and RTP processes) depending on what parameters were changed. Note that this may cause service disruption.

The configuration parameters are grouped as follows:

- AWS Parameters
- Backup Parameters
- CDR Parameters
- Conference Parameters
- Eventbeat Parameters
- Event Parameters
- Firewall Parameters
- Lawful Interception Parameters
- Low-level Parameters
- Miscellaneous Parameters
- PCAP Parameters
- RTP handling Parameters
- SEMS Parameters
- SIPREC Parameters
- SIP Parameters
- SRTP Parameters
- Signaling SSL
- Syslog Parameters
- System Monitoring Parameters
- Probe Parameters

5.2.1 AWS Parameters

These parameters are used when SBC is deployed on Amazon AWS.

They are currently applied during the initial AWS configuration when using High Availability (HA) under AWS.

> **Info**
>
> Anyone in possession of an AWS IAM User Access Key can impersonate the key's owner. It is therefore recommended to create a dedicated IAM user with limited permissions and configure the SBC to access AWS using this user's identity.

For more details, see the AWS documentation on IAM user identities: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html

Parameter Name	Description
Region for AWS requests	AWS region used for requests. Default value: empty API: aws_region
AWS access KEY ID	Access Key ID of an AWS user with permissions for the AWS service. Default value: empty API: aws_key_id
AWS secret access KEY	Secret key associated with the AWS user's Access Key ID. The secret is only visible when the key is created. If forgotten, the key must be regenerated. If compromised, anyone in possession of the key can impersonate the user. Default value: empty API: aws_key

Table 23: AWS Parameters

5.2.2 Backup Parameters

These parameters configure SBC daily backups. For more details, see Backup and Restore Operations.

Parameter Name	Description
Equivalent settings as for CCM	If enabled, the settings on this Backup tab will not be applied to SBC nodes. Instead, the same settings as configured for the CCM node (under *CCM -> CCM Config -> Backup* page) will be applied to SBC nodes. Default value: disabled API: backup_ccm_equivalent
Create daily Sbc configuration backups	If enabled, a daily snapshot of the SBC configuration will be created and stored as a gzipped tarball file. Default value: disabled API: backup_enable
Include provisioned tables in daily backups	If enabled, the daily backup will also include the content of all provisioned tables. Default value: enabled API: backup_prov
Number of days to keep backups	Sets the retention period for backup files. All files named 'sbc-backup-*' in the backup directory that are older than the specified number of days will be deleted during each daily backup run. Use '0' to disable automatic deletion of old backup files. Default value: 7 API: backup_keep
Destination directory for backups	Specifies the destination directory for the daily backup files. Default value: /data/backups API: backup_dir

Parameter Name	Description
Full path to extra files or dirs to include in backup	Extra custom files or directories to be included in the backup, using full paths. Multiple entries can be separated by commas. A '*' wildcard may be used. Paths must not contain commas. Default value: empty API: backup_extra_files

Table 24: Backup Parameters

5.2.3 CDR Parameters

These parameters define how and where CDRs are stored. For more details, see Call Data Records (CDRs).

Parameter Name	Description
Enable CDRs	Enables writing of Call Detail Records (CDRs). **Changing this value may restart signaling process.** Default value: enabled API: cdr_enabled
Number of daily CDR files to keep	CDR retention policy. The SBC produces CDRs for all completed calls in CSV format. This setting specifies the number of CDR files to retain. **Changing this value may restart signaling process.** Default value: 93 API: cdr_days_keep
Directory for exported CDR files	Directory in the filesystem where the CSV CDR files are stored. Default value: /data/cdr/export API: cdr_export_dir
CDR files rotation frequency	Sets the frequency of CDR file rotation. Supported values: "daily", "weekly", or "monthly". The number of rotated files to keep before deletion is defined by the *Number of CDR files to keep* parameter. Default value: daily API: cdr_rotate_freq Possible values over API: • hourly • daily • weekly • monthly • yearly
Enable new version of CDRs (CDR-NG)	Enables the new version of CDRs, called CDR-NG. This feature is in an experimental state. **Changing this value may restart signaling process.** Default value: disabled API: cdr_ng_enabled

Parameter Name	Description
CDR format	Defines the structure of CDR records. See Customized CDR Records. **Changing this value may restart signaling process.** Available since: 5.6 Default value: $srcrlm.name, $srcca.name, $dstrlm.name, $dstca.name, caller_id_user, caller_id_host, caller_id_name, callee_id_user, callee_id_host, callee_id_name, $ltag, $start_tm, $connect_tm, $end_tm, $duration, $setup_duration, $bill_duration, sip_req_uri, sip_from_uri, sip_to_uri API: cdr_legacy_fmt

Table 25: CDR Parameters

5.2.4 Conference Parameters

This section contains the configuration parameters for meet-me Conference feature.

Parameter Name	Description
Keep participant's name recordings for	Duration (in hours) for which files containing web conference participants' names are retained on the file system. These files are not subject to replication. **Changing this value may restart signaling process.** Default value: 90 API: wc_preserve_participants_files_for
Play the number of participants in the room	If enabled, the number of participants is announced when a participant joins or leaves the conference room. Alternatively, a participant can press the star (*) key during a call to hear this information. **Changing this value may restart signaling process.** Default value: disabled API: wc_echo_nb_participants
Use room security pin value for the admin pin	If both *Use security pin* and *Use admin pin* options are enabled for a room, the admin pin value is set to the same value as the security pin. **Changing this value may restart signaling process.** Default value: disabled API: wc_use_secu_as_admin_pin
Path to directory holding digits wav files	Path to the directory containing audio files used to announce numbers (e.g., "one", "twenty", "(seven-)teen"). By default, SBC ships two versions of this directory: `/usr/lib/sems/audio/webconference/digits/` for English prompts, and `/usr/lib/sems/audio/webconference/de/digits/` for German prompts. **Changing this value may restart signaling process.** Default value: /usr/lib/sems/audio/webconference/digits/ API: wc_path_to_digit_wav

Parameter Name	Description
List of provisioned tables to watch for expired room	List of provisioned tables to watch for expired rooms. Generated web conference names and PINs are stored in the CCM provisioned tables. The CCM attempts to remove expired rooms from these tables every day at 02:00. Default value: empty API: wc_generated_table_to_watch
Generated rooms validity	Number of days a generated conference room remains active. Once closed, the room's PIN is blocked for a fixed period. **Changing this value may restart signaling process.** Default value: 30 API: wc_block_generated_room_after
Keep expired generated rooms	Number of days before a closed generated conference room's PIN is unblocked. **Changing this value may restart signaling process.** Default value: 90 API: wc_block_generated_room_for

Table 26: Conference Parameters

5.2.5 Eventbeat Parameters

These parameters are used to tweak and debug event communications between an SBC node and an ABC Monitor node. They also enable generating statistics, which can be exposed through the application interface on TCP ports :4247 and :4248.

Parameter Name	Description
Event batching size	Maximum number of events sent at once to the monitor. Default value: 10 API: eventbeat_batch
Enable eventbeat statistic reporting	Exposes live event processing metrics on a TCP port (`:4247` for the primary ABC Monitor and `:4248` for the secondary ABC Monitor). Default value: enabled API: eventbeat_enable_stats
Interval between each statistic	Interval in seconds at which a single statistics entry is recorded. Default value: 5 API: eventbeat_stat_interval
How many statistic entries per payload	Number of statistics entries to be returned in a single payload. Example: To collect statistics for the last minute in 5-second intervals: • Set *Interval between each statistic* to `5` • Set *How many entries per payload* to `12` Default value: 10 API: eventbeat_stat_report

Table 27: Eventbeat Parameters

5.2.6 Event Parameters

These parameters are used to define how and where events are stored.

Parameter Name	Description
Number of days to keep old traffic log files	Defines the local retention policy for traffic log files. Particularly useful when no ABC Monitor is attached to the SBC. Must be shorter than the retention policy on ABC Monitor to avoid copying files already expired on ABC Monitor. Default value: 7 API: events_days_keep
ABC Monitor address	Specifies the IP address or DNS name of ABC Monitor. Empty if no ABC Monitor is attached to the SBC. **Changing this value may restart signaling process.** Default value: empty API: events_remote_console
Secondary ABC Monitor address	Specifies the IP address or DNS name of the secondary ABC Monitor. Empty if no secondary ABC Monitor is attached to the SBC. **Changing this value may restart signaling process.** Default value: empty API: events_remote_console2
Replicate traffic logs to ABC Monitor	If enabled, collected PCAP files are sent to a Monitor server using the rsync protocol. Files are deleted from the SBC after transfer. Default value: enabled API: replicate_traffic_log
Replicate traffic logs to secondary ABC Monitor	If enabled, collected PCAP files are sent to the secondary Monitor server using the rsync protocol. Files are deleted from the SBC after transfer. Default value: enabled API: replicate_traffic_log2
Traffic Log Chunk Seconds	The duration at which the PCAP is split into chunks. Default value: 10 API: traffic_log_chunk_secs
Replicate recordings to ABC Monitor	If enabled, recorded audio files (see Section Audio Recording) are sent to a Monitor server using the rsync protocol. Files are deleted from the SBC after transfer. Default value: enabled API: replicate_recordings
Replicate recordings to secondary ABC Monitor	If enabled, recorded audio files (see Section Audio Recording) are sent to the secondary Monitor server using the rsync protocol. Files are deleted from the SBC after transfer. Default value: enabled API: replicate_recordings2

Parameter Name	Description
Replication rsync password	Specifies the rsync password used for replicating traffic logs and recorded audio files. Available up to: 5.5 Default value: empty API: replication_rsync_password
Replication rsync password for secondary ABC Monitor	Specifies the rsync password used for replicating traffic logs and recorded audio files to the secondary Monitor. Available up to: 5.5 Default value: empty API: replication_rsync_password2
Use TLS secure connection to ABC Monitor	If enabled, events, traffic logs, and recording files are sent to ABC Monitor over a TLS-secured connection. It is recommended to install a trusted certificate on ABC Monitor instead of the default self-signed certificate. On the SBC side, the TLS profile of the IMI interface is used. Default value: disabled API: events_use_tls
Number of hours to keep old recordings (0 to not delete)	Defines the retention policy for recorded WAV files. A value of 0 disables deletion. Default value: 168 API: recordings_hours_keep
Generate an event if a SIP transaction reaches the defined number of retransmissions	Generates a “notice” event when the configured number of SIP retransmissions is reached. These events appear in the ABC Monitor Transport Dashboard. Use with care — low values may cause excessive events. If used, recommended value is 4. Use ‘0’ to disable this functionality. **Changing this value may restart signaling process.** Default value: 0 API: sip_retrans_event_trigger
Maximum number of events buffered in local Redis	Defines the retention policy for locally buffered events in Redis. Use ‘0’ for no limit. **Changing this value may restart signaling process.** Default value: 0 API: redis_max_local_events
List of call variables added into events	Specifies a list of call variables to include in call events. Format: `<var_name>:<flag>`, where `<flag>` is 0 or 1, indicating whether the value can be overwritten. Use * as a wildcard to include all. **Changing this value may restart signaling process.** Default value: empty API: cvars_allowed

Parameter Name	Description
Generate an event on UDP receive buffer errors	If enabled, an alert is generated when UDP receive buffer errors are detected on the system network interface. Default value: enabled API: checknet_udp_recv_err
Generate an event on UDP send buffer errors	If enabled, an alert is generated when UDP send buffer errors are detected on the system network interface. Default value: enabled API: checknet_udp_send_err
Generate an event on UDP packet receive errors	If enabled, an alert is generated when UDP packet receive errors are detected on the system network interface. Default value: enabled API: checknet_udp_pacerr_err
Generate an event on IP incoming packet receive errors	If enabled, an alert is generated when IP incoming packet receive errors are detected on the system network interface. Default value: enabled API: checknet_ip_incdis_err
Generate an event on outgoing packets dropped errors	If enabled, an alert is generated when outgoing packets dropped errors are detected on the system network interface. Default value: enabled API: checknet_ip_outdrop_err
Alarm when number of calls reaches % of the license	Issues a warning when the number of sessions reaches X% of the license limit. An info-level event is generated when the session count drops back below X%. **Changing this value may restart signaling process.** Default value: 75 API: session_nearly_limit_percent
Privacy monitor mode	If enabled, prevents sending call-attempt, call-start, and call-end events to ABC Monitor. Can be overridden using the “Disable privacy monitor mode” action. **Changing this value may restart signaling process.** Default value: disabled API: privacy_monitor
Destination monitor event interval (sec)	Sets the interval, in seconds, at which destination monitor events are generated. **Changing this value may restart signaling process.** Default value: 300 API: interval_dest_monit

Parameter Name	Description
Threshold of number of events buffered on Sbc to set warning	Sets the maximum number of events allowed in the Redis queue on the SBC before the node status is set to “warning” on the System Monitoring page. Default value: 500 API: events_buffer_treshold
Enable events redis disk persistence	If enabled, enables Redis disk persistence for events, using the `/data/redis` directory. Ensure sufficient disk space before enabling. Default value: disabled API: events_redis_persistence
Periodic RTP Statistics	Enables sending periodic RTP statistics per call leg at 10-second intervals. **Changing this value may restart signaling process.** Available since: 5.4 Default value: disabled API: periodic_rtp_stats

Table 28: Event Parameters

5.2.7 Firewall Parameters

This section contains the configuration parameters for Firewall.

Parameter Name	Description
Enable SBC firewall	If enabled, the firewall chains are populated with SBC firewall rules. If deployed in a container or on a system not supporting nftables or nfsets, disable this option to avoid SBC node errors in System status. Note: on SBC < 5.4 the firewall uses iptables and ipsets. **Changing this value may restart signaling process.** Available since: 5.0 Default value: enabled API: fw_enable
Reject packets instead of dropping	If enabled, packets not allowed by the firewall are rejected and an ICMP admin-prohibited message is sent. If disabled, packets are silently dropped. Note: on SBC < 5.4 the firewall uses reject; starting with 5.4 it uses drop by default. Available since: 5.4 Default value: disabled API: fw_reject

Parameter Name	Description
Do not accept any ICMP packets	If enabled, incoming ICMP packets are dropped and not answered. Use with caution, as blocking ICMP may cause connectivity issues. Available since: 5.5 Default value: disabled API: fw_block_icmp
Blacklist IP addr for repeated signaling failures	If enabled, IP address of request that failed authentication, exceeded limit, failed sanity check, was dropped by "Drop" action or "Log message / Event for replies" action was used, will be added to the blacklist, silently dropping all packets from it. Note: the specific blacklist reasons must also be enabled in CA settings or in the "Drop" or "Log message/Event for replies" actions. See Section Automatic IP Address Blocking for details. Default value: disabled API: fbl_sig_enable
Signaling failures blacklist: IP address start score before any offense	Sets the starting score used before any offenses are registered. This value decreases with each offense until it reaches 0 or less, at which point the IP address is blacklisted. See Section Automatic IP Address Blocking for more details. **Changing this value may restart signaling process.** Default value: 2.8 API: fbl_start_score
Signaling failures blacklist: rate per second used to calculate a time-related bonus between offenses	Sets the allowed rate of offenses in events per second. Allows the score to recover slightly over time, acting as a bonus for good behavior. See Section Automatic IP Address Blocking for more details. **Changing this value may restart signaling process.** Default value: 0.0005 API: fbl_allowed_rate
Signaling failures blacklist: time in seconds to remove entries for which no event has occurred from score calculation	Sets the number of seconds after which, if no offense is seen from an IP address, it is removed from the scoring table. New offenses from removed IPs use the start score. See Section Automatic IP Address Blocking for more details. **Changing this value may restart signaling process.** Default value: 600 API: fbl_inactive_timeout
Time in seconds to blacklist IP addr for signaling failures	Sets how long an IP address remains blacklisted for signaling failures before automatic removal (for drop, failed auth, limit, sanity). See Section Automatic IP Address Blocking for more details. Default value: 3600 API: fbl_sig_bantime

Parameter Name	Description
Greylist: time delay in seconds to give IP a chance to prove validity	If the traffic from IP address proves validity during this probation period, the source IP addr will be added to whitelist. Note that the corresponding action options like "Greylist IP address" or "Log to greylist" have to be used. See Section Automatic Proactive Blocking: Greylisting for more details. Default value: 600 (seconds) API: fgl_delaytime
Greylist: time period in seconds when IP can be blacklisted if repeats and did not prove validity	If traffic from IP address did not prove validity during the probation time period, and new packet comes during this time period since first packet, the source IP addr will be added to blacklist. Note that the "Greylist" flag has to be enabled on SBC signaling interface for this to work. All traffic from the IP addresses on blacklist will be silently dropped. See Section Automatic Proactive Blocking: Greylisting for more details. Default value: 1800 (seconds) API: fgl_learntime
Greylist: time in seconds to keep IP on blacklist	Sets how long to keep the IP address on blacklist. After this time it is removed from blacklist and has a chance to prove validity again. See Section Automatic Proactive Blocking: Greylisting for details. Default value: 86400 (seconds) API: fgl_bltime
Greylist: time in seconds to keep IP on whitelist	Sets how long to keep IP address on whitelist. After this time it is removed from whitelist and has to prove validity again. See Section Automatic Proactive Blocking: Greylisting for details. Default value: 86400 (seconds) API: fgl_wltime
Greylist: additional ports or port ranges (a:b) to check in addition to signaling ports, space separated	Sets additional ports to ports defined on SBC signaling interfaces. If used, traffic coming to this port(s) will be also subject to the greylisting procedure. You can specify single port(s) or port ranges (in format lower:higher), space separated. See Section Automatic Proactive Blocking: Greylisting for details. Default value: empty API: fgl_extraports
Log blacklisted IP addresses to syslog	If enabled, logs blacklisted IP addresses to syslog. Entries are stored in `/var/log/frafos/sems-blacklist.log`. **Changing this value may restart signaling process.** Default value: disabled API: fbl_syslog_enable

Parameter Name	Description
Log greylisted IP addresses to syslog	If enabled, logs greylisted IP addresses to syslog. Entries are stored in `/var/log/frafos/sems-greylist.log`. **Changing this value may restart signaling process.** Default value: disabled API: fgl_syslog_enable
Overall limit in packets per second from not approved IP addresses	This option can be used to set overall packets per second limit on all IP addresses, that did not prove validity using "Greylist IP address" or "Log to greylist" action options. Use with caution. Use '0' to disable any rate limiting. Default value: 0 API: ratelimit_notwl
Filter incoming signaling and media traffic by input interface name	Allow incoming signaling and media traffic only from corresponding interfaces. It may be needed to be turned off in some cases like more interfaces connected to the same segment or macvlan interfaces used, depending also on ARP related settings. Available since: 5.5

Table 29: Firewall Parameters

5.2.8 Lawful Interception Parameters

This section contains the configuration parameters for Lawful Interception (LI).

Parameter Name	Description
Lawful Interception enabled	If enabled, activates the lawful interception feature. This must also be configured under the corresponding action to take effect. **Changing this value may restart signaling process.** Default value: disabled API: li_enabled
Operator ID	Specifies the Operator ID value. **Changing this value may restart signaling process.** Default value: empty API: li_operator_id
Delivery Country Code (DCC)	Specifies the Delivery Country Code (DCC) value. **Changing this value may restart signaling process.** Default value: empty API: li_deliv_cc
Network Element Identifier	Specifies the Network Identifier (NID), which consists of the operator ID and, optionally, the network element. **Changing this value may restart signaling process.** Default value: empty API: li_network_element_id

Parameter Name	Description
Interception Point ID	Specifies the Interception Point ID value. **Changing this value may restart signaling process.** Default value: empty API: li_interception_point_id

Table 30: Lawful Interception Parameters

5.2.9 Low-level Parameters

These settings take effect only after a server reboot. Additional information can be found in Hardware Specific Configurations.

> **Warning**
>
> Changing these parameters may significantly alter system behavior. Their impact largely depends on the equipment in use.

Parameter Name	Description
Interfaces where to enable RPS	Network interfaces on which a "receive packet steering" kernel feature should be enabled, separated by spaces. While the kernel leaves this option by default off, turning it on can increase media throughput. Default value: empty API: tunings_rps_interfaces
Interfaces where to set ethtool options	Network interfaces, separated by spaces, where to apply coalesce and ringbuffer *ethtool* options specified in "Coalesce ethtool options" and "Ringbuffer ethtool options" attributes. Default value: empty API: tunings_ethtool_interfaces
Coalesce ethtool options	Ethernet adapter coalescing options (syntax of *ethtool*). Applied to interfaces listed in "Interfaces where to set ethtool options". Tunes the trade-off between less-CPU-intensive and more-real-time packet processing latency; results depend on used network card. Default value: rx-usecs 200 tx-usecs 100 API: tunings_ethtool_coalesce
Ringbuffer ethtool options	Ethernet adapter RX/TX ring parameters (syntax of *ethtool*). Applied to listed interfaces in "Interfaces where to set ethtool options". Fine-tuning this parameter is specific to used network card. Increasing buffer sizes allows to deal with temporary packet bursts, while latency may increase. Default value: rx 1000 tx 1000 API: tunings_ethtool_ringbuffer

Parameter Name	Description
Interfaces where to bind irqs to CPUs	Network interfaces on which the individual interrupts for receive and transmit queues should be statically bound to individual CPUs / CPU cores. This option may increase media throughput on network cards with multiple queues. Default value: empty API: tunings_irqs_interfaces
Run db check on boot	If enabled, runs `mysqlcheck` during boot to help recover from unexpected shutdowns. May slow startup. Default value: enabled API: tunings_run_mysqlcheck
Clean tmp files on boot	If enabled, clean-up system directory for temporary files at boot. Default value: enabled API: tunings_clean_tmp
Sems memory limit in % from total memory	Limits maximum memory usage of the Sems process as a percentage of total system memory. Set to '0' for no limit. Default value: 75 API: sems_mem_limit
Provisioned tables redis disk persistence time interval (in seconds)	Time interval (seconds) after which provisioned tables data on an SBC backup node is saved from Redis to disk for persistence. Saving occurs only if changes exceed the next setting's threshold. Default value: 600 API: redis_prov_savetime
Provisioned tables redis disk persistence number of records to trigger save	Minimum number of provisioned tables records changes that trigger save to disk. The data will be saved when both the number of changed records and the time interval conditions are met. Default value: 1 API: redis_prov_saverecords
Use real-time priority on provisioned tables redis	If enabled, real-time process priority will be used on provisioned tables redis db, which helps performance. It can only be used when operating system or container permissions support this. For podman installations please make sure the `--cap-add=CAP_SYS_NICE` is used if redis real-time priority is required. Default value: disabled API: redis_prov_enable_rt

Parameter Name	Description
Session processor threads	Threads that process the SIP signaling and the B (routing) and C (outbound) rules of the ABC rule set. They are created in a thread pool among which all SIP sessions are distributed. Recommended: CPU threads x 2, minimum 8 threads. If the SBC needs to process a lot of external data in the routing or C rules, e.g. needs to query provisioned tables or external API server via REST, then it is recommended to set this to a high number. **Changing this value may restart signaling process.** Default value: 16 API: sems_session_processor_threads
Media processor threads	Threads for RTP transcoding and media apps (conferencing, announcements). In normal SBC operation, when those functionalities are not used, these threads will be idle. Like the session processor threads, the number configured here sets the number of threads created in a thread pool among which all media sessions are distributed. If transcoding or media applications are used, it is recommended to set this number to CPU threads x 2, otherwise it is recommended to leave them to the default or even less. **Changing this value may restart signaling process.** Default value: 16 API: sems_media_processor_threads
SIP server threads	These threads receive SIP messages from the network and initially parse them for later processing by the Session processor threads, immediately reply e.g. if the reply is given by the SIP dialog state (e.g. errors). They also process the A rules of the ABC rule sets. The number of threads configured here is started for every signaling interface (SI), and one set for udp and one for TCP; so e.g. if five SI interfaces are configured, and this is set to 4, then 5*4*2=40 threads are started. The recommended number depends on the number of signaling interfaces; e.g. on a setup with two signaling interfaces, the recommended number would be equal to the number of CPU cores (e.g. 8, 16 or 32). On a setup with many signaling interfaces, this should be set to e.g. 2 or 4. **Changing this value may restart signaling process.** Default value: 16 API: sems_sip_server_threads
Out-of-dialog requests threads	Threads handling REGISTER, SUBSCRIBE/NOTIFY and MESSAGE requests. It is recommended to increase this number, if processing many registrations or subscriptions. **Changing this value may restart signaling process.** Default value: 1 API: sems_ood_requests_threads

Parameter Name	Description
RTP receiver threads	Threads that receive and relay RTP (including SRTP decrypt when enabled). As with the thread pools above, this number is a global number of threads for a thread pool. The recommended number to set this to is 2-4 times the usable CPU hardware threads. **Changing this value may restart signaling process.** Default value: 16 API: sems_rtp_receiver_threads
Call restore threads (HA)	Thread pool used only for call restoration after HA failover. It is recommended to set it to the number of usable CPU hardware threads. **Changing this value may restart signaling process.** Default value: 4 API: restore_threads
HA interval to send adverts to peer, in seconds	Keepalived advert interval (seconds). Lower values shorten failover detection time but may cause stability issues. It is recommended keeping the default value for typical setup. Default value: 1 API: ha_advert_int
HA priority	VRRP priority (1-255). The `255` has special meaning of address owner. It is recommended to keep at default value for typical deployments. Default value: 250 API: ha_priority

Parameter Name	Description
HA mode	Configures the HA mode to use. This option should be modified only if required by a specific setup. **Warning:** Changing this option may trigger an HA switchover. The default mode uses multicast on the IMI interface (or a custom interface where the HA application is enabled) for HA VRRP advertisements. The second mode uses unicast on the IMI interface (or a custom interface where the HA application is enabled) for HA VRRP advertisements. The third mode uses multicast on both IMI and SIG interfaces for HA VRRP advertisements, providing greater robustness. For this mode to work, the system interface corresponding to the signaling interface must also have a non-VIP IP address. The fourth mode uses multicast only on SIG interfaces for HA VRRP advertisements and employs macvlan sub-interfaces for VIP IP addresses. This mode is typically not needed, but it can improve switchover performance in some cases. For this mode to work, the system interface corresponding to the signaling interface must also have a non-VIP IP address. In some cases, it may be needed to turn off the global config option “Filter incoming signaling and media traffic by input interface name” uder Firewall tab. Available since: 5.5 Default value: use multicast, vrrp on IMI interface only API: ha_mode Possible values over API: • “default”: “use multicast, vrrp on IMI interface only” • “unicast”: “use unicast, vrrp on IMI interface only” • “imisig”: “use multicast, IMI and SIG interfaces for vrrp” • “vmac”: “use multicast, use only SIG interfaces for vrrp and use macvlan for VIPs”

Table 31: Low-level Parameters

5.2.10 Miscellaneous Parameters

This section contains miscellaneous configuration parameters, including GeoIP and other general-purpose settings.

Parameter Name	Description
Blacklist timeout for IP addresses from external sources	Timeout (seconds) for IP addresses blacklisted by RESTful requests. Default value: 86400 API: externalbl_timeout
Enable sending important syslog entries to ABC monitor	If enabled, sends syslog entries with severity `critical` through `emergency` as alerts to the ABC Monitor. Default value: disabled API: enable_syslog_event

Parameter Name	Description
Session Management enable	If enabled, advanced load balancing is used. **Changing this value may restart signaling process.** Default value: disabled API: session_management_enable
GeoIP - account ID for geoipupdate command	Account ID used with the `geoipupdate` command, which periodically updates the GeoLite2 database if a license is provided. The license must be created with a MaxMind account by user. Default value: empty API: geoip_account_id
GeoIP - license key for geoipupdate command	License key used with the `geoipupdate` command, which periodically updates the GeoLite2 database if a license is provided. The license must be created with a MaxMind account by user. Default value: API: geoip_license_key

Table 32: Miscellaneous Parameters

5.2.11 PCAP Parameters

These parameters define how the most recent SIP traffic is recorded on the system for troubleshooting purposes. The SBC stores SIP traffic in PCAP files of a configured size and automatically deletes the oldest files when the retention limit is reached. The recorded PCAP files can be accessed and inspected in the administrative interface, as described in Section User Recent Traffic.

Parameter Name	Description
File size in MB for one pcap file	Maximum size of a PCAP file before a new file is created. Default value: 50 API: pcap_file_size
Number of pcap files to keep	Defines the PCAP retention policy. Files are rotated, and only the configured number of files is kept. Least recent files are deleted. Use 0 to disable SIP traffic capture (not recommended). Note: PCAP filenames use ".pcapXX", where XX is the file number. Changing the number of files deletes all existing traffic.pcap* files once configuration changes are activated. Default value: 0 API: pcap_file_count

Table 33: PCAP Parameters

5.2.12 RTP handling Parameters

This section contains configuration parameters that control how RTP streams are handled and processed.

Parameter Name	Description
Drop unnegotiated RTP payloads	If enabled, RTP packets with payload types not negotiated in the SDP are dropped. **Changing this value may restart signaling process.** Default value: disabled API: drop_unnegotiated_payloads
Do not drop unnegotiated comfort noise RTP payload	If enabled, RTP packets with static comfort noise payload type (13) are not dropped even if not negotiated in SDP and *Drop unnegotiated RTP payloads* is enabled. **Changing this value may restart signaling process.** Default value: disabled Available since: 5.3 API: dont_drop_unnegotiated_cn
Force symmetric RTP for media server apps	If enabled, embedded media processing actions ignore IP addresses in the caller's SDP and send RTP to the source of received RTP packets. **Changing this value may restart signaling process.** Default value: enabled API: sems_force_symmetric_rtp
RTP keep-alive frequency	Defines how often (in seconds) the SBC sends RTP keep-alive packets to its peers. See Setting RTP Inactivity Timer and Keepalive Timer. **Changing this value may restart signaling process.** Default value: 0 API: sems_rtp_keepalive_freq
RTP keep-alive method	Defines which RTP keep-alive method is used. **Changing this value may restart signaling process.** Default value: "RTP version equals 0" API: sems_rtp_keepalive_method Possible values over API: • "0": RTP version equals 0 • "1": All RTP header fields equal 0
RTP timeout	Defines the time (in seconds) after which a call is terminated if RTP packets stop arriving. See Setting RTP Inactivity Timer and Keepalive Timer. **Changing this value may restart signaling process.** Default value: 0 API: sems_rtp_timeout

Parameter Name	Description
Learn remote media address interval	Interval (ms) after the first RTP packet is received during which the remote address may still change and be re-learned. After this interval SEMS locks the remote address. Especially for re-learning after re-Invite, this may prevent locking on the old address due to some late RTP packets from the old remote address. Use 0 to disable (lock on the first packet). **Changing this value may restart signaling process.** Default value: 1000 API: learn_raddress_interval
Recording playout buffer type	Defines the playout buffer type used when recording WAV files. Possible values: • adaptive: Recommended, especially for higher jitter and packet loss showing in the RTP stream. • simple: Basic buffer, may fail on lossy links. **Changing this value may restart signaling process.** Default value: adaptive API: recording_playout_buffer Possible values over API: • adaptive • simple
MOS Average S Coefficient	Coefficient for short-term MOS calculation. The value is updated with the formula `value=value*coefficient+new_value*(1-coefficient)` on each call end where new_value is the MOS average on call end. Initial value for `value` is `4.409405954387269`. **Changing this value may restart signaling process.** Default value: empty API: rtp_mos_average_s
MOS Average M Coefficient	Coefficient for medium-term MOS calculation. The value is updated with the formula `value=value*coefficient+new_value*(1-coefficient)` on each call end where new_value is the MOS average on call end. Initial value for `value` is `4.409405954387269`. **Changing this value may restart signaling process.** Default value: empty API: rtp_mos_average_m
MOS Average L Coefficient	Coefficient for long-term MOS calculation. The value is updated with the formula `value=value*coefficient+new_value*(1-coefficient)` on each call end where new_value is the MOS average on call end. Initial value for `value` is `4.409405954387269`. **Changing this value may restart signaling process.** Default value: empty API: rtp_mos_average_l

Table 34: RTP handling Parameters

5.2.13 SEMS Parameters

These parameters define the behavior of the ABC-SBC engine, specifically the SEMS signaling and media processor. They are primarily intended for troubleshooting and performance tuning, and should only be modified when necessary.

Parameter Name	Description
Use raw sockets	Performance optimization technique for sending RTP packets on Linux systems with a slow UDP stack. **Changing this value may restart signaling process.** Default value: disabled API: use_raw_sockets
Default Destination Blacklist TTL	Defines how long (ms) unavailable IP destinations are maintained on a blacklist to which no SIP traffic is sent by default. For Call Agent, a specific value may be entered in the CA parameters. See IP Blacklisting: Adaptive Availability Management. **Changing this value may restart signaling process.** Default value: 0 API: default_bl_ttl
Load q850_reason call control module	If enabled, the module for processing Q.850 reasons will be loaded. The file cc_q850_reason.conf is empty by default and must be provided as a custom local template (/data/local-templates/etc/sems/cc_q850_reason.conf). **Changing this value may restart signaling process.** Default value: disabled API: q850_reason
General RT priority	General thread real-time priority value. **Changing this value may restart signaling process.** Default value: 5 API: sems_rt_prio
SIP thread RT priority	Real-time priority for SIP threads. **Changing this value may restart signaling process.** Default value: 50 API: sems_sip_server_rt_prio
RTP RT priority	Real-time priority for RTP threads. **Changing this value may restart signaling process.** Default value: 90 API: sems_rtp_receiver_rt_prio
Session processor RT priority	Real-time priority for session processor threads. **Changing this value may restart signaling process.** Default value: 50 API: sems_session_processor_rt_prio

Parameter Name	Description
Media processor RT priority	Real-time priority for media processor threads. **Changing this value may restart signaling process.** Default value: 90 API: sems_media_processor_rt_prio
Websocket ping-pong interval in seconds	Interval in seconds for keepalive ping-pong messages on WebSocket signaling interfaces. Use 0 to disable. **Changing this value may restart signaling process.** Default value: 20 API: ws_ping_pong
Soft limit for out-of-dialog transactions (event logging only)	Maximum number of out-of-dialog (OOD) server transactions before triggering an alert event. Applies when creating a new OOD transaction not linked to an existing dialog. Use 0 to disable. **Changing this value may restart signaling process.** Default value: 0 API: tr_soft_limit_ood
Hard limit for out-of-dialog transactions (event + reply 503)	Maximum number of active server transactions not related to existing dialogs. When exceeded, new requests are rejected with "503 Overloaded" and a monitoring event is generated. Use 0 to disable. See section Server Transaction limits for details. **Changing this value may restart signaling process.** Default value: 0 API name: tr_hard_limit_ood
Event throttling for soft/hard OOD limit	Limits the number of events generated by soft/hard out-of-dialog transaction limits to one per type (soft/hard) per configured time interval (in seconds). Use 0 to disable. See section Server Transaction limits for details. **Changing this value may restart signaling process.** Default value: 0 API name: tr_limit_ood_throttle
Soft limit for in-dialog transactions (event logging only)	Number of active in-dialog server transactions that, if exceeded, will trigger an alert event. Applies only to transactions tied to an existing dialog. Use 0 to disable. See section Server Transaction limits for details. **Changing this value may restart signaling process.** Default value: 0 API name: tr_soft_limit_dlg

Parameter Name	Description
Hard limit for in-dialog transactions (event + reply 503)	Maximum number of in-dialog server transactions allowed. When exceeded, new requests are rejected with "503 Overloaded" and a monitoring event is generated. Use 0 to disable. See section Server Transaction limits for details. **Changing this value may restart signaling process.** Default value: 0 API name: tr_hard_limit_dlg
Event throttling for soft/hard DLG limit	Limits the number of events generated by soft/hard DLG transaction limits to one per type (soft/hard) per configured time interval (in seconds). Use 0 to disable. See section Server Transaction limits for details. **Changing this value may restart signaling process.** Default value: 0 API name: tr_limit_dlg_throttle
Loop detection secret	Used to create a special branch tag. When a new request is received containing this tag in the first Via-HF, it is rejected with "482 Loop Detected". Accepted values: • empty: disables the feature • auto: generates an automatic secret string • string: use provided string **Changing this value may restart signaling process.** Default value: auto API name: loop_detection_secret
Strict URI user charset check	If enabled, the user part of a URI is restricted to the characters defined by RFC3261 (ABNF rules). If disabled, SBC allows more characters as long as URI parsing is not broken. **Changing this value may restart signaling process.** Default value: enabled API name: strict_uri_user_charset
TCP connection idle timeout in milliseconds	Sets the TCP connection idle timeout (in ms) on signaling interfaces. Use 0 to disable. **Changing this value may restart signaling process.** Default value: 0 API name: tcp_idle_timeout
TCP send timeout for signaling interfaces	Corresponds to TCP_USER_TIMEOUT from tcp(7). **Changing this value may restart signaling process.** Default value: 20000 API name: tcp_user_timeout

Parameter Name	Description
Delay after startup to ignore limits	Delay (in seconds) during which CAPS and other limits are ignored after SBC startup. **Changing this value may restart signaling process.** Default value: 30 API name: inhibit_caps_timer
RESTful interface - verify https peer	If enabled, HTTPS peer certificate validity is checked in RESTful queries. **Changing this value may restart signaling process.** Default value: enabled API name: rest_verify_https_peer
Custom CA path for request queries	Custom CA certificate path for SEMS REST module. If provided, every outgoing HTTPS request uses this CA. SBC automatically adds the CA to the trusted chain. Process: • Copy CA to `/etc/pki/ca-trust/source/anchors/` • Run `update-ca-trust` **Changing this value may restart signaling process.** Default value: empty API name: rest_ca_file
User-agent string	Sets the SIP User-Agent header value if specified. **Changing this value may restart signaling process.** Default value: empty API name: user_agent_string
Unprocessed events limit	Maximum size of Redis write queue when Redis is offline. If exceeded, new writes are ignored. Set 0 to disable limit. **Changing this value may restart signaling process.** Default value: 20000 API name: redis_write_queue_limit
Unprocessed events limit warning threshold	Threshold of queued writes (see `Unprocessed events limit` option) at which SEMS logs WARN syslog messages. Use 0 to disable warnings. **Changing this value may restart signaling process.** Default value: 5000 API name: redis_warn_threshold
Log every monitored destination state	If enabled, every state of a monitored destination (changed or not) is logged at INFO level. **Changing this value may restart signaling process.** Default value: disabled API name: dest_monit_log_every_state

Parameter Name	Description
Terminate calls on SBC shutdown or restart	If enabled, SBC terminates calls during shutdown or restart. **Info** When using HA, it may be better to keep calls alive across SBC restarts. In such cases, keep this option disabled. Default value: disabled API name: terminate_calls
DNS Resolver Timeout	Timeout (in ms) for DNS query responses. **Changing this value may restart signaling process.** Default value: 100 API name: dns_resolver_timeout
DNS Cache Renew Period	Time (in s) to refresh DNS cache entries before expiry. **Changing this value may restart signaling process.** Default value: 300 API name: dns_cache_early_renew
DNS cache grace period	Time (in s) to keep expired DNS cache entries before removal. **Changing this value may restart signaling process.** Default value: 300 API name: dns_cache_grace_period

Table 35: SEMS Parameters

5.2.14 SIPREC Parameters

This section contains the configuration parameters related to SIP Recording (SIPREC).

Parameter Name	Description
SIPREC media interface	SBC interface to be used for sending RTP media towards the SIPREC server. **Changing this value may restart signaling process.** Default value: empty API name: siprec_outbound_rtp_interface
SIPREC outbound interface	SBC interface to be used for sending SIP signaling towards the SIPREC server. **Changing this value may restart signaling process.** Default value: empty API name: siprec_outbound_interface
SIPREC SIP timer A (ms)	SIP Timer A used towards the SIPREC server. **Changing this value may restart signaling process.** Default value: 500 ms API name: siprec_timer_A

Parameter Name	Description
SIPREC SIP timer B (ms)	SIP Timer B used towards the SIPREC server. **Changing this value may restart signaling process.** Default value: 32000 ms API name: siprec_timer_B
SIPREC SIP timer C (ms)	SIP Timer C used towards the SIPREC server. **Changing this value may restart signaling process.** Default value: 180000 ms API name: siprec_timer_C
SIPREC SIP timer F (ms)	SIP Timer F used towards the SIPREC server. **Changing this value may restart signaling process.** Default value: 32000 ms API name: siprec_timer_F
SIPREC SIP timer L (ms)	SIP Timer L used towards the SIPREC server. **Changing this value may restart signaling process.** Default value: 32000 ms API name: siprec_timer_L
SIPREC SIP timer M (ms)	SIP Timer M used towards the SIPREC server. **Changing this value may restart signaling process.** Default value: 8000 ms API name: siprec_timer_M
SIPREC SIP timer T2 (ms)	SIP Timer T2 used towards the SIPREC server. **Changing this value may restart signaling process.** Default value: 4000 ms API name: siprec_timer_T2

Table 36: SIPREC Parameters

5.2.15 SIP Parameters

This panel allows configuring default SIP timer values in the SBC including some extra parameters. For more details, see section SIP Timers.

Parameter Name	Description
SIP timer A (ms)	SIP Timer A used towards SIP peers. **Changing this value may restart signaling process.** Default value: 500 ms API name: sip_timer_A
SIP timer B (ms)	SIP Timer B used towards SIP peers. **Changing this value may restart signaling process.** Default value: 32000 ms API name: sip_timer_B

Parameter Name	Description
SIP timer C (ms)	SIP Timer C used towards SIP peers. **Changing this value may restart signaling process.** Default value: 180000 ms API name: sip_timer_C
SIP timer D (ms)	SIP Timer D used towards SIP peers. **Changing this value may restart signaling process.** Default value: 32000 ms API name: sip_timer_D
SIP timer E (ms)	SIP Timer E used towards SIP peers. **Changing this value may restart signaling process.** Default value: 500 ms API name: sip_timer_E
SIP timer F (ms)	SIP Timer F used towards SIP peers. **Changing this value may restart signaling process.** Default value: 32000 ms API name: sip_timer_F
SIP timer G (ms)	SIP Timer G used towards SIP peers. **Changing this value may restart signaling process.** Default value: 500 ms API name: sip_timer_G
SIP timer H (ms)	SIP Timer H used towards SIP peers. **Changing this value may restart signaling process.** Default value: 32000 ms API name: sip_timer_H
SIP timer I (ms)	SIP Timer I used towards SIP peers. **Changing this value may restart signaling process.** Default value: 5000 ms API name: sip_timer_I
SIP timer J (ms)	SIP Timer J used towards SIP peers. **Changing this value may restart signaling process.** Default value: 32000 ms API name: sip_timer_J
SIP timer K (ms)	SIP Timer K used towards SIP peers. **Changing this value may restart signaling process.** Default value: 5000 ms API name: sip_timer_K

Parameter Name	Description
SIP timer L (ms)	SIP Timer L used towards SIP peers. **Changing this value may restart signaling process.** Default value: 32000 ms API name: sip_timer_L
SIP timer M (ms)	SIP Timer M used towards SIP peers. **Changing this value may restart signaling process.** Default value: 8000 ms API name: sip_timer_M
SIP timer T2 (ms)	SIP Timer T2 used towards SIP peers. **Changing this value may restart signaling process.** Default value: 4000 ms API name: sip_timer_T2
Add Q850 header to timer expiration's CANCEL	If enabled, a Q850 header is added to the CANCEL generated by a timer expiration. Currently, only Timer C is supported. **Changing this value may restart signaling process.** Default value: disabled API name: sip_Q850_expiration
Terminate dialog upon failure replies for in-dialog OPTIONS	Terminates the dialog if in-dialog OPTIONS request fails with reply that should cause dialog termination. Affects only INVITE-based dialogs (i.e., calls). RFC 5057 replies causing termination: 404, 410, 416, 482, 483, 484, 485, 502, 604. Additionally, SBC handles replies 408 and 480 the same way. Purpose: cope with poorly implemented SIP UAs that cannot handle in-dialog OPTIONS correctly. **Changing this value may restart signaling process.** Default value: enabled (terminate the dialog) API name: remote_disappeared_on_options
Remove filtered m-lines	If enabled, removes media lines filtered out by media whitelist/blacklist. If disabled, they remain in the SDP but are marked as inactive instead. This option is applied globally on all calls with active media whitelist or blacklist (see Media Type Filtering). Purpose: cope with poorly implemented SIP UAs that cannot handle inactive media streams correctly. **Changing this value may restart signaling process.** Default value: disabled (i.e. mark media lines as inactive) API name: remove_filtered_mlines

Parameter Name	Description
Try filling missing rtpmap in disabled media	On SDP answers with disabled media, SBC tries to fill in missing a=rtpmap and a=fmtp lines by copying them from the offer if absent. **Changing this value may restart signaling process.** Default value: disabled Available since: 5.3 API name: guess_missing_rtpmap
Filter forced transports	If enabled, removes media lines that do not match outbound transport forced by the Force RTP/SRTP action (see RTP and SRTP Interworking). These lines are left in SDP but converted to the required transport if not enabled. For example: Caller is sending one audio stream over RTP and another audio stream over SRTP (commonly used when SRTP is configured as optional on a phone). SRTP is forced in outbound rules on SBC. If Filter forced transports option is "off" SBC forwards SDP with two audio streams to the callee both of them over SRTP. If this option is "on" SBC forwards SDP with just one audio stream over SRTP to the callee. This option is applied globally on all calls using Force RTP/SRTP action. Purpose: cope with interoperability issues of SIP UAs that cannot handle multiple media streams of the same type. **Changing this value may restart signaling process.** Default value: disabled (i.e. convert the media lines to the forced transport) API name: filter_forced_transports
Call transfers using late offer-answer	If enabled, uses offer-less INVITE when generating a new call leg during call transfer (unattended call transfer or call transfer replacing non-local call). Reliable, but many SIP UAs fail to implement late offer-answer correctly. **Changing this value may restart signaling process.** Default value: disabled API name: xfer_using_late_oa

Parameter Name	Description
Predefined payloads for call transfers	Comma-separated list of codecs to add into SDP of INVITE generated during call transfer (unattended call transfer or call transfer replacing non-local call). If empty, only codecs used in the call are included, which may cause issues if the peer does not support them. Only simple codecs can be used (no parameters can be specified). Example: PCMU,PCMA **Changing this value may restart signaling process.** Default value: empty API name: xfer_static_payloads
Force outbound interface	If enabled, UDP packets are forced to use the system interface attached to the outbound call agent. **Info** This option relies on operating system capabilities that have heavy limitations. Especially, when forcing the outbound interface, the Linux IP stack will set the source IP on its own, which might lead to unwanted effects (invalid source IP that e.g. SEMS might not be using at all). In many cases, this option will not effect the desired functionality and is not recommended. Manually configured source IP based policy routing is the preferred method. **Changing this value may restart signaling process.** Default value: disabled API name: force_outbound_if

Table 37: SIP Parameters

5.2.16 SRTP Parameters

These parameters define the security handshake for Secure RTP (SRTP). SRTP is always used for WebRTC and may also be used with certain SIP devices that support encryption.

Parameter Name	Description
DTLS certificate file	Certificate file. Optional. Leave empty for self-signed certificate. This is the recommended configuration, since other certificates may cause DTLS packets to become too large and fail to traverse NATs due to IP fragmentation. **Changing this value may restart signaling process.** Default value: empty API: dtls_certificate
DTLS private key file	Private key file. Optional. **Changing this value may restart signaling process.** Default value: empty API: dtls_key

Parameter Name	Description
DTLS CA list file	CA list file. Optional. Leave empty if don't want to verify the remote DTLS certificates with a CA. **Changing this value may restart signaling process.** Default value: empty API: dtls_ca_list
DTLS handshake timeout (ms)	Duration in milliseconds for the DTLS handshake to complete before terminating the call. Use 0 to disable. **Changing this value may restart signaling process.** Default value: 5000 API: dtls_timeout
SRTP crypto-suite AES_CM_128_HMAC_SHA1_32	Enables or disables the crypto suite AES_CM_128_HMAC_SHA1_32. It should remain enabled unless required otherwise for interoperability. **Changing this value may restart signaling process.** Default value: enabled API: srtp_support_hmac_sha1_32_suite
SRTP crypto-suite AES_CM_128_HMAC_SHA1_80	Enables or disables the crypto suite AES_CM_128_HMAC_SHA1_80. It should remain enabled unless required otherwise for interoperability. **Changing this value may restart signaling process.** Default value: enabled API: srtp_support_hmac_sha1_80_suite
SRTP crypto-suite AES_256_CM_HMAC_SHA1_80 (SDES only)	Enables or disables the crypto suite AES_256_CM_HMAC_SHA1_80 (SDES only). It should remain enabled unless required otherwise for interoperability. **Changing this value may restart signaling process.** Default value: enabled API: srtp_support_aes_cm_256_suite
SRTP crypto-suite AEAD_AES_128_GCM	Enables or disables the crypto suite AEAD_AES_128_GCM. It should remain enabled unless required otherwise for interoperability. **Changing this value may restart signaling process.** Default value: enabled Available since: 5.2 API: srtp_support_aead_aes_128_gcm_suite

Parameter Name	Description
SRTP crypto-suite AEAD_AES_128_GCM_8 (SDES only)	Enables or disables the crypto suite AEAD_AES_128_GCM_8 (SDES only). It should remain enabled unless required otherwise for interoperability. **Changing this value may restart signaling process.** Default value: enabled Available since: 5.2 API: srtp_support_aead_aes_128_gcm_8_suite
SRTP crypto-suite AEAD_AES_256_GCM	Enables or disables the crypto suite AEAD_AES_256_GCM. It should remain enabled unless required otherwise for interoperability. **Changing this value may restart signaling process.** Default value: enabled Available since: 5.2 API: srtp_support_aead_aes_256_gcm_suite
SRTP crypto-suite AEAD_AES_256_GCM_8 (SDES only)	Enables or disables the crypto suite AEAD_AES_256_GCM_8 (SDES only). It should remain enabled unless required otherwise for interoperability. **Changing this value may restart signaling process.** Default value: enabled Available since: 5.2 API: srtp_support_aead_aes_256_gcm_8_suite
SRTP crypto-suite preference order	Comma-separated list of crypto suites, e.g., 'AEAD_AES_256_GCM, AEAD_AES_128_GCM'. Suites offered by the remote endpoint take precedence. Suites supported but not listed here are appended at the end according to default order. **Changing this value may restart signaling process.** Default value: empty Available since: 5.2 API: srtp_support_suite_order
MTU to use with DTLS (ms)	MTU value (in milliseconds) to use with DTLS. Use 0 for the interface default. **Changing this value may restart signaling process.** Default value: 0 API: dtls_config_mtu
DTLS fallback MTU (ms)	MTU value (in milliseconds) used only if the normal MTU is too large and causes socket errors. **Changing this value may restart signaling process.** Default value: 1200 API: dtls_config_fallback_mtu

Parameter Name	Description
DTLS legacy server connect	Enables or disables support for connections to legacy DTLS servers that do not support secure renegotiation extension. By default, the DTLS handshake is terminated if this is not enabled. See RFC5746 sections 3.4 and 4.1 for details. WARNING: Enabling this reduces security and should only be used when the DTLS servers the SBC connects to have renegotiation disabled. **Changing this value may restart signaling process.** Default value: disabled Available since: 5.4 API: dtls_legacy_server_connect

Table 38: SRTP Parameters

5.2.17 Signaling SSL

This section contains the configuration parameters related to signaling SSL.

Parameter Name	Description
Revoked certificates (CRL) file	CRL file holding a list of revoked certificates. Used only by the SEMS signaling process. **Changing this value may restart signaling process.** Default value: empty API: ssl_crl_file
Minimal supported TLS version	The minimal supported TLS version on signaling interfaces. Supported values: tls1, tls1.1, tls1.2, tls1.3. **Changing this value may restart signaling process.** Default value: tls1.2 API: tls_min_version Possible values over API: • tls1 • tls1.1 • tls1.2 • tls1.3
TLS cipher list	The supported TLS cipher list for signaling interfaces. OpenSSL syntax is used. **Changing this value may restart signaling process.** Default value: HIGH:!aNULL:!MD5:!DH:!SHA1:!SHA256:!SHA384:!RSA API: tls_cipher_list
TLS EC curves list	Colon-separated list of EC curves used with TLS for signaling interfaces. Each item is a curve NID or name, for example “P-521:P-384:P-256”. **Changing this value may restart signaling process.** Default value: empty API: tls_ec_curves_list

Parameter Name	Description
Dump TLS session keys to file	If enabled, TLS session keys will be dumped to a file for diagnostics (/data/pcap/tls_keys). Requirements: This option must be enabled if downloading a diagnostic bundle from the GUI that includes both pcap files and TLS keys. Otherwise, only pcap files are included. Limitations: WebRTC interface is not supported. **Changing this value may restart signaling process.** Default value: disabled API: dump_tls_keys

Table 39: Signaling SSL Parameters

5.2.18 Syslog Parameters

These parameters allow fine-tuning of the syslog daemon behavior. They are primarily useful when syslog messages are configured to be sent to an external system.

Parameter Name	Description
Log level	Sets the global SEMS syslog log level. See section Reference of Log Level Parameters for the full list of options. **Changing this value may restart signaling process.** Default value: 2 API: sems_log_level
Syslog facility	Syslog facility to use for logs generated by the main SBC processes. Possible values: 'daemon', 'user', 'local0' … 'local7'. Default value: daemon API: syslog_facility Possible values over API: • daemon • user • local0 • local1 • local2 • local3 • local4 • local5 • local6 • local7
Enable remote syslog servers	If enabled, syslog messages will also be sent to an external syslog host in addition to being written to the local file- system. Default value: disabled API: syslog_remote_enabled
Remote syslog server address	IP address of the external syslog server. Default value: 127.0.0.1 API: syslog_remote_ip

Parameter Name	Description
Remote syslog server port	Port number on which the external syslog server listens. Default value: 514 API: syslog_remote_port
Remote syslog transport	Transport protocol for the external syslog server. Supported values: 'udp', 'tcp' or 'tls'. Default value: udp API: syslog_remote_transport Possible values over API: • udp • tcp • tls
Log level for remote syslog server	Minimum severity of messages sent to the external syslog server. Supported values: 'emergency', 'alert', 'critical', 'error', 'warning', 'notice', 'info', 'debug'. Default value: error API: syslog_remote_loglevel Possible values over API: • emergency • alert • critical • error • warning • notice • info • debug
Log files rotation frequency	Frequency of rotating local log files. Supported values: "daily", "weekly", or "monthly". Default value: daily API: syslog_logrotate_freq Possible values over API: • daily • weekly • monthly
Number of old log files to keep	Number of rotated log files to retain before deletion. Default value: 7 API: syslog_logrotate_times
Secondary remote syslog server address	IP address of a secondary external syslog server. Use empty to disable. Default value: empty API: syslog_remote_ip2
Secondary remote syslog server port	Port number of the secondary external syslog server. Default value: 514 API: syslog_remote_port2

Parameter Name	Description
Secondary remote syslog transport	Transport protocol for the secondary external syslog server. Supported values: 'udp', 'tcp' or 'tls'. Default value: udp API: syslog_remote_transport2 Possible values over API: • udp • tcp • tls
Log level for secondary remote syslog server	Minimum severity of messages sent to the secondary external syslog server. Supported values: 'emergency', 'alert', 'critical', 'error', 'warning', 'notice', 'info', 'debug'. Default value: error API: syslog_remote_loglevel2 Possible values over API: • emergency • alert • critical • error • warning • notice • info • debug
Send CDRs to remote syslog server	If enabled, Call Detail Records (CDRs) will be included in log messages sent to the remote syslog server. Default value: disabled API: syslog_remote_cdr_enabled Possible values over API: • daily • weekly • monthly

Table 40: Syslog Parameters

5.2.19 System Monitoring Parameters

These parameters configure email alerts that are triggered when system resources are excessively used. The same email address is also used for setting up automatic Let's Encrypt certificate renewal.

Parameter Name	Description
Email address for sending alerts	Email address to which important alerts such as excessive CPU usage reports are sent. Leave empty to disable email alerts. This address is also used for Let's Encrypt TLS certificate auto-renewal. Default value: empty API: alert_email

Parameter Name	Description
From email address for sending alerts	Available since: 4.3 Email address used as the “From” field in alert emails. If empty, the system default is used. Default value: empty API: alert_emailfrom
SMTP mail server address for sending alerts	Address of the SMTP server used to send alert emails. When SBC runs in a container, a local mail relay is unavailable, so an external mail server must be configured. Default value: empty API: alert_mailserver
SMTP mail server port	Port number of the SMTP mail server. Default value: 25 API: alert_mailserver_port
Use secure connection to SMTP mailserver	Defines whether the SMTP connection should be encrypted, and if so, whether TLS or STARTTLS is used. Default value: no API: alert_mailserver_tls Possible values over API: • no • TLS • STARTTLS
SMTP mail server authentication	Enables or disables SMTP authentication. If set to “on”, the authentication type is chosen automatically. Default value: off API: alert_mailserver_auth Possible values over API: • off • on
Username for SMTP authentication	Username used for SMTP authentication, if authentication is enabled. Default value: empty API: alert_mailserver_user
Password for SMTP authentication	Password used for SMTP authentication, if authentication is enabled. Default value: empty API: alert_mailserver_pass
Send system monitoring data to ABC Monitor	If enabled, system monitoring data is sent to the remote ABC Monitor together with signaling events. Default value: enabled API: collectd_to_console
Send extended system info emails when threshold is exceeded	If enabled, detailed system information is included in alert emails when monitoring thresholds are exceeded. Default value: enabled API: alert_email_info

Parameter Name	Description
Extended info emails frequency	Defines the minimum interval between sending extended system info emails. The value must include a suffix: “min”, “hour”, or “day” (e.g., 10min, 1hour, 1day). Default value: 1hour API: alert_freq
Check status of system interfaces	If enabled, system network interfaces are periodically checked and alert events are generated if errors are detected. Specific check types can be configured using related parameters. Default value: enabled API: checknet_interfaces_chk
Include only MI/SI interfaces in status check	Defines whether the system network interface checks include only media and signaling interfaces or all interfaces. Default value: enabled API: checknet_interfaces_misi
Interval for the system interfaces checks	Interval in seconds at which system network interfaces are checked. Default value: 300 API: checknet_wait

Table 41: System Monitoring Parameters

5.2.20 Probe Parameters

This section contains the configuration parameters related to Probe operation and management.

Parameter Name	Description
Log level	General log level for probe logging. Default value: info API: probe_log_level Possible values over API: • “-3”: bug • “-2”: critical • “-1”: error • “0” : warning • “1” : notice • “2” : info • “3” : debug
Log message options	Log message options (bitmask of flags): • 1: short filename:line • 2: long format • 4: timestamp • 8: backtrace short • 16: backtrace long Default value: 1 API: probe_log_opt

Parameter Name	Description
Parse / capture log level	Log level for parsing/capturing SIP messages. Default value: notice API: probe_parse_log_level Possible values over API: • "-3": bug • "-2": critical • "-1": error • "0" : warning • "1" : notice • "2" : info • "3" : debug
Parse / capture log message options	Log message options for parsing/capturing SIP messages (bitmask of flags): • 1: short filename:line • 2: long format • 4: timestamp • 8: backtrace short • 16: backtrace long Default value: 4 API: probe_parse_log_opt
Extra REGISTER expiration delta	Extra REGISTER expiration delta (in seconds) used to absorb delayed re-REGISTERs. Default value: 30 API: probe_reg_exp_delta
Ignore port number when comparing contacts	If enabled, the port number is ignored when comparing contacts (but not AORs). Default value: disabled API: probe_contact_ignore_port
Interval for generating statistics/counters events	Interval in seconds for generating statistics and counters events. Default value: 300 API: probe_stats_interval
Event rate max values list	List of event rate maximum values. If any is exceeded, the event is blacklisted. Use 0 to disable a rate. Example for 2 rates: 20,240,0. Default value: 16,240,600 API: probe_evr_limits
Event rate time intervals	Time intervals corresponding to event rate limits, used to calculate the rates. Use 0 to disable. Default value: 1s,1m,10m API: probe_evr_intervals

Parameter Name	Description
Event rate exponential back-off min	Report blacklisted events only if the number of "repetitions" is a multiple of this value and 2^k and n < Event rate exponential back-off max. To disable exponential back-off for the blacklist re-reporting, set the value to 0 or to Event rate exponential back-off max. Default value: 100 API: probe_evr_conseq_report_min
Event rate exponential back-off max	Maximum repetition count for re-reporting blacklisted events. Use 0 to disable. Default value: 10000 API: probe_evr_conseq_report_max

Table 42: Probe Parameters

5.3 Reference of Log Level Parameters

In several ABC SBC configuration places, the log reporting levels may be configured. The ABC SBC allows to set the logging levels both globally and by functional areas. The increase log level may help with troubleshooting however caution is advised. Increased log level can dramatically degrade system performance.

This reference provides explanation how to set the proper logging level. Log levels are represented with an integer value and have the following possible values:

- 0 / ERROR
- 1 / WARNING
- 2 / INFO
- 3 / DEBUG

If only log-level is set, it is used globally. The log level can be changed however for only some specific functional area by preceding the value with "Category:Subcategory=" expression. Multiple such expressions can be combined with each other using semicolon as shown in the following example:

```
1;SIP:Transaction=3;SDP:Parser=3;RTP:*=3;PLUGIN:sbc=3
```

This example sets the default log level to 1, whereas SIP transaction machine, SDP parser, RTP engine and SBC logic reports at log level 3.

Category	Subcategory
Core	• Main • Config • Thread • Timer • Events • SessionContainer • SessionProcessor • SessionWatcher • MediaProcessor • Plugin • Utils

SIP	• Ctrl • Parser • Transport • Transaction • Dialog • OfferAnswer • Session • Registration • Subscription • DNS • Blacklist
B2B	• B2BSession • B2BMedia
SDP	• Parser • MimeBody
RTP	• Stun • RtpPacket • RtcpPacket • RtpTransport • RtpStream • RtpAudio • DTMF
SRTP	• SRTP • SDES • DTLS • ZRTP • Socket
AUDIO	• Audio • AudioFile • AudioMixer • Conference • Playlist • Prompt • Jitter
PLUGIN	• all • dsm • sbc • redis_store • websock • reg_agent • cc_gui • cc_gui_rules • cc_gui_xfer • cc_gui_dest • sbc_replication • rest • webconference • xmlrpc2di

Table 43: Log Level categories

5.4 Reference of Call Agent Configuration Parameters

This reference lists all Call Agent configuration parameters used in ABC SBC. These parameters take effect on any traffic that is specific to a Call Agent without need to place any additional action into the Call Agent's rulebase.

The actions are grouped as follows:

- Destination Monitor Parameters
- Failover Parameters
- Registration Agent Parameters
- Topology Hiding Parameters
- Firewall Blacklisting Parameters
- Security Parameters
- SIP Timer Parameters
- Resolver Parameters

5.4.1 Destination Monitor Parameters

These parameters configure health checks on Call Agents by sending OPTIONS requests at regular intervals.

Depending on whether the Call Agent responds to these OPTIONS requests, its destinations can be added to a destination blacklist, thereby removing them from the pool of potential target destinations.

If blacklisting is enabled, it is also possible to configure a list of SIP reply codes that, if received, will also mark the destination as unavailable.

Parameter Name	Description
Monitoring interval (sec)	Interval between sending OPTIONS-based health-checks to the monitored Call Agent. If zero, no monitoring takes place.
Max-Forwards	Value of Max-Forwards header field in the health checking OPTIONS requests.
Blacklist TTL (seconds)	The period of time an unresponsive address remains on the blacklist. If zero, blacklisting is not used.
Unavailable on Reply Codes	Comma separated list of SIP Response codes

5.4.2 Failover Parameters

These parameters allow to define when a new destination is tried. By default, this occurs if a destination fails to respond within a predetermined timeframe. However, it is possible to configure a list of SIP Response codes that will produce the same effect, triggering a failover to the next available destination.

It is also possible to add destinations that have been found to be unresponsive (either through a timer or due to a specific SIP reply code) to s destination blacklist.

See the Section IP Blacklisting: Adaptive Availability Management for additional information.

Parameter Name	Description
On Reply Codes	Comma separated list of SIP Response codes
Blacklist TTL (seconds)	The period of time an unresponsive address remains on the blacklist. If zero, blacklisting is not used.
Blacklist grace timer (milliseconds)	Additional period of time to provide a safety buffer in case that conflicting timers occur along a SIP path.

5.4.3 Registration Agent Parameters

Registration agent allows to register the ABC SBC with a third-party SIP service be sending pre-defined REGISTER requests as described in the Section Registration Agent. The following Call Agent parameters define if such a registration agent shall be active and how its registration parameters shall be formed.

Registration agent credentials set on the call agent will also be used for authorization requests for calls as well, unless overridden by auth actions.

Parameter Name	Description
Enabled	Turns a registration agent on or off.
URI domain.	Domain name to be used in REGISTER requests URIs
URI name.	User name to be used in REGISTER request URIs
Display name	Display names as included in the From header-field of the REGISTER requests
auth name	SIP User id as used in the authentication header fields. May be different from user names in URIs.
auth password	SIP user password used in the digest authentication
Contact	Content of the Contact header-field in the REGISTER requests. Specific usernames may be chosen to make it easier to identify incoming requests coming to addresses registered using the registration agent.
Contact HF Params	Semi-colon separated header parameters to add to the Contact header.
Additional headers	\r\n-separated headers to add to the requests. I.e. 'x-my-hdr: v1\r\nx-my-hdr2: v2'.
Registration interval (seconds)	Time between subsequent registrations are sent
Retry interval (seconds)	Period of time to keep till the next attempt when the previous failed
Next Hop (IP address)	Address of a destination to which a request will be sent
Registrar affinity	Binding of the registrar. `Sticky` mode records the reply IP/Port/Transport and initially tries that for refreshing the registration. `Lazy` is same as sticky except that it does a lookup of the recorded reply IP address in the SBC's internal reverse-dns cache table and discards the record if it is not found in the cache. `Active` does not record the reply address at all. Limitations: • In `Lazy` mode, only the IP address is checked for existance in the cache and not port & transport. • Items in the reverse-dns cache are still considered valid after their expiry, until the duration specified in the `DNS Cache Grace Period` global configuration passes. Available since: 5.2
Bulk Contact	Turn on to support the SIP bulk contact registration form as described in RFC3680.

Table 44: Registration Agent Parameters

5.4.4 Topology Hiding Parameters

The Section Topology Hiding discussed purpose and use of Topology Hiding. The following options enable/disable this functionality for the respective Call Agents.

Parameter Name	Description
Enabled	Turning this option replaces occurrences of IP addresses in well-known header-fields of SIP signaling with those of the ABC SBC .
Cross-Realm	If enabled, topology hiding is used even when signaling ingress and egress realms are the same.

5.4.5 Firewall Blacklisting Parameters

Automated IP address blocking is discussed in the Section Automatic IP Address Blocking. Several attributes defined what kind of Call Agent behavior adds to the score that may eventually lead to blacklisting of the source IP address.

Parameter Name	Description
Sanity	If turned on, invalid SIP messages add to the auto-blocking score and may lead to blocking of their originator. Otherwise they are silently ignored.
Auth	If enabled, failed authentication add to the auto-blocking score and may lead to blocking of their originator. Otherwise only events are reported but no further action is taken.

5.4.6 Security Parameters

Parameter Name	Description
Don't expect authentication	Don't expect any authentication on this call agent. Drops any 401/407 replies from this agent. Removes 'Authorization' and 'WWW-Authorization' sent towards this agent, 'Proxy-Authenticate' and 'WWW-Authenticate' headers received from this agent.

5.4.7 SIP Timer Parameters

Parameter Name	Description
SIP Timer [X]	Allows setting SIP timers per agent. Each SIP timer set overrides the global configuration.
Failover reduce factor	`Failover reduce factor` is used to divide B, F & M timers when the destination call agent has a backup CA. This allows for a faster failover. Leaving it empty uses the default value of 4.

Please refer to section SIP Timers for more details.

5.4.8 Resolver Parameters

Parameter Name	Description

Nameserver IP addresses (comma-separated)	DNS nameservers to use while communicating through this call-agent. Each unique nameserver configuration has its own reverse-dns-cache. If parameters of two configurations are the same (i.e. regardless of the order, same set of nameservers & bind-to-ip address flag resolves to the same physical interface) then they share a common reverse-dns cache. This rule covers the DNS configuration in the signaling interfaces as well. If this is set, it will get used as soon as this call agent is chosen. Until the CA is chosen, either the signaling interface's configuration will be effective, or if that does not exist, system's configured nameservers will be used. When trying to find a source call-agent that is identified by DNS, a DNS reverse-cache lookup is done using the source IP. This look-up follows these steps until a match is found: 1. A reverse-cache search is done on the resolver of each call-agent that is assigned to the signaling interface that the SIP message came from. If such a call-agent does not have a nameserver configuration, then the look-up is done on the system-level resolver for that call-agent. 2. A reverse-cache search is done on the resolver of the signaling interface. If the signaling interface does not have a nameserver configuration, then the look-up is done on the system-level resolver. The same look-up logic applies to finding a destination call-agent as well. This configuration is per-leg. Registration Agent also makes use of this configuration.
Bind to signaling interface	Strictly use the underlying physical interface of the signaling interface of this call-agent.

5.5 Default Audio Files

Most of the prompts' sample rate is 8000. It isn't necessarily required, as `sems` resample them. Note that wideband samples may sounds nicer.

All of the meet-me actions' offer two sets of defaults audio prompts:
- */usr/lib/sems/audio/webconference* (English)
- */usr/lib/sems/audio/webconference/de* (German)

Multi-lingual support can be used in conjuncture with those 2 directories. See Multi lingual conferencing announcements for more information about that feature.

5.5.1 Join meet-me conference

The following prompts are used by multiple meet-me conference configuration.

Audio file	Content
General audio files	
contact_support	Please contact support.
enter_pin	Please enter your code, then press the pound key.
entering_conference	You are now entering your conference room.
first_participant	**ton** Welcome, you are the first participant in the conference.

Audio file	Content
max_attempt_reached	We are sorry you are having problems. Please try later or contact customer support.
please_enter_room	Please enter your conference room, then press the pound key.
please_enter_your_code	Please enter your code, then press the pound key.
short_pin	This PIN is too short. Please try again.
simple_pin	This PIN is too simple. Please try again.
room_created	Room created.
timeout_enter_pin	This input unfortunately took to long. Please try again later.
yourcodeis	Your code is
yourroomnumberis	Your room number is
welcome	Welcome. This is FRAFOS conference.
wrong_pin	This code is not correct. Please try again.
wrong_pin_bye	This code is not correct. Please try again later or contact customer support.
x_welcome_and_prompt	Welcome this is FRAFOS' conference. Please enter your code, then press the pound key.
join_sound / drop_sound	**biip / buup**
Security PIN audio files	
andpinis	And the PIN is
create_secu_pin	Please enter a PIN for the new room, followed by the pound key.
enter_secu_pin	Please enter the PIN of the room, followed by then pound key.
repeat_secu_pin	Please repeat the new PIN, followed by the pound key.
secu_pin_set_to	PIN set to
secu_pin_3_digits	Sorry, security PIN must be at least 3 digits.
Record username audio files	
current_participants_are	**The current participants in the conference** are...
just_joined_conf	... just joined the conference
just_leaved_conf	... just leaved the conference
recording_1_2_3	To keep this recording, please press 1, To replay the recording, please press 2. To record your name again, please press 3.
say_ur_name	Please, say your name after the tone. Then, press the pound key.
timeout_record	Username recording timed out. Please try again later.
ur_name_is	Your recorded name is
Generate room audio files	
ask_if_gen	To enter a conference room, please press 1. To create a new room, please press 2
error_persist_room	An error occurred while saving the new room and PIN.
generating_room	We are now creating a conference room

Audio file	Content
repeat_or_enter	Press 1 to hear room number an pin again. Press 2 to go into your room.
timeout_generate_room	This input unfortunately took to long. Please try again later.
Multi lingual support audio files	
select_lang	To continue in English, press one. Um auf Deutsch vor zu fahren, drücken Sie bitten bis zwei

Please note that digits prompts are also needed. When multi-lingual isn't used, files are expected to be found in the same directory as the matching Conferencing' global config. In case of multi-lingual, files are expected to be found in the *digits/* sub-directory.

SBC support two kind of number echoing:
- left to right: Forty two
- right to left: Zwei Und Vierzig

LtR expected files are the following:
- digits: 0.wav, 1.wav, 2.wav, 3.wav, 4.wav, 5.wav, 6.wav, 7.wav, 8.wav, 9.wav
- multiple of 10: 10.wav, 20.wav, 30.wav, 40.wav, 50.wav, 60.wav, 70.wav, 80.wav, 90.wav
- tens: 11.wav, 12.wav, 13.wav, 14.wav, 15.wav, 16.wav, 17.wav, 18.wav, 19.wav
- 21 to 99: x2.wav, x3.wav, x4.wav, x5.wav, x6.wav, x7.wav, x8.wav x9.wav

RtL expected files are the following:
- digits: 0.wav, 1.wav, 2.wav, 3.wav, 4.wav, 5.wav, 6.wav, 7.wav, 8.wav, 9.wav
- multiple of 10: 10.wav, 20.wav, 30.wav, 40.wav, 50.wav, 60.wav, 70.wav, 80.wav, 90.wav
- tens: 11.wav, 12.wav, 13.wav, 14.wav, 15.wav, 16.wav, 17.wav, 18.wav, 19.wav
- 21 to 99: 2x.wav, 3x.wav, 4x.wav, 5x.wav, 6x.wav, 7x.wav, 8x.wav 9x.wav

5.5.2 Meet-me set PIN audio prompts

Audio file	Context	Content
setPin_welcome	Use for welcome	'welcome, you can set a pin for your personal conference room with the number' …
setPin_welcome_set	Used to welcome when the security PIN is already set.	'welcome, your personal conference room with the number' …
setPin_enter_pin	Used to prompt the security PIN	'please enter the security pin of the room number' …
setPin_change_pin	Used to prompt the security PIN	'please hang up if you want to keep it, otherwise'
setPin_repeat_pin	Used to confirm the security PIN user	'please repeat the pin and and press the pound key'
setPin_pin_set	Used in case of success	'your pin was successfully set, thanks you.'
setPin_pin_dont _match	Used when user PIN don't match	'the pin numbers you've enter does not match. Please try again, and enter a new PIN, followed by the pound key.'
setPin_failed	Used in case of failure	'please hang up if you want to keep it, otherwise'

Table 45: Audio prompts

5.5.3 Two-Factor authentication

Audio file	Context	Content
2fa_greeting	Use for welcome	'Please enter the two factor authentication PIN number that was set for this line
2fa_pin_correct	Used in case of success	'that is correct, thanks you. Please hold the line to be connected'
2fa_failed	Used to prompt the security PIN	'I'm sorry you're having entering the pin number. Please hold the line to be connected to the help desk.'
2fa_pin_wrong	Used to prompt the security PIN	'sorry this is not correct. Please enter the 2 factor authentication pin number that we set for this line.'

Table 46: Audio prompts

5.6 Reference of Default Port Numbers

The reference lists port numbers the ABC SBC, Cluster Config Manager and ABC Monitor uses. It is particularly useful when considering firewall policies for firewalls placed in front of the ABC SBC. The reference lists default port numbers, transport protocols, container opening the port, service listening on that port and the interface on which the respective applications are permitted. In addition to the SBC interfaces (see SBC Interfaces), some applications may be listening on all interfaces while some management applications are using the loopback interface for internal communication.

Note that while the ABC SBC only accepts traffic on the ports and interfaces specified in the following specification, further restrictions may apply. Signaling is only accepted from well-defined Call Agents and certain traffic may be blacklisted (see Manual SIP Traffic Blocking).

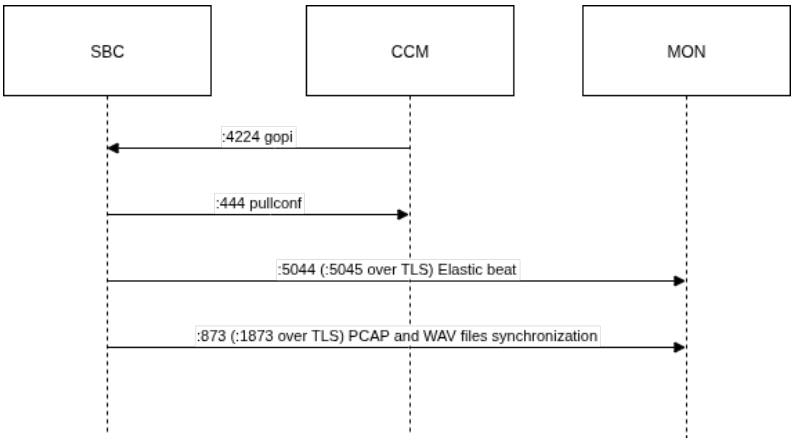

Port	Container	Description
22	ABC SBC	(ssh / TCP) Secure shell server. Used for remote management. Value 0 can be used for default port, which is 24. It can be set using ssh app on SBC interface.
25	ABC SBC	(SMTP / TCP) Local Email relay. Used to forward email alerts. From outside perspective it acts as a client.
161	ABC SBC	(SNMP / UDP) Internal SNMP management.

443	Cluster Config Manager	(HTTPS / TCP) Administrative GUI.
444	Cluster Config Manager	(HTTPS / TCP) Allow ABC SBC to download new configuration and upload status file to the Cluster Config Manager.
1443	ABC SBC	(HTTPS / TCP) XML-RPC provisioning.
3306	Cluster Config Manager	(TCP) MySQL database.
5060	ABC SBC	(sip / UDP, TCP) SIP signaling.
5061	ABC SBC	(sip / TLS)SIP signaling over TLS.
6379	ABC SBC	(TCP) redis replication, if HA is used.
8080, 8081	ABC SBC	(TCP) SIP over Websocket WebRTC.
8090	ABC SBC	(TCP) XML-RPC remote programming interface
10000 to 60000	ABC SBC	(UDP) Audio/video media.
15441, 4443		(TCP) webconference demo. available only on request.
1444	ABC SBC	(TCP) RESTful port for AWS SNS, disabled by default.
4247, 4248	ABC SBC	(TCP) sbc-eventbeat-[1,2] Expose live metrics and statistics about the redis queues event processing. Local use on localhost, SBC node only.
4224	ABC SBC	(HTTPS) sbc-gopi unified RESTful json API, allowing various interactions with the ABC SBC node. By default the 4224 is protected by firewall and accessible only from CCM. However if IMI interface is configured with a valid TLS certificate and *Verify peer certificate* is enabled, the firewall rule is removed and port protection relies on TLS verification only.

Additional fixed source port numbers shall be opened for the ABC SBC acting as client reaching outside servers as listed in the following table:

SBC Client Port	**Description**
NTP/123/UDP	Time Synchronization
domain/53/UDP	DNS Resolver

Other applications running on the ABC SBC use external applications while locally binding to ephemeral ports.

Remote Server Port	**Description**
syslog/514	remote syslog facility if configured under Global Config / syslog-ng
rsync/873	remote PCAP/WAV storage if enabled under Global Config / replicate recordings / traffic log
rsync/1873	remote PCAP/WAV storage if enabled under Global Config / replicate recordings / traffic log, using TLS if secure connection to ABC Monitor enabled
6379,redis	redis replication and event generation to a ABC Monitor
16379,redis	redis replication and event generation to a ABC Monitor over TLS if enabled

5.7 Reference Interface Parameters

The following parameters can be defined at interface level:

Parameter Name	Description
force_via_address	When enabled, incoming requests are replied to the address shown in their Via header field. This conforms to the RFC3261 specification but often fails to traverse NATs and also permits a reflection attack through the ABC SBC.
no_rport	When enabled, the *rport* parameter is not added to the first *Via* header field.
wspath_xxx	The option, where xxx can be set as needed, sets up an HTTP proxy from path /xxx on HTTPS 443 port (or other port number if using a non-standard one) to the Websocket port on localhost . (It has to be used only on interface using system interface "lo".

5.8 Reference Application Interface Options

Starting 4.5, the ABC SBC offers the possibility to configure some application option per logical interface, allowing a better control over which process is listening on which port.

Some applications require a TLS profile assigned to corresponding SBC or applied interface.

From SBC release 5.2 and up to 5.4, the following applications were available:

- SSH
- Media
- Signaling
- WebSocket signaling
- SNMP
- Prometheus Pull Service
- TURN server for websocket
- Local monitoring query service
- PCAP query service
- HA vrrp and call state replication
- Local webconf API
- Management for host
- HTTP proxy
- HTTP redirect
- frafos-logprovider
- Log files provider (`frafos-logprovider` has been replaced by it in 5.1).
- Local packet classifier

Starting 5.5, the following applications have been made available, but not configurable:

- Unified SBC management service

It act as a replacement for:

- Local monitoring query service
- PCAP query service
- Local webconf API
- Management for host
- Log files provider
- Local packet classifier
- Prometheus Pull Service

As such, those applications are deprecated, marked as `For nodes up to version 5.4 only` in the UI. If added to an interface, they'll be highlighted in beige. Please note that this UI behavior can be tweaked depending on the selected value for the CCM compatibility mode.

Following the migration from a systemd based environment to the new-generation s6 one, the following applications were dropped:

- SSH
- SNMP
- TURN server for websocket
- HTTP proxy
- HTTP redirect

The following key words in this document are to be interpreted as:

- `IMI`: internal management interface
- `SI`: signaling interface
- `MI`: media interface
- `WS`: websocket interface
- `CX`: custom interface

See Legacy application for legacy applications.

5.8.1 Supported by the current release

Unified SBC management service

> **Info**
>
> Application available since SBC release 5.5.

The SBC node now ships a single RESTful API listening on the port **4224**: `gopi`.

This API unifies the functionality of the legacy APIs, aka: Local monitoring query service, PCAP query service, Local webconf API, Management for host, Log files provider, Local packet classifier and `goconf`.

As such, the API allows to:

- fetch metrics from various sources (redis list, SEMS's xmlrpc API)
- generate and serve PCAP files based on an aggregation of the one available on the SBC node file system
- expose information and actions related to SEMS's web-conferencing features
- run various administrator tasks from RESTful endpoints
- access the node's file system log files
- partial interaction with the node's firewall
- publish new configuration and provisioned tables to the node
- fetch the node status

The API will by default listen on any IP (`0.0.0.0`) over HTTPS, allowing the Cluster Config Manager to push the initial configuration to it. The default listening IP address can be changed by setting the environment variable `SERVER_IP` to the desired IP.

When the ABC SBC container is started in push mode (`CONFIG_MODE=push`), it normally still performs an initial configuration pull from the Cluster Config Manager at startup, using `MASTER` as the Cluster Config Manager address and `CONFIG_USER` / `CONFIG_PASS` for HTTP basic authentication on the pullconf API. Setting the environment variable `NO_BOOTSTRAP_PULL` to `1` skips this initial pull, so that the node simply waits for the Cluster Config Manager to push the first configuration.

Pushing a configuration to the ABC SBC API requires a valid API token. To avoid a chicken-and-egg situation on a freshly started node that has not yet received any configuration, `NO_BOOTSTRAP_PULL=1` additionally causes token authorization on the ABC SBC API to be disabled until a configuration containing tokens is activated on the node.

See Configuring the SBC container for a complete description of the environment variables that control how the ABC SBC container connects to the Cluster Config Manager at startup.

For every valid configuration published to the SBC node, the process will (re-)start listening on any `IMI` interface, using the attached TLS profile. The API's swagger documentation can be accessed from any browser where the SBC node's IP can be reached, at `https://[SBC IMI IP]:4224/`.

Please note that the API is **not** configurable, albeit one may enable/disable it from an `IMI` interface, allowing the reduce the bloated UI. One may also hide the application from the CCM UI screen by setting the compatibility mode to something **lower than** `5.5` (ex: `5.4`). In any ways, please note that having the application enabled on a 5.4 or earlier node **won't** affect that particular node.

Media

The media application impacts SBC communication handling. Note that this application only has effect on SBC node.

The application is exclusive and mandatory to `MI` interface.

The port range specifies a UDP port range used for media traffic, and does not use TLS.

Parameter Name	Description
Ports	Port range on which SBC may open a socket for media communications.
TOS	This sets "type of service" field in IP packets header. Default value: 184

Signaling

The signaling application impacts SBC communication handling. Note that this application only has effect on SBC node.

If "TLS Port" is not empty, a TLS profile is required.

The application is exclusive and mandatory to `SI` interface.

Parameter Name	Description
Port	Ports on which SBC will open a signaling socket.
TLS Port	(optional) TLS port on which SBC opens a socket for secured signaling communication. SEMS will not use TLS for signaling, if this value is left empty.
Interface Options	Special interface options. Note: allowed values are `force_via_address` and `no_rport`.
TOS	This sets "type of service" field in IP packets header. Default value: 104
Greylist	Enables usage of greylist filter.

Resolver Nameservers	DNS nameservers to use while communicating through this interface. Each unique nameserver configuration has its own reverse-dns cache. If the parameters of two configurations are the same (i.e. regardless of the order, the same set of nameservers & bind-to-ip addr. flag resolves to the same physical interface), then they share a common reverse-dns cache. This rule covers the resolver configuration in the call-agents as well. Requests inbound from this interface will attempt to use the resolver configuration of this interface for DNS requests, until a call-agent is chosen. After that, if the call-agent has a resolver configuration, it will override this. When trying to find a source call-agent that is identified by DNS, a DNS reverse-cache lookup is done using the source IP. This look-up follows these steps until a match is found: 1. A reverse-cache search is done on the resolver of each call-agent that is assigned to the signaling interface that the SIP message came from. If such a call-agent does not have a nameserver configuration, then the look-up is done on the system-level resolver for that call-agent. 2. A reverse-cache search is done on the resolver of the signaling interface. If the signaling interface does not have a nameserver configuration, then the look-up is done on the system-level resolver. The same look-up logic applies to finding a destination call-agent as well.
Resolver Bind To Interface	Strictly use the underlying physical interface to send the DNS requests.

WebSocket signaling

The websocket application allows signaling communication over websocket interface.

If “TLS enabled” is set, a TLS profile is required.

The application is exclusive and mandatory to `WS` interface.

Parameter Name	Description
Port	Listening port of the websocket server.
TLS enabled	Enable secure communications.
Interface Options	Special interface options. Note: value must start by `wspath_`.
Greylist	Enables usage of greylist filter.
TCP keep-alive	Set TCP keep-alive value (seconds) on WS. 0 disables it. I.e. if it is set to `120`, then the SBC will try to send a TCP keep-alive after 120 seconds of inactivity and wait another 120 seconds for a response. This will happen `probes` (below) times before timing out the connection.
TCP keep-alive probes	How many times to try to send keep-alive message without getting a response.

HA vrrp and call state replication

Please note that the application only have effect if HA is configured and used.

The redis HA replication application uses internal redis protocol for it’s communications.

The application is exclusive and mandatory to `IMI` interface.

Parameter Name	Description
Port	Port on which call state redis will be listening. Note: value not editable (`6379`).
Enable TLS	Make use of the interface' TLS profile to authenticate and secure redis HA. Redis internal protocol is used for communications. Please note that, if used, the TLS certificate **must** either be loaded with a matching CA certificate or be registered by the node's system CA (currently latest debian:12). Note 1: disable by default. Note 2: incompatible with the "default certificate" due to the CA certificate requirement. Note 3: TLS profiles' "Verify peer certificate" option isn't taken into account.

Probe management

FRAFOS is preparing to launch a new product: a passive monitoring probe designed for VoIP environments. This new solution delivers VoIP monitoring capabilities comparable to those offered by the ABC SBC, generating detailed events based on observed SIP signaling and RTP traffic without actively interfering with network flows. Designed for maximum deployment flexibility, the passive monitoring probe can be operated as a lightweight container anywhere within the VoIP infrastructure. It captures live SIP messages and RTP packets, producing familiar event streams and traffic captures (including PCAP files), consistent with the output of the ABC SBC and Monitor systems.

To support the integration of probe nodes, the CCM has been extended to manage them in the same manner as existing SBC nodes. A new "Node Role" field has been introduced, allowing users to specify either "sbc" or "probe" when configuring a node. At this stage, the introduction of the "probe" node type is purely preparatory and does not alter the behavior or configuration of any existing SBC nodes.

Parameter Name	Description
Port	Port for the internal http server, provides API for stats and probe's mgmt settings.
Capturing interface	Interface to capture packets from - one specific or "any" Default value: "any"
Capturing filter	Berkley packet filter for capture Default value: "port 5060"
Enable PCAP traces	Writes individual PCAP files for each call.
Blacklisted events	List of event types that should be blacklisted: event1, event2, … e.g.: msg-probe, other-failed, other-timeout, other-ok, parse-error. Blacklisted events will not be generated.
VXLAN ports	UDP destination ports used for vxlan (tunnel endpoint). Packets arriving on these ports will be automatically decapsulated. An empty lists means disabled. The standard vxlan port is 4789.

5.9 Command Line Reference

The administrative GUI is the preferred way of the ABC SBC. However there are cases like the initial configuration and/or automation when accessing the ABC SBC via Command Line is useful.

5.9.1 Configuration Management

CLI	Purpose	Reference
ccm-backup	backup ABC SBC configuration	ABC SBC Recovery Procedure
ccm-restore	restore a configuration backup	ABC SBC Recovery Procedure
sbc-init-config	Interactively configures the IP address or DNS name of the main configuration node from which the ABC SBC node will pull its configuration. The preferred method for new deployments is to use environment variables instead; see Configuring the SBC container.	Configuring the SBC container
ccm-publish-config	Activate the current SBC configuration and make it available for all nodes.	
cluster-config-export	Export configuration in JSON format.	
cluster-config-import	Import the configuration exported by cluster-config-export command.	
ccm-config	Manage CCM configuration options. It is equivalent to CCM->CCM config GUI screen.	
sbc-cfg-cli • fetch • convert • generate • apply • compile • defaults • template	Script which is used to generate and apply config on SBC side. It also can be used to get a config template for different services like sems or keepalived for example.	

5.9.2 User Management

CLI	Purpose	Reference
sbc-add-user	Add new GUI user or add a user to a group.	CLI User Management
sbc-del-user	Remove a GUI user or remove a user from a group.	CLI User Management
sbc-list-groups	Get list of existing user groups	CLI User Management
sbc-list-users	Get list of SBC users	CLI User Management
sbc-user-passwd	Change password of SBC user, unlock user locked by too many login attempts or reset two-factor authentication secret.	CLI User Management

5.9.3 Low-Level CLI

CLI	Purpose	Reference
sbc-loglevel action [loglevel]	Shows or sets the logging level for the ABC SBC signaling process. Action is either 'get' to retrieve current value or 'set' to set it. Loglevel takes category and level. Log files are stored in the directory /var/log/frafos	Reference of Log Level Parameters

5.9.4 HA CLI

In previous ABC SBC releases up to 4.1, the high availability solution used was based on Pacemaker. The ABC SBC 4.2 was a transitional release that removed the Pacemaker based HA solution, before new Keepalived based HA solution was introduced in 4.3 release.

CLI	Purpose	Reference
sbc-ha-offline	Forces the node when run to be put forcibly into HA FAULT state	
sbc-ha-online	Clears the forcibly set HA FAULT state set by sbc-ha-offline	
sbc-ha-status	Shows the node's current HA status, which can be Unknown, MASTER, BACKUP, FAULT and STOP.	High Availability statuses

5.10 Reference of Used Open-Source Software

The key components of ABC SBC are built as commercial software fully owned by FRAFOS GmbH and its subsidiaries. Additionally it relies on the Linux operating systems and numerous accompanying libraries and components provided by third parties under the following license terms:

- bash , GPLv3+
- boost: Boost Software License & similar (http://www.boost.org/users/license.html)
- cronie , MIT and BSD and ISC and GPLv2+
- crontabs , Public Domain and GPLv2
- dialog , LGPLv2
- dmidecode , GPLv2+
- ethtool , GPLv2
- expat (XML parser): MIT https://sourceforge.net/p/expat/code_git/ci/master/tree/expat/COPYING
- fence-agents-all , GPLv2+ and LGPLv2+
- flite , X11-like http://www.festvox.org/flite/doc/flite_2.html
- hiredis , BSD https://github.com/redis/hiredis/blob/master/COPYING
- iLBC: BSD-like
- js , GPLv2+ or LGPLv2+ or MPLv1.1
- json-c: MIT (https://github.com/json-c/json-c/blob/master/COPYING)
- jsonxx: MIT? (https://github.com/hjiang/jsonxx/blob/master/LICENSE)
- libbcg729: GPLv3 (https://github.com/BelledonneCommunications/bcg729/blob/master/LICENSE.txt)
- libcap , LGPLv2+
- libcurl: MIT/X derivate license https://curl.haxx.se/docs/copyright.html
- libevent: BDS-like http://libevent.org/LICENSE.txt

- libisac: WebRTC license
- libopus: BSD
- libosip2 , LGPLv2+
- libpcap , BSD with advertising
- librsvg2 , LGPLv2+
- libsrtp , BSD-like https://github.com/cisco/libsrtp/blob/master/LICENSE
- libtiff , BSD-like (http://www.libtiff.org/misc.html)
- libxml2 , MIT http://www.xmlsoft.org/FAQ.html
- mailx , BSD with advertising and MPLv1.1
- mariadb-server , GPLv2 with exceptions and LGPLv2 and BSD
- monit , AGPLv3
- mysql++ , LGPLv2
- mysql-connector-c++ , GPLv2 with exceptions
- MySQL-python , GPLv2+
- nginx, BSD-like
- net-snmp , BSD http://www.net-snmp.org/about/license.html
- net-snmp-utils , BSD
- ntp , (MIT and BSD and BSD with advertising) and GPLv2
- opencore-amr: Apache V2.0
- openssh-clients , BSD
- openssl, BSD-like https://www.openssl.org/source/license.html
- opus , BSD
- pciutils , GPLv2+
- pcmisc , GPLv2+
- pcs , GPLv2
- perl-Net-SSLeay , OpenSSL
- php-cli , PHP and Zend and BSD
- php-db , PHP
- php-log , PHP
- php-mysql , PHP
- php-pear-XML-RPC , PHP
- php-pecl-runkit , PHP
- php-xmlrpc , PHP and BSD
- python , Python
- python-jinja2 , BSD
- redis , BSD
- rsync , GPLv3+
- sems-gsm , public domain
- sems-speex , modified BSD
- serweb-frmwrk , GPL
- silk: BSD-like
- spandsp (g722, DTMF): LGPL
- speex , BSD
- sqlite , Public Domain
- stunnel, GPL
- syslog-ng , GPLv2+
- sysstat , GPLv2+
- tcpdump , BSD with advertising
- vconfig , GPLv2+
- yajl (JSON): ISC license https://en.wikipedia.org/wiki/ISC_license
- wireshark , GPL+

5.11 Reference Userdata Parameters for AWS Instances

The behavior of the ABC SBC can be altered by Userdata passed to it during instance launch. See the following link for more information about Userdata: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html#instancedata-add-user-data

The ability to alter the instance behavior is often useful when instances are started using a CloudFormation template. The parameters passed through Userdata must be encoded as attribute name:value pair; name and value are separated by comma and so are the pairs.

The following table shows reserved attribute names and how they are used.

Attribute Name and Value	Description
configurl <URL>	Download an ABC SBC backup configuration (only applicable when ismater:TRUE, if the instance is backup it retrieves the configuration from its master.
cwgroup <NAME>	an additional CloudWatch Dimension to which the ABC SBC sends CloudWatch metrics; this can be used to group metrics from multiple instances; note that proper CloudWatch permissions must be set
cwregion <REGION>	if CloudWatch metrics is to be gathered in a different region than instance's own, set the CloudWatch region using this parameter
ismaster TRUE	Enforce configuration master role
master <IP address>	Run this instance as configuration backup of a master identified by an IP address.
remotebootscript <URL>	URL of a bash script that will be downloaded and sourced during instance launch. The script must be finite because the boot process doesn't continue until it completes.
rtcecdns <IP address>	Address of the primary Monitor

Note that any attribute names including custom ones can be passed via Userdata. When a remotebootscript is used and started, all the attributes are passed to it as shell variables.

An example of UserData may look like this:

```
rtcecdns,172.12.1.1, configurl, https://s3-eu-west-1.amazonaws.com/frafos-abcconfig/40014-honeypot.
sql
```

5.12 Reference XML-RPC functions

In the case that the ABC SBC administrator needs to configure large data sets to CCM GUI, it will be easier to provision those data automatically with a script as opposed to typing it in using the web-interface. This can be accomplished using the ABC SBC's XML-RPC data provisioning interface.

The following example shows a python code fragment for accessing the built-in XML-RPC provisioning server:

```
#!/usr/bin/python
from xmlrpc import client
server = client.Server('https://username:password@10.0.0.10:1443/rpc.php')
```

Note that for the python client, a question mark (?) in the password does not work. The user accessing XML-RPC interface has to be either member of **SBCrpc** group or member of another group having **XML-RPC** privilege.

For the XML-RPC access the IP address of the configuration master node has to be used. The XML-RPC is accessible by default on port 1443.

The XML-RPC interface is self documented via function `rpc.help()`. When the function is called without any argument it prints list of all available function. When function name is given as an argument to this function (`rpc.help(<function name>)`) it will return detailed help of the specified function.

For example try following calls in python:

```
print(server.rpc.help())
print(server.rpc.help('rpc.help'))
```

As of now functions for manipulate following entities are available:

- Provisioned Tables
- Call agents
- TLS profiles
- Nodes
- Logical interfaces
- System interfaces
- Maintenance mode

Bellow is list of all available XML RPC functions. Call `rpc.help(<function name>)` to get detailed help of specified function.

5.12.1 Provisioned Tables

Functions for define provisioned tables and manipulate data in them.

Function Name	Description
tables.fetch_rules($table_name, $start, $count, $key_values)	Get all rules from specified provisioned table.
tables.fetch_rule($table_name, $key_values)	Get a rule matching the key from the specified provisioned table.
tables.insert_rule($table_name, $data)	Insert rule into specified provisioned table.
tables.insert_rules($table_name, $rules)	Insert multiple rules into specified provisioned table.
tables.update_rule($table_name, $data)	Update rule of specified provisioned table.
tables.update_rules($table_name, $rules)	Update multiple rules of specified provisioned table.
tables.insert_update_rule($table_name, $data)	Try update rule of specified provisioned table. If rule with matching UUID or key columns does not exists, new rule is inserted.
tables.delete_rule($table_name, $uuids)	Delete rule(s) from specified provisioned table.
tables.delete_all_rules($table_name)	Delete all rules from specified provisioned table.
tables.commit($table_name, $msg)	Commit working version of provisioned table into use by signaling and create new working version by copying the current one.
tables.fetch()	Get all provisioned table definitions.
tables.insert($payload)	Insert provisioned table.
tables.update($payload)	Update provisioned table.
tables.delete($table_name)	Delete provisioned table.
tables.delete_room($room_name)	Delete a conference room (PIN provtable type).

For example, to introduce a new entry to the blacklist and check the outcome, the following three RPC commands must be called: *insert_rule*, *commit* and *fetch_rules*:

```
data = {"key_value":"sip:restricted@abc.com"}
print(server.tables.insert_rule('test_uri_bl',data))
print(server.tables.commit('test_uri_bl', 'new restricted used introduced'))
print(server.tables.fetch_rules('test_uri_bl'))
```

This script will result in the following list of URIs shown on the command-line output:

```
[{'key_value': 'sip:banned@abcsbc.com', 'uuid': '6c01a834-9d32-df09-0217-000000f074ee'},
{'key_value': 'sip:forbidden@abcsbc.com', 'uuid': '54d15a12-62bc-73c9-8313-000012f8ae1b'},
{'key_value': 'sip:restricted@abc.com', 'uuid': '6d831a12-88bc-7fa9-7483-000083ff992a'}]
```

Note that the routing tables have several predefined mandatory elements that must use the following conventions:

- *cagent* takes name or UUID of a call-agent
- *outbound_proxy* and *next_hop* is passed as string
- boolean parameters *next_hop_1st_rq*, *upd_ruri_host*, and *upd_ruri_dns_ip* take either 0 or 1 as value
- the enumerative parameters *route_via* takes one of the following values: *outbound_proxy*, *next_hop* or *ruri*

5.12.2 Call agents

Function Name	Description
cagents.fetch($filter)	Get call agents
cagents.insert($payload)	Insert call agent
cagents.update($payload)	Update call agent
cagents.delete($realm_name, $cagent_name)	Delete call agent
cagents.add_target($realm_name, $cagent_name, $payload)	Add target destination to call agent
cagents.del_target($realm_name, $cagent_name, $payload)	Remove target destination from call agent

5.12.3 TLS profiles

Function Name	Description
tls_profile.fetch($filter)	Get TLS profiles
tls_profile.insert($payload)	Insert TLS profile
tls_profile.update($payload)	Update TLS profile
tls_profile.delete($name)	Delete TLS profile

5.12.4 Nodes

Function Name	Description
node.fetch($filter)	Get SBC nodes
node.insert($payload)	Insert SBC node

node.update($payload)	Update SBC node
node.delete($name)	Delete SBC node

5.12.5 Logical interfaces

Function Name	Description
log_interface.fetch($filter)	Get logical interfaces
log_interface.insert($payload)	Insert logical interface
log_interface.update($payload)	Update logical interface
log_interface.delete($name)	Delete logical interface
log_interface.help_app_list()	Return list of available applications
log_interface.help_app($application)	Return detailed info about an application

5.12.6 System interfaces

Function Name	Description
sys_interface.fetch($filter)	Get system interfaces
sys_interface.insert($payload)	Insert system interface
sys_interface.update($payload)	Update system interface
sys_interface.delete($log_if_name, $owner_type, $owner_name)	Delete system interface

5.12.7 Maintenance mode

If the "maintenance mode" is activated, the SBC answers 503 to any request.

The XMLRPC interface allows to toggle a "maintenance mode" for a given node. Please use `sems-stats -c "set_shutdown 1"` to trigger the maintenance mode, or `sems-stats -c "set_shutdownmode 0"` to disable it. At any time, one may use `sems-stats -c "get_shutdownmode"` to fetch the current node status.

One may also trigger the maintenance mode via the `gopi` API (:4224), using either the `/api/v1/enable/shutdownmode` or the `/api/v1/disable/shutdownmode` endpoints.

Finally, a helper script `sbc-shutdownmode` exists. Please refer to `sbc-shutdownmode -h` for more information about it.

5.13 Reference of CCM Configuration Parameters

This reference lists all CCM configuration parameters. The configuration parameters are grouped as follows:

- Login
- LDAP Parameters
- Backup Parameters
- Management access Parameters
- SBC security Parameters
- Email Parameters
- Certbot Parameters
- Miscellaneous Parameters

Each section specifies an API category and each attribute has the *API* name listed in its description. Use these two names together to set or retrieve a value through the CCM REST API.

5.13.1 Login

Parameters related to login/logout.

API category: login

Parameter Name	Description
GUI auto-logout time	Timeout (in minutes) of inactivity after which the GUI user is automatically logged out. If changed, relogin is required. Use '0' to disable auto-logout. Default value: 15 API: gui_logouttime
Max failed login	Maximum number of failed login attempts before the user account is blocked. This protects against brute-force attacks. Use '0' to disable account blocking. Default value: 3 API: login_max_failed
Blocking period	Duration (in seconds) that the user account remains blocked after reaching the "Max failed login" limit. Default value: 300 API: login_blocking_period
Allow concurrent login	Allows a single GUI user to log in from multiple devices simultaneously. Disabled by default. Default value: enabled API: login_allow_concurent
Garbage collect timeout	Timeout (in days) after which brute-force protection data is removed from the database. Default value: 180 API: login_gc
Do not allow re-use passwords - history length	Prevents users from setting a password identical to any of the last N passwords used. This field specifies N. Use '0' to disable history length. Default value: disabled API: password_history
Password expiration	Number of days before the user password expires and must be changed. Use '0' to disable expiration. Default value: disabled API: password_expire
Minimum password length	Minimum required length of a user password. Default value: 8 API: password_min_length

Parameter Name	Description
Password strength policy	Defines the set of character types that must be included in a user password. Default value: At least one uppercase letter, one lowercase letter, and one number. API: password_strength Possible values over API: • "": No policy • "aA": At least one uppercase and one lowercase letter • "aA0": At least one uppercase, one lowercase letter and one number • "aA0.": At least one uppercase, one lowercase letter, one number and one special character
Enable two-factor authentication for LDAP users (initial value)	Enables two-factor authentication for LDAP users upon their first connection to the CCM. Default value: disabled API: login_ldap_totp_enabled
Relying Party ID for passkeys	Identifier of the WebAuthn Relying Party on whose behalf registration or authentication is performed. A credential can only be used with the same RP ID it was registered with. By default, the RP ID is set to the caller's origin's effective domain. For more information see w3.org. Default value: empty API: login_webauth_rp_id

Table 47: Login Parameters

5.13.2 LDAP Parameters

The Cluster Config Manager GUI supports two-step authentication against an LDAP server.

The first authentication step verifies a user against the LDAP server. In this step, the user DN (e.g., `uid=john,ou=People,dc=example,dc=org`) and password (e.g., `johnldap`) are used.

The second step ensures that at least one of the LDAP user's groups matches one of the existing GUI groups.

Once configured, users can log in to the Cluster Config Manager using their LDAP UID and password on the login page.

API category: ldap

Parameter Name	Description
LDAP enabled	Enables or disables LDAP authentication. Default value: disabled API: ldap_enabled
LDAP server address	Address of the LDAP server in the format `ldap://IP:PORT`, `ldap://IP`, or `ldap://my.domain`. Default value: empty API: ldap_host

Parameter Name	Description
LDAP bind distinguished name	Distinguished Name (DN) used to bind to the LDAP server for lookups. Default value: empty API: ldap_user
LDAP bind credentials	Credentials (password) for the LDAP bind user. Default value: empty API: ldap_password
Base DN of the LDAP server	Default search DN for the LDAP directory. Example: for `cn=admin,dc=example,dc=org`, the base DN is `dc=example,dc=org`. Default value: empty API: ldap_base_dn
Extra group	Additional bind DN part for user entries. For example, if a user `john` exists as `uid=john,ou=People,dc=example,dc=org`, set this parameter to `ou=People`. The system will build the full DN as `uid=[username],[extra_group],[base_dn]`. You can enter multiple groups by separating them with a semicolon. Default value: empty API: ldap_extra_group
Enable compatibility with Microsoft Active Directory LDAP	Enables compatibility mode for Microsoft Active Directory LDAP servers. Default value: disabled API: ldap_win
Set the user fetching template	Select the appropriate LDAP attribute for usernames. For Microsoft Active Directory, use `sAMAccountName`. For typical OpenLDAP setups, use `uid`. Some setups may require `cn`. Default value: 'uid' user indexing (usually OpenLDAP) API: ldap_user_template Possible values over API: • "(cn=%(username)s)" • "(uid=%(username)s)" • "(sAMAccountName=%(username)s)": Microsoft Active Directory
Set the group fetching template	Select the appropriate LDAP attribute for groups. For Microsoft Active Directory, use `memberOf`. For typical OpenLDAP setups, use `gidNumber`. Some may require `memberUid`. Default value: 'gidnumber' group indexing (usually OpenLDAP) API: ldap_group_template Possible values over API: • "gidNumber" • "memberUid" • "memberOf": usually for Microsoft Active Directory

Parameter Name	Description
Verify certificate of LDAP server	If enabled, the LDAP server's TLS certificate is verified. Default value: disabled API: ldap_reqcert
Trusted CA certificates file	Path to a PEM-formatted file containing trusted CA certificates used to verify the LDAP server's certificate. Default value: empty API: ldap_tlscacert

Table 48: LDAP Parameters

OpenLDAP configuration example

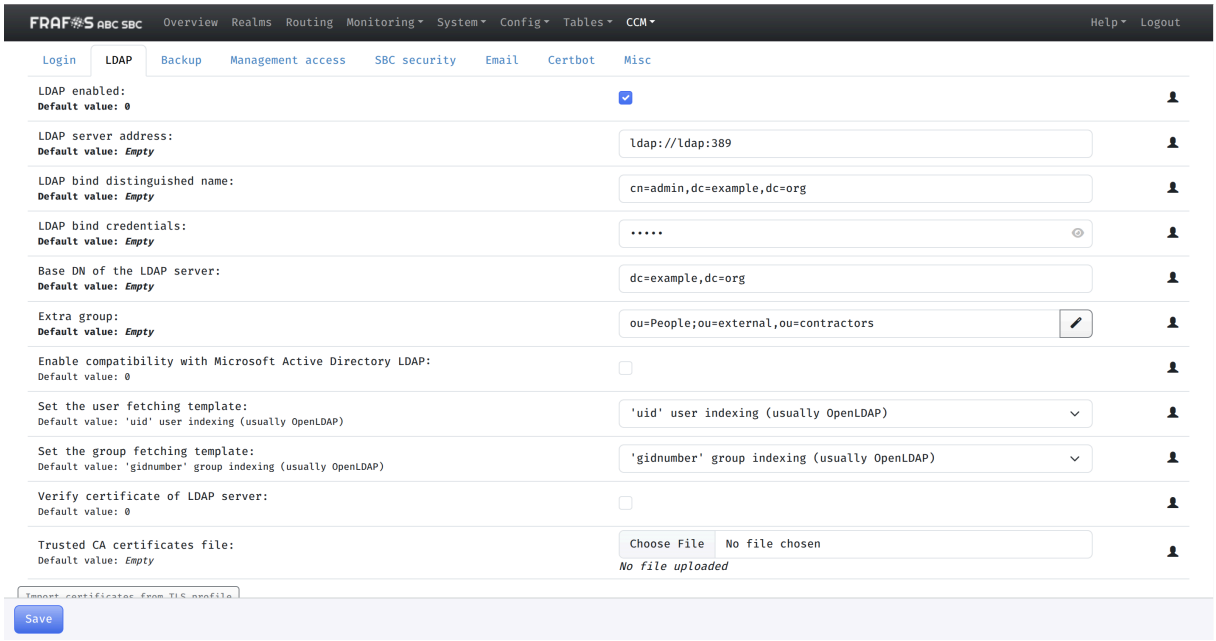

There is a docker container available on github that match the screenshot configuration : https://github.com/frafos/docker-ldap.

The image come in with 3 users (+ admin) :

User	dn	pwd	note
john	`uid=john,ou=People, dc=example,dc=org`	`johnldap`	The following example works for this user.
jane	`uid=jane,ou=People, dc=example,dc=org`	`janeldap`	The following example **doesn't** work for this user. John and Jane belong to different groups.
june	`uid=june,ou=external, ou=contractors, dc=example,dc=org`	`juneldap`	The following example works for this user.

Table 49: OpenLDAP users

In that following ldap, user `john` can be authenticated against the ldap via `uid=john,ou=People,dc=example,dc=org`. To allow an ldap user to access the ABC SBC GUI, a **GUI group name** with access to the GUI **must** match one of the primary group of the ldap user.

So we create GUI group named after the full dn of one `john` LDAP group (`cn=GUI,ou=Groups,dc=example,dc=org`) :

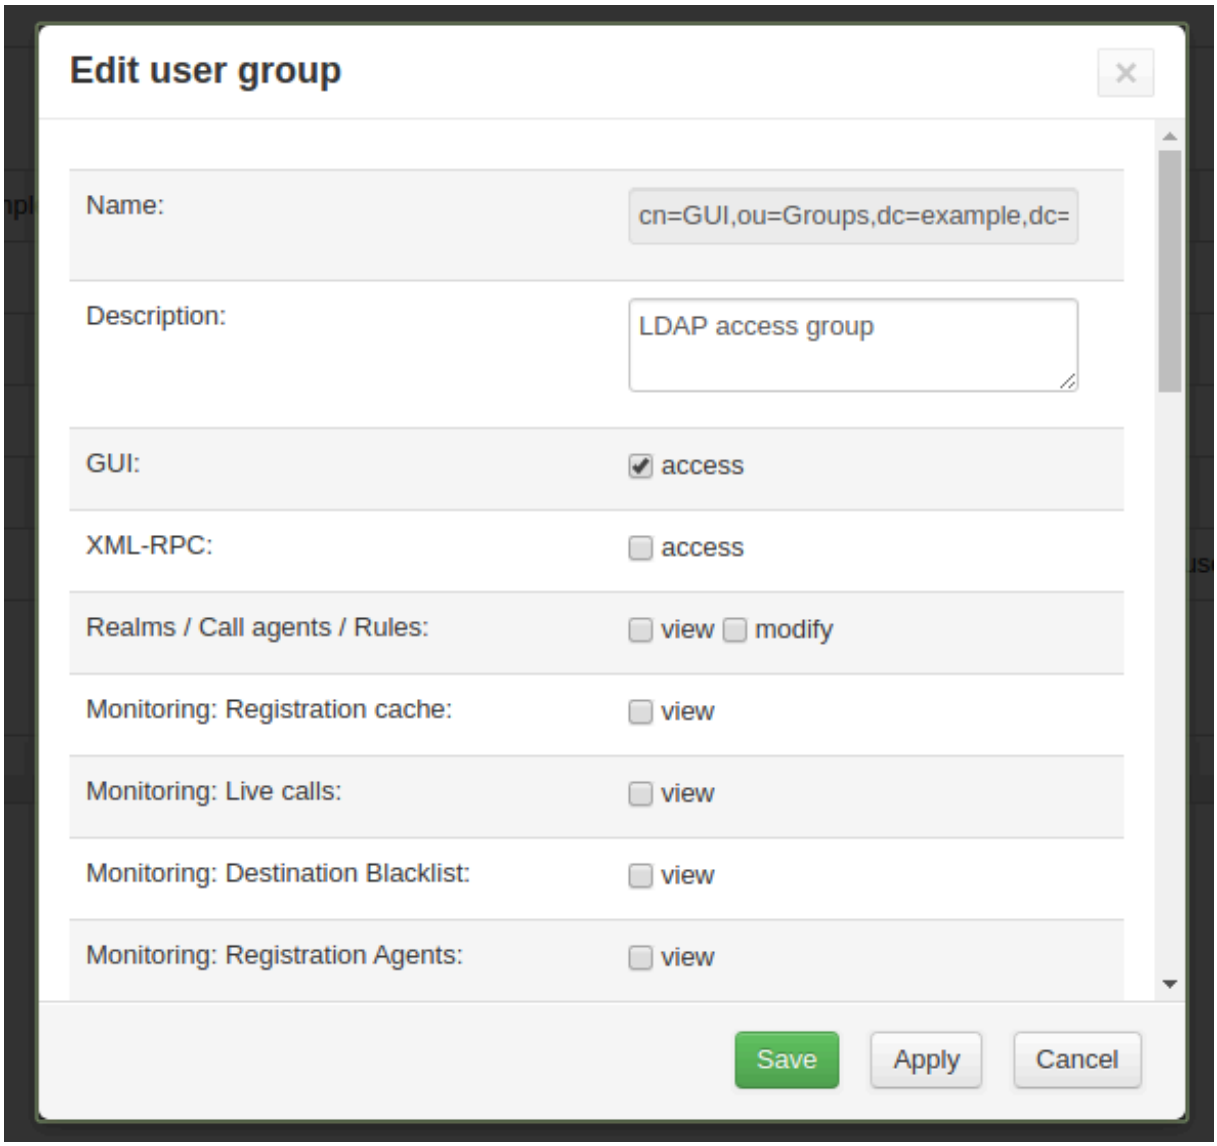

You can then login with the credential `john` and the password `johnldap`.

Note that the following credentials are also valid:

- `june/juneldap`

Note: If we want Jane to be able to access the GUI, we'll need to define another ABC SBC GUI groups, matching one of Jane ldap groups name (`cn=Mistyc,ou=Groups,dc=example,dc=org` in this case).

FreeIPA LDAP configuration example

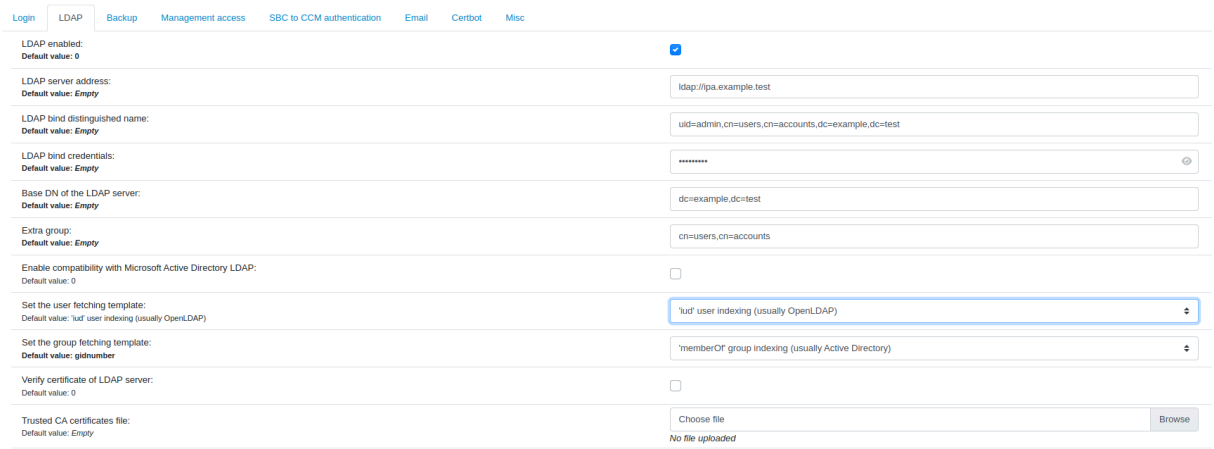

We'll skip the server configuration part for sanity reasons. We recommend to have a look at https://www.freeipa.org/page/Docker for easy setups.

In the following case, the Free IPA server was configured with default values, generating the following configuration:

```
The IPA Master Server will be configured with:
Hostname:       ipa.example.test
IP address(es): 172.42.0.142
Domain name:    example.test
Realm name:     EXAMPLE.TEST

The CA will be configured with:
Subject DN:   CN=Certificate Authority,O=EXAMPLE.TEST
Subject base: O=EXAMPLE.TEST
Chaining:     self-signed

Client hostname: ipa.example.test
Realm: EXAMPLE.TEST
DNS Domain: example.test
IPA Server: ipa.example.test
BaseDN: dc=example,dc=test
```

On the FreeIPA side, we've created an `sbcgui` group and a `john` user belonging to that group. We can query them over the ldap with the following:

```
$ ldapsearch \
  -D "uid=admin,cn=users,cn=accounts,dc=example,dc=test" \
  -w [admin password] \
  -H ldap://ipa.example.test \
  -b dc=example,dc=test 'uid=john'
(...)

# john, users, accounts, example.test
dn: uid=john,cn=users,cn=accounts,dc=example,dc=test
givenName: John
sn: Doe
uid: john
cn: John Doe
displayName: John Doe
initials: JD
gecos: John Doe
krbPrincipalName: john@EXAMPLE.TEST
gidNumber: 681800003
objectClass: top
objectClass: person
objectClass: organizationalperson
objectClass: inetorgperson
objectClass: inetuser
objectClass: posixaccount
objectClass: krbprincipalaux
objectClass: krbticketpolicyaux
objectClass: ipaobject
objectClass: ipasshuser
objectClass: ipaSshGroupOfPubKeys
objectClass: mepOriginEntry
objectClass: ipantuserattrs
loginShell: /bin/sh
homeDirectory: /home/john
mail: john@example.test
krbCanonicalName: john@EXAMPLE.TEST
ipaUniqueID: c81b0b0e-950a-11ee-8471-0242ac2a008e
uidNumber: 681800009
krbPasswordExpiration: 20231207141322Z
krbLastPwdChange: 20231207141322Z
krbExtraData:: AAIC03Flcm9vdC9hZG1pbkBFWEFNUExFLlRFU1QA
mepManagedEntry: cn=john,cn=groups,cn=accounts,dc=example,dc=test
ipaNTSecurityIdentifier: S-1-5-21-1615603866-3760360139-3083941652-1009
memberOf: cn=ipausers,cn=groups,cn=accounts,dc=example,dc=test
memberOf: cn=sbcgui,cn=groups,cn=accounts,dc=example,dc=test
```

On the CCM side, we've created a new `cn=sbcgui,cn=groups,cn=accounts,dc=example,dc=test` group with GUI permissions. Please note that the CCM group matches the LDAP full CN.

Microsoft Active Directory configuration example

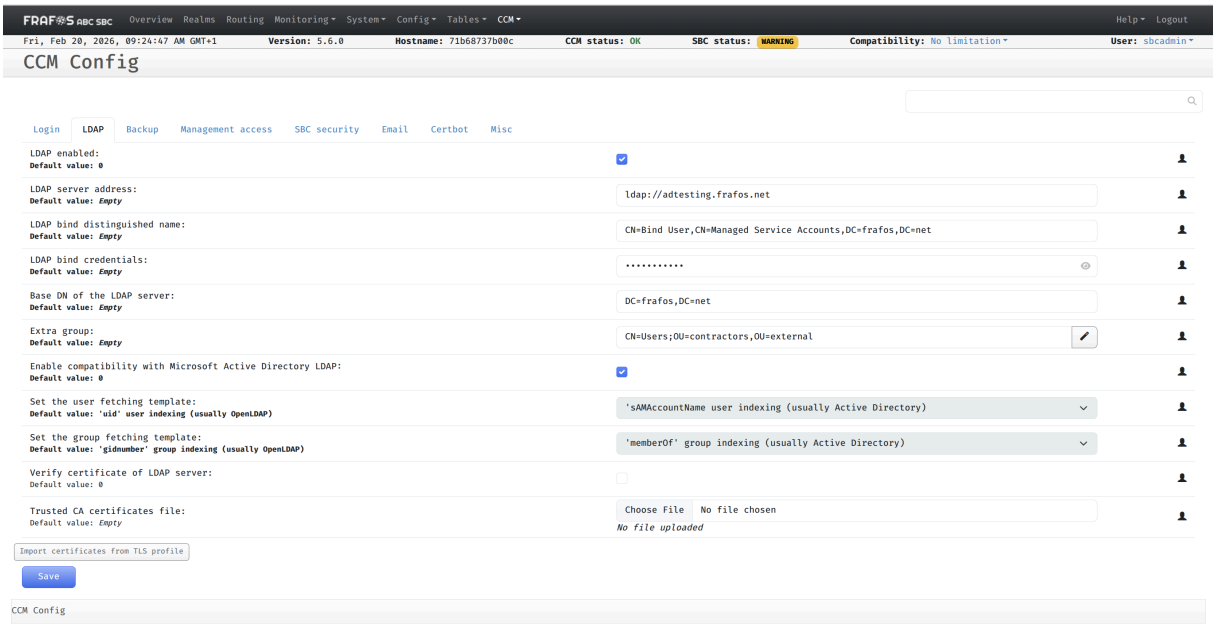

Active Directory configuration slightly differs from the OpenLDAP one, please note that all value are **case-sensitive**!

On the CCM side, we've created a new `CN=sbcadmins,CN=Users,DC=frafos,DC=net` group (this is the full name of the group) with appropriate GUI permissions. Please note that the CCM group matches the LDAP full CN.

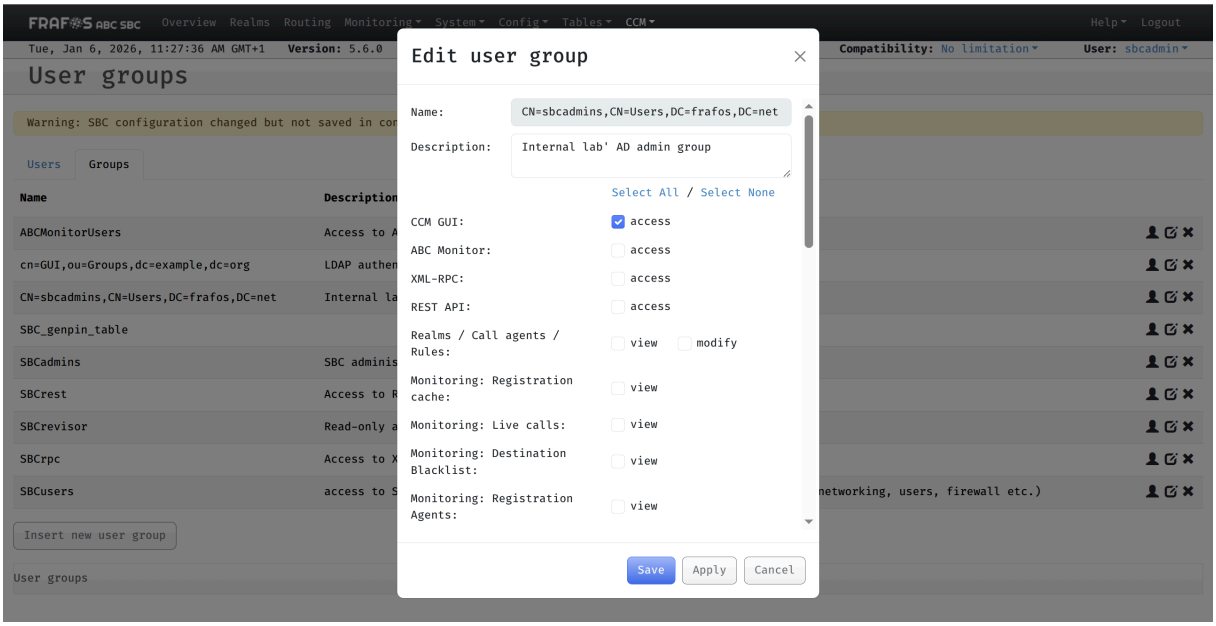

5.13.3 Backup Parameters

These parameters configure ABC SBC daily backups. For more information, see Backup and Restore Operations.

API category: backup

Parameter Name	Description
Create daily SBC configuration backups	If enabled, a daily snapshot of the SBC configuration is created as a gzipped tarball file. Automatic backups are created when a new container starts and before database upgrade is performed. This backup file allows a rollback to a previous container if required. Default value: disabled API: backup_enable
Include provisioned tables in daily or automatic backups	If enabled, daily or automatic backups also include the full content of provisioned tables. Default value: enabled API: backup_prov
Number of days to keep backups	Defines the retention period (in days) for backup files. All files matching `sbc-backup-*` in the backup directory that are older than the specified number of days are deleted during the daily backup process. Use `0` to disable automatic deletion. Default value: 7 API: backup_keep
Destination directory for backups	Specifies the target directory for daily backup files. Default value: `/data/backups` API: backup_dir

Table 50: Backup Parameters

5.13.4 Management access Parameters

API category: xmi

Parameter Name	Description
SSL certificate file for GUI, REST API and XML-RPC interface	A file containing SSL certificate in PEM format, used for the GUI, REST API, and XML-RPC interfaces. Default value: empty API: tlscert
SSL private key file for GUI, REST API and XML-RPC interface	A file containing a private key in PEM format, corresponding to the SSL certificate, used for the GUI, REST API, and XML-RPC interfaces. Default value: empty API: tlskey
Minimal supported TLS version for GUI, REST API and XML-RPC interface	Specifies the minimum TLS version to be supported on the GUI, REST API, and XML-RPC interfaces. Default value: tls1.2 API: tls_min_version

Table 51: Management access Parameters

5.13.5 SBC security Parameters

These parameters are used for securing the connection between SBCs and the CCM for services running on the IMI interface.

API category: pullconf

Parameter Name	Description
HTTP Basic Authentication username for configuration pull or status push	Username that SBC nodes use to pull the configuration from the CCM or to push status. The same value must be provided to the SBC container via the `CONFIG_USER` environment variable (see Configuring the SBC container). API: user
HTTP Basic Authentication password for configuration pull or status push	Password that SBC nodes use to pull the configuration from the CCM or to push status. The same value must be provided to the SBC container via the `CONFIG_PASS` environment variable (see Configuring the SBC container). API: pass
SSL certificate file for pullconf	A file containing the SSL certificate in PEM format. Default value: empty API: tlscert
SSL private key file for pullconf	A file containing the private key for the SSL certificate in PEM format, unencrypted (no passphrase protection). Default value: empty API: tlskey
Trusted CA certificates file	A file containing the list of trusted CA certificates against which clients' TLS certificates are verified. If intermediate CAs are used, include the full chain. Certificates must be in PEM format. Default value: empty API: tlscacert
Enable mTLS	If enabled, the CCM verifies the peer's TLS certificate against the trusted CA certificates. Default value: disabled API: mandate_client_cert
SBC API authentication token	Bearer token used by the Cluster Config Manager to authenticate against the ABC SBC node API. This token is auto-generated by the Cluster Config Manager and distributed to all ABC SBC nodes as part of their configuration. Use the **Renew** button to invalidate and regenerate the token. After confirmation, the Cluster Config Manager automatically pushes the new configuration to nodes configured in push mode. Nodes in pull mode will pick up the new token on their next pull. The previous token is kept as a fallback for nodes that have not yet received the update. Auto-generated if not set. Not accessible via API.

Table 52: SBC security Parameters

5.13.6 Email Parameters

These parameters are used to configure sending emails from the Cluster Config Manager.

API category: email

Parameter Name	Description
Email address for sending certificate and other alerts	Email address to which important alerts, such as certificate renewal failure or acquisition success, are sent. This field must be set if any Let's Encrypt certificate is expected to be used. Default value: empty API: alert_email
From email address for sending alerts	Email address used as the "From" address in alert emails. If empty, the system default is used. Default value: empty API: alert_email_from
SMTP email server address for sending alerts	SMTP server address used by CCM to send alert emails. Default value: empty API: alert_mail_server
SMTP mail server port	SMTP mail server port. Default value: 25 API: alert_mail_server_port
Use secure connection to SMTP mail server	Controls whether the SMTP connection should be encrypted using TLS or STARTTLS. Default value: no API: alert_mail_server_tls
SMTP mail server authentication	Set to `off` to disable authentication, or `on` to enable it and select the auth type automatically. Default value: off API: alert_mail_server_auth
Username for SMTP authentication	Username for SMTP authentication, if authentication is enabled. Default value: empty API: alert_mail_server_user
Password for SMTP authentication	Password for SMTP authentication, if authentication is enabled. Default value: empty API: alert_mail_server_pass

Table 53: CCM Email Parameters

5.13.7 Certbot Parameters

Cluster Config Manager certbot acts like the standard Let's Encrypt certbot. For more information, see the TLS chapter: Let's encrypt gocertbot.

API category: certbot

Parameter Name	Description
Query Let's Encrypt staging environment	If testing, query the Let's Encrypt staging environment to avoid hitting the Let's Encrypt 168-hour rate limit. Note: Staging certificates are not suitable for production use. Default value: disabled API: certbot_query_stagging
Attempt renewal X days before certificate expiration	Number of days before certificate expiration when certbot attempts renewal. This setting does not affect automatic email notifications about certificate expiration sent by Let's Encrypt. Default value: 15 API: certbot_days_renew
CRON job interval	CRON schedule (in CRON format) defining how often certbot runs to attempt renewal of certificates near expiration. Default value: 0 1 * * * API: certbot_cron_interval (the valid value is for example: `"0 1 * * *"`)
Preferred CA chain	If the CA provides multiple certificate chains, prefer the one whose issuer matches the specified Subject Common Name. If no match is found, the default offered chain will be used. Default value: ISRG Root X1 API: certbot_preferred_chain

Table 54: Certbot Parameters

The certbot is invoked automatically under the following conditions:
- By a CRON job, every night at 01:00.
- When a node successfully pulls a new configuration.
- When a configuration is successfully pushed to a node.

You can also invoke certbot manually from within a Cluster Config Manager shell by running:

```
% sbc-gocertbot -d
```

Note: For testing purposes, to avoid reaching the Let's Encrypt 168-hour rate limit, enable the *Query Let's Encrypt staging environment* option in the Cluster Config Manager configuration.

5.13.8 Miscellaneous Parameters

API category: misc

Parameter Name	Description
Automatically add new nodes	If enabled, records for new nodes that pull configuration from the configuration master will be automatically added. If disabled, the CCM will provide configuration only to nodes that are already registered in the CCM. Not yet registered nodes will be refused. Default value: enabled API: auto_add_nodes

Parameter Name	Description
Compatibility mode	If using CCM with older SBCs, the SBC version can be selected here. CCM will hide settings (e.g., rule conditions and actions, interface applications, global config values, or entire screens) that are not available in the selected SBC version. Selecting the "No limitation" option displays all settings for all supported SBC versions. Default value: <the latest SBC version> API: compatibility_mode (use "" for "No limitation", or SBC version in the format "X.Y" - for example "5.6")
Compatibility mode with secunet SBC	If enabled, the firewall control and HA configuration screens will be hidden. Default value: disabled API: secunet_compat_mode
Allow overlap of Call Agent IP ranges	If enabled, the GUI will not check whether IP address ranges of call agents overlap. Default value: disabled API: allow_ca_overlap
Address of ABC Monitor GUI	If set, a link to the ABC Monitor GUI will be added to the top bar menu under the Monitoring tab. If not set, the value of the ABC Monitor address from Event Parameters is used. Default value: empty API: monitor1_gui_url
Address of secondary ABC Monitor GUI	If set, a link to the secondary ABC Monitor GUI will be added to the top bar menu under the Monitoring tab. If not set, the value of the Secondary ABC Monitor address from Event Parameters is used. Default value: empty API: monitor2_gui_url
Name of the CCM instance	Name of this CCM instance that will be displayed in the browser tab. If not set, "CCM" will be used. Default value: empty API: ccm_name

Table 55: Miscellaneous Parameters

5.14 Reference of Token Capabilities

CCM API tokens and SBC API tokens use separate sets of capabilities. This reference lists the capabilities available for each token type and describes how SBC token capabilities map to Cluster Config Manager user permissions.

5.14.1 CCM API token capabilities

CCM API tokens authenticate against the Cluster Config Manager REST API. Their capabilities mirror the Cluster Config Manager user permission model. When creating or editing an API token, only capabilities and actions that the current user holds are available for selection.

The `restapi:access` capability is always added to API tokens automatically.

For the full list of Cluster Config Manager user capabilities, see the user group configuration in *CCM –> CCM config –> Groups*.

5.14.2 SBC API token capabilities

SBC API tokens authenticate against the ABC SBC node API. They use a dedicated set of capabilities that control access to SBC node operations.

Capability	Action	Description
sbc-config	read	Read SBC configuration
sbc-config	push	Push configuration changes to the SBC node
sbc-provtables	read	Read provisioning tables
sbc-provtables	push	Push provisioning table changes to the SBC node
sbc-status	read	Read SBC status and monitoring data
sbc-system	exec	Execute system administration commands

Table 56: SBC token capabilities

Permission mapping

When creating or editing an SBC token, the selectable capabilities are limited by the Cluster Config Manager user's own permissions. The following table shows which Cluster Config Manager user capability and action is required to assign each SBC token capability.

SBC Capability	SBC Action	Required CCM User Capability	Required CCM Action
sbc-config	read	sbc-config (Activate SBC configuration)	activate
sbc-config	push	sbc-config (Activate SBC configuration)	activate
sbc-provtables	read	prov-tables-tbl (Tables: values)	view
sbc-provtables	push	prov-tables-tbl (Tables: values)	activate
sbc-status	read	mon-node-status (Monitoring: System status)	view
sbc-system	exec	sys-restart (System: Administration)	access

Table 57: SBC token to CCM user permission mapping

The required permission for each ABC SBC API endpoint is documented in the ABC SBC OpenAPI specification as the `x-required-permission` extension on each operation.

5.15 CCM configuration API

Starting 5.5, the CCM ships a RESTful API allowing SBC nodes' JSON configuration interactions.

The API can be accessed at *http://localhost:1444*. All requests to the CCM's *:444* port, with an URI prefixed by */exportconf/*, are forward to the API.

A non exhaustive list of the available endpoints actions are:

- fetch an SBC JSON configuration for a given set of parameters (node's release/uuid, config group etc …)
- list nodes using a specific TLS profile
- list nodes using any Let's Encrypt TLS profile

The API does **not** offer any configuration to the user.

5.16 Reference of Supported Codecs

This reference lists all supported codecs by ABC SBC.

- PCMU/8000
- G721/8000
- GSM/8000
- PCMA/8000
- g722/8000
- L16/32000
- L16/16000
- L16/8000
- G726-32/8000
- G726-24/8000
- G726-40/8000
- G726-16/8000
- G729/8000
- opus/48000
- isac/16000
- iLBC/8000
- speex/32000
- speex/16000
- speex/8000
- AMR/8000
- AMR-WB/16000

5.17 SIP Timers

5.17.1 Timer Definitions

The SBC SIP stack uses timers to control the transaction state machine and parts of the dialog state machine.

In general, timers A to J are used as defined in RFC3261.

Timer L

Timer L is used to determine for how long a UAC transaction will stay in memory after a 2xx reply has been received. It corresponds roughly to the transaction state machine fixes described in RFC6026 as timer M.

Timer M

Timer M is a reduced transaction timer used to implement tranport layer failover. It defaults to a quarter (1/4) of the client transaction timer (timer B), such that up to 4 destinations can be probed before the SIP client transaction will be considered expired.

Please note that timer M is only used when there are indeed multiple destinations to be probed.

Timer BL

Timer BL is used to prolong the transaction timer for the purpose of properly handling destination blacklisting.

Once a transaction has expired and no reply has been received, timer BL is started. If a reply is received before it expires, the timer is cleared and the destination will not be placed into the destination blacklist.

If however no reply is received and the timer expires, the destination will be placed into the destination blacklist.

Please note that timer BL is only used if destination blacklisting is enabled for that particular transaction / destination.

5.17.2 Customizing Timer Values

It is to be noted that the SBC uses by default the timer values as advised by the respective RFCs. However it does not enforce that the values are set relatively to each other when setting custom values.

This means that great care needs to be taken to respect correct proportions when using values that differ from the defaults. In particular, these timers are not set dependent to T1, which is not used when configuring custom timer values.

5.17.3 SIP Timer Precision

The SBC uses so called hierarchical timers to implement SIP timers effeciently. This together with the inherent imprecision of timers and latency of non-real-time operating systems introduces a few limitations on the precision achievable by the timer settings.

SIP timers are triggered in slots of 20ms each, meaning that a timer expiry routine might be executed anywhere inside that 20ms slot, depending on many factors. It is thus common that very short timers suffer from this imprecision and might display a huge variance.

For example, settings any re-transmission timer to a value of 40ms may lead to the first re-transmission happening anywhere between 20 and 60 milliseconds, depending on the exact slot the timer has been placed in, the time at which it has been started, or even the overall number of timers hitting around that time.

5.17.4 Destination Monitoring

The destination monitoring feature uses by default a different set of SIP timers to allow for higher reactivity.

The timers used are as follows:
- Re-transmission timer (E): 2 second
- Transaction timer (F): 5 second

The values are chosen to avoid interfering with the normal traffic, however still allowing for a shorter transaction timeout to react faster to unavailability.

5.18 Legacy application

5.18.1 SSH

> **Info**
>
> Support dropped starting 5.5.

The ssh application allows a shell access via the associated interface on the configured port options. The application may be enabled on all interface types.

Parameter Name	Description
Port	Port allowing ssh access.

5.18.2 SNMP

> **Info**
>
> Support dropped starting 5.5. Please use Prometheus via Unified SBC management service instead.

The snmp application enables SNMP daemon listening. Note that this application only has effect on SBC node.

It does not require TLS profile, as TLS is not used.

The application may be enabled on `CX` interface.

Parameter Name	Description
Port	Port on which the SNMP server listens.

5.18.3 TURN server for websocket

> **Info**
>
> Application available since SBC release 4.5.

> **Info**
>
> Support dropped starting 5.2.

It enables the TURN server on given node. It is possible to configure one TURN server per node but it can be configured for more than one node.

It does not require a TLS profile.

The application may be enabled on `CX` interface.

Note well: using the TURN server application might expose the SBC to certain security risks. Indeed, the TURN server application makes use of static credentials for compatibility purposes, such that these well known credentials might be misused. It is therefor important to limit the use of the TURN server

application to the use case where it is absolutely required (support TCP media transport). Enabling this application is absolutely not necessary to supporting WebRTC in general.

Parameter Name	Description
Listening port	Listening port of the TURN server.
Aux server	Auxiliary server address in the format `IP:port`.
Relay IP	Note: mandatory.
External IP	TURN Server public/private address mapping, if the server is behind NAT. In that situation, the External IP will be reported as relay IP address of all allocations. This scenario works only in a simple case when one single relay address is be used, and no RFC5780 functionality is required. That single relay address must be mapped by NAT to the 'external' IP. The External IP value, if not empty, is returned in XOR-RELAYED-ADDRESS field. For that 'external' IP, NAT must forward ports directly (relayed port 12345 must be always mapped to the same 'external' port 12345).
UDP port range min port	Sets the UDP range that is used for relaying media start port. Note: mandatory.
UDP port range max port	Sets the UDP range that is used for relaying media end port. Note: mandatory.
Auth user	Sets the username used for TURN server authentication. Note: mandatory.
Auth password	Sets the password used for TURN server authentication. Note: mandatory.
Realm for users	Realm passed, which is usually domain name.
Media IP to allow UDP on firewall	Sets the IP address that will be allowed on SBC firewall to talk to the TURN.
UDP port range min port for media IP	Sets the UDP range that is used for media IP, start port.
UDP port range max port for media IP	Sets the UDP range that is used for media IP, end port.

5.18.4 Local monitoring query service

> **Info**
>
> Application available since SBC release 4.2.

> **Info**
>
> Replaced by Unified SBC management service starting 5.5.

The `sbc-xmloredis` API serves some metrics issued from different sources. The API will by default listen on the `localhost` interface, reachable via `http`. For every other interface application enabled, the API will listen exclusively via `https`, serving the configured TLS profile, which is required.

A non expose list of the available endpoints information and actions are:

- read and delete various data about registration cache
- read and delete various data about live calls
- fetch various data about monitored destination
- fetch various data about process's statistics
- support for Let's Encrypt HTTP01 challenge
- read and write various data about blacklists' call agent
- read and write various data about blacklists' destinations

- fetch various data about registration agent

The application only exists on SBC node. It is also exclusive and mandatory to `IMI` interface.

Parameter Name	Description
Port	Port on which the API server listens. Note: value not editable (`4242`).

5.18.5 PCAP query service

> **Info**
>
> Application available since SBC release 4.5.

> **Info**
>
> Replaced by Unified SBC management service starting 5.5.

The `sbc-pkapman` API generates and serves pcap files based on an aggregation of the pcap files available on the file system. The API will by default listen on the `localhost` interface, reachable via `http`. For every other interface application enabled, the API will listen exclusively via `https`, serving the configured TLS profile, which is required.

Requirements: SEMS's global option "Dump TLS session keys to file" Signaling SSL must be enabled if one wishes to download both pcap files and session TLS keys into a zip'ed bundle. Otherwise, the bundle may only contain pcap files.

Limitations: WebRTC interface don't support dump of the TLS keys.

A non expose list of the available endpoints information and actions are:
- fetch SBC node' file system PCAP files' timestamps
- merge SBC node' file system PCAP files as one
- merge SBC node' file system PCAP files as a ZIP with TLS keys

The application only exists on SBC node and it is mandatory and exclusive to `IMI` interface.

Parameter Name	Description
Port	Port on which the API server listens. Note: value not editable (`4243`).

5.18.6 Local webconf API

> **Info**
>
> Application available since SBC release 4.6.

> **Info**
>
> Replaced by Unified SBC management service starting 5.5.

The `sbc-webconf` API expose information and actions related to SEMS's web-conferencing features.

A non expose list of the available endpoints information and actions are:
- create and read rooms
- kick / mute / unmute room's active participants
- create room's dialouts
- create and edit room's pin
- read server information

The API will by default listen on the `localhost` interface, reachable via `http`. For every other interface application enabled, the API will listen exclusively via `https`, serving the configured TLS profile, which is required.

The application is exclusive and mandatory to `IMI` interface.

Parameter Name	Description
Port	Port on which the API server listens. Note: value not editable (`4244`).

5.18.7 Management for host

> **Info**
>
> Application available since SBC release 4.5.

> **Info**
>
> Replaced by Unified SBC management service starting 5.5.

The `sbc-gominstrator` API run various administrator tasks from RESTful endpoints.

A non exhaustive list of the available endpoints actions are:
- toggle the node' HA mode
- toggle the node' maintenance node
- restart/halt/shutdown the node

The API will by default listen on the `localhost` interface, reachable via `http`. For every other interface application enabled, the API will listen exclusively via `https`, serving the configured TLS profile, which is required.

The application may be enabled on `IMI` interface.

Parameter Name	Description
Port	Port on which the API server listens. Note: value not editable (`4249`).

5.18.8 Log files provider

> **Info**
>
> Application available since SBC release 5.1.

Info

Replaced by Unified SBC management service starting 5.5.

The `sbc-goplog` API allow HTTP interactions with the SBC node' file system log files.

A non exhaustive list of the available endpoints actions are:
- list the SBC node' file system log files
- read the SBC node' file system log files by streaming the lnav application trough websockets

The API will by default listen on the `localhost` interface, reachable via `http`. For every other interface application enabled, the API will listen exclusively via `https`, serving the configured TLS profile, which is required.

The application is exclusive to `IMI` interface.

Parameter Name	Description
Port	Port on which the API server listening. Note: value not editable (`4250`).

Application is available since ABC SBC' release 5.1.

5.18.9 Local packet classifier

Info

Application available since SBC release 5.4.

Info

Replaced by Unified SBC management service starting 5.5.

The `sbc-gopacla` API allow to partially interact with the ABC SBC firewall. Currently, the API allow to list nftable sets entries, add entry to nftable set or prune the entry / sets.

A non exhaustive list of the available endpoints actions are:
- feed and prune the node firewall' sets and their entries
- test if entries exist in the node firewall' sets

The API will by default listen on the `localhost` interface, reachable via `http`. For every other interface application enabled, the API will listen exclusively via `https`, serving the configured TLS profile, which is required.

The application is exclusive to `IMI` interface.

Parameter Name	Description
Port	Port on which the API server listens. Note: value not editable (`4252`).

5.18.10 HTTP proxy

> **Info**
>
> Application available since SBC release 4.6.

> **Info**
>
> Support dropped starting 5.5.

Setup an HTTP proxy, based on nginx reverse proxy: ProxyDoc.

The application adds the `X-Real-IP`, `Upgrade` and `Connection` headers. The template (`/etc/frafos/templates/nginx/proxy.tmpl`) may be overloaded, as described in Command Line Reference.

If "TLS enable" is set, a TLS profile is required.

The application may be enabled on `CX` interface.

Parameter Name	Description
Source Port	Port from which the proxy should operate.
Source Path	Path from which the proxy should operate.
Target IP address	IP to which the proxy redirect. Note: mandatory.
Target port	Port to which the proxy redirect. Note: mandatory.
TLS enable	Proxy over TLS.

5.18.11 HTTP redirect

> **Info**
>
> Application available since SBC release 4.6.

> **Info**
>
> Support dropped starting 5.5.

Setup an HTTP redirect pattern, using nginx rewrite directive: RewriteDoc.

The template (`/etc/frafos/templates/nginx/http_redirect.tmpl`) may be overloaded, as described in Command Line Reference.

If "TLS enable" is set, a TLS profile is required.

The application may be enabled on `CX` interface.

Parameter Name	Description
Port	Port from which the redirect should operate.
Path	Path from which the redirect should operate. Path is a regex to which we prefix ^ (start of line).

Target URL	URL to where be redirected. Note: mandatory.
TLS enable	Redirect over TLS.

5.18.12 Prometheus Pull Service

> **Info**
>
> Application available since SBC release 5.2.

> **Info**
>
> Starting 5.5 Prometheus statistics are available via Unified SBC management service application and Prometheus Pull Service application is deprecated.

Enable the prometheus pull service application on SBC node, allowing external prometheus scrapers to query the pull service to get statistics on the SBC.

The application may only be enabled on `CX` interfaces.

Parameter Name	Description
Port	The http(s) port of prometheus pull service.
Path	The url path part which to serve the statistics on.
TLS enabled	Use TLS on the pull service. Plain http will not be allowed. This can follow the configuration of the TLS profile (i.e. auth with trusted clients).
HTTP Auth. Username	Whether or not to use HTTP basic authentication on the pull service.
HTTP Auth. Password	Whether or not to use HTTP basic authentication on the pull service.
Threads	Number of threads to use while serving the requests.
Update interval	Interval in milliseconds to update the served statistics.

Chapter 6

Glossary

3GPP	3rd Generation Partnership Project
AoR	Address of Record
B2BUA	Back to Back User Agent
BLF	Busy Lamp Field
CA	Call Agent
CDR	Call Data Record / Call Detail Record
CPU	Central Processing Unit
DNS	Domain Name System
DoS	Denial of Service
ENUM	Electronic Number Mapping System
FQDN	Fully Qualified Domain Name
HA	High availability
IMI	Internal Management Interface
IMS	IP Multimedia Subsystem
IP	Internet Protocol
ISUP	ISDN User Part
MI	Media Interface
NAT	Network Address Translator
NIC	Network Interface Card
NNI	Network-Network Interface
PBX	Private Exchange
PSTN	Public Switched Telecommunication Network
REST	Representational state transfer
RLM	Realm
RTP	Real-Time Transport Protocol
RTCP	Real-Time Transport Control Protocol
SAP	Session Announcement Protocol
SBC	Session Border Controller
SCTP	Stream Control Transport Protocol
SDP	Session Description Protocol
SI	Signaling interface
SIP	Session Initiation Protocol
SNMP	Simple Network Management Protocol
SST	SIP Session Timers
SRV	Service Record

STUN	Session Traversal Utilities for NAT
TISPAN	Telecommunications and Internet converged Services and Protocols for Advanced Networking
TCP	Transport Control Protocol
TLS	Transport Level Security
UAC	User Agent Client
UAS	User Agent Server
UDP	User Datagram Protocol
UNI	User-Network Interface
URI	Universal Resource Indicator
VIP	Virtual IP Address
VoIP	Voice over IP
XMI	External Management Interface

Additionally, the ABC SBC documentation is available in PDF format either as a single PDF or separate Installation guide, User guide, Reference book and Release notes.

Internal ABC SBC API description and CCM API description can be found in API reference.