iptel.org SIP Express Router v0.8.10 -- Admin's Guide

Jiri Kuthan

Jan Janak

Yacine Rebahi

iptel.org SIP Express Router v0.8.10 -- Admin's Guide
by Jiri Kuthan, Jan Janak, and Yacine Rebahi

Copyright © 2001, 2002 by FhG Fokus

The document describes the SIP Express Router and its use in SIP networks. It is intended as an aid to server administrators.

This documentation is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.
This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.
You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
For more details see the file COPYING in the source distribution of SER.

Table of Contents

1. General Information 1
- About SIP Express Router (SER) 1
- About iptel.org 1
- Feature List 1
- Use Cases 2
  - Added-Value ISP Services 2
  - PC2Phone 3
  - PBX Replacement 3
- About SIP Technology 3
- Known SER Limitations 4
- Licensing 4
- Obtaining Technical Assistance 9
- More Information 10

2. Introduction to SER 11
- Request Routing and SER Scripts 11
- Conditional Statements 11
  - Operators and Operands 12
  - URI Matching 14
- Request URI Rewriting 16
- Destination Set 18
- External Modules 19
- Writing Scripts 21
  - Default Configuration Script 21
  - Stateful User Agent 24
  - Redirect Server 25
  - Executing External Script 26
  - Reply Processing (Forward on Unavailable) 27

3. Server Operation 29
- Recommended Operational Practices 29
- HOWTOs 32
- Troubleshooting 41

4. Application Writing 43
- Application FIFO Server 43

5. Reference 47
- Core Options 47
- Core Commands 48
- Command Line Parameters 51
- `serctl` command 52
- Modules 53
- FIFO Commands Reference 56
- Used Database Tables 57

Chapter 1. General Information

About SIP Express Router (SER)

SIP Express Router (SER) is an industrial-strength, free VoIP server based on the Session Initiation Protocol (SIP, RFC3261). It is engineered to power IP telephony infrastructures up to large scale. The server keeps track of users, sets up VoIP sessions, relays instant messages and creates space for new plug-in applications. Its proven interoperability guarantees seamless integration with components from other vendors, eliminating the risk of a single-vendor trap. It has successfully participated in various interoperability tests in which it worked with the products of other leading SIP vendors.

The SIP Express Router enables a flexible plug-in model for new applications: Third parties can easily link their plug-ins with the server code and provide thereby advanced and customized services. In this way, plug-ins such as SNMP support, RADIUS accounting, or SMS gateway have already been developed and are provided as advanced features. Other modules are underway: Presence server, firewall control, and more.

Its performance and robustness allows it to serve millions of users and accommodate needs of very large operators. With a $3000 dual-CPU PC, the SIP Express Router is able to power IP telephony services in an area as large as the Bay Area during peak hours. Even on an IPAQ PDA, the server withstands 150 calls per second (CPS)! The server has been powering our iptel.org free SIP site withstanding heavy daily load that is further increasing with the popularity of Microsoft's Windows Messenger.

The SIP Express Router is extremely configurable to allow the creation of various routing and admission policies as well as setting up new and customized services. Its configurability allows it to serve many roles: network security barrier, application server, or PSTN gateway guard for example.

`ser` can be also used with contributed applications. Currently, `serweb`, a `ser` web interface, and `SIPSak` diagnostic tool are available. Visit our site, http://www.iptel.org/, for more information on contributed packages.

About iptel.org

iptel.org is a know-how company spun off from Germany's national research company FhG Fokus. One of the first SIP implementations ever, low-QoS enhancements, interoperability tests and VoIP-capable firewall control concepts are examples of well-known FhG's work.

iptel.org continues to keep this know-how leadership in SIP. The access rate of the company's site, a well-known source of technological information, is a best proof of interest. Thousands of hits come every day from the whole Internet.

The iptel.org site, powered by SER, offers SIP services on the public Internet. Feel free to apply for a free SIP account at http://www.iptel.org/user/ [2]

Feature List

Based on the latest standards, the SIP Express Router (SER) includes support for registrar, proxy and redirect mode. Further it acts as an application server with support for instant messaging and presence including a 2G/SMS and Jabber gateway, a call control policy language, call number translation, private dial plans and accounting, authorization and authentication (AAA) services. `SER` runs on Sun/Solaris, PC/Linux, PC/BSD, IPAQ/Linux platforms and supports both IPv4 and IPv6. Hosting multiple domains and database redundancy is supported.

Other extensions are underway: presence server, firewall control and more.

ser has been carefully engineered with the following design objectives in mind:

- *Speed* - With ser, thousands of calls per seconds are achievable even on low-cost platforms. This competitive capacity allows setting up networks which are inexpensive and easy to manage due to low number of devices required. The processing capacity makes dealing with many stress factors easier. The stress factors may include but are not limited to broken configurations and implementations, boot avalanches on power-up, high-traffic applications such as presence, redundancy replications and denial-of-service attacks.

  The speed has been achieved by extensive code optimization, use of customized code, ANSI C combined with assembly instructions and leveraging latest SIP improvements. When powered by a dual-CPU Linux PC, ser is able to process thousands of calls per second, capacity needed to serve call signaling demands of Bay Area population.

- *Flexibility* - SER allows its users to define its behavior. Administrators may write textual scripts which determine SIP routing decisions, the main job of a proxy server. They may use the script to configure numerous parameters and introduce additional logic. For example, the scripts can determine for which destinations record routing should be performed, who will be authenticated, which transactions should be processed statefully, which requests will be proxied or redirected, etc.
- *Extensibility* - SER's extensibility allows linking of new C code to ser to redefine or extend its logic. The new code can be developed independently on SER core and linked to it in run-time. The concept is similar to the module concept known for example in Apache Web server. Even such essential parts such as transaction management have been developed as modules to keep the SER core compact and fast.
- *Portability.* ser has been written in ANSI C. It has been extensively tested on PC/Linux and Sun/Solaris. Ports to BSD and IPAQ/Linux exist.
- *Interoperability.* ser is based on the open SIP standard. It has undergone extensive tests with products of other vendors both in iptel.org labs and in the SIP Interoperability Tests (SIPIT). ser powers the public iptel.org site 24 hours a day, 356 days a year serving numerous SIP implementations using this site.
- *Small size.* Footprint of the core is 300k, add-on modules take up to 630k.

Use Cases

This section illustrates the most frequent uses of SIP. In all these scenarios, the SIP Express Router (SER) can be easily deployed as the glue connecting all SIP components together, be it soft-phones, hard-phones, PSTN gateways or any other SIP-compliant devices.

Added-Value ISP Services

To attract customers, ISPs frequently offer applications bundled with IP access. With SIP, the providers can conveniently offer a variety of services running on top of a single infrastructure. Particularly, deploying VoIP and instant messaging and presence services is as easy as setting up a SIP server and guiding customers to use Windows Messenger. Additionally, the ISPs may offer advanced services such as PSTN termi-

nation, user-driven call handling or unified messaging all using the same infrastructure.

SIP Express Router has been engineered to power large scale networks: its capacity can deal with large number of customers under high load caused by modern applications. Premium performance allows deploying a low number of boxes while keeping investments and operational expenses extremely low. ISPs can offer SIP-based instant messaging services and interface them to other instant messaging systems (Jabber, SMS). VoIP can be easily integrated along with added-value services, such as voicemail.

PC2Phone

Internet Telephony Service Providers (ITSPs) offer the service of interconnecting Internet telephony users using PC softphone or appliances to PSTN. Particularly with long-distance and international calls, competitive pricing can be achieved by routing the calls over the Internet.

SIP Express Router can be easily configured to serve pc2phone users, distribute calls to geographically appropriate PSTN gateway, act as a security barrier and keep track of charging.

PBX Replacement

Replacing a traditional PBX in an enterprise can achieve reasonable savings. Enterprises can deploy a single infrastructure for both voice and data and bridge distant locations over the Internet. Additionally, they can benefit of integration of voice and data.

The SIP Express Router scales from SOHOs to large, international enterprises. Even a single installation on a common PC is able to serve VoIP signaling of any world's enterprise. Its policy-based routing language makes implementation of numbering plans of companies spread across the world very easy. ACL features allow for protection of PSTN gateway from unauthorized callers.

SIP Express Router's support for programmable routing and accounting efficiently allows for implementation of such a scenario.

About SIP Technology

The SIP protocol family is the technology which integrates services. With SIP, Internet users can easily contact each other; figure out willingness to have a conversation and couple different applications such as VoIP, video and instant messaging. Integration with added-value services is seamless and easy. Examples include integration with web (click-to-dial), E-mail (voice2email, UMS), and PSTN-like services (conditional forwarding).

The core piece of the technology is the Session Initiation Protocol (SIP, RFC3261) standardized by IETF. Its main function is to establish communication sessions between users connected to the public Internet and identified by e-mail-like addresses. One of SIP's greatest features is its transparent support for multiple applications: the same infrastructure may be used for voice, video, gaming or instant messaging as well as any other communication application.

There are numerous scenarios in which SIP is already deployed: PBX replacement allows for deployment of single inexpensive infrastructure in enterprises; PC-2-phone long-distance services (e.g., Deltathree) cut callers long-distance expenses; instant messaging offered by public severs (e.g., iptel.org) combines voice and text services with presence information. New deployment scenarios are underway: SIP is a part

of UMTS networks and research publications suggest the use of SIP for virtual home environments or distributed network games.

Known SER Limitations

The following items are not part of current distribution and are planned for next releases:

- TCP transport
- Script processing of multiple branches on forking

> **Warning**
>
> `ser`'s request processing language allows to make request decisions based on current URI. When a request if forked to multiple destinations, only the first branch's URI is used as input for script processing. This might lead to unexpected results. Whenever a URI resolves to multiple different next-hop URIs, only the first is processed which may result in handling not appropriate for the other branch. For example, a URI might resolve to an IP phone SIP address and PSTN gateway SIP address. If the IP phone address is the first, then script execution ignores the second branch. If a script includes checking gateway address in request URI, the checks never match. That might result in ignoring of gateway admission control rules or applying them unnecessarily to non-gateway destinations.

List of known bugs is publicly available at http://developer.berlios.de/bugs/?group_id=480 [3].

Licensing

`ser` is freely available under terms and conditions of the GNU General Public License.

```
------------------------------------------------------------------------
----
IMPORTANT NOTES

1) The GPL applies to this copy of SIP Express Router software (ser).
   For a license to use the ser software under conditions
   other than those described here, or to purchase support for this
   software, please contact iptel.org by e-mail at the following addresses:

    info@iptel.org

   (see http://www.gnu.org/copyleft/gpl-faq.html#TOCHeardOtherLicense
    for an explanation how parallel licenses comply with GPL)

2) ser software allows programmers to plug-in external modules to the
   core part. Note that GPL mandates all plug-ins developed for the
   ser software released under GPL license to be GPL-ed as well.

   (see http://www.gnu.org/copyleft/gpl-faq.html#GPLAndPlugins
    for a detailed explanation)
```

3) Note that the GPL bellow is copyrighted by the Free Software Foundation,
   but the ser software is copyrighted by FhG

GNU GENERAL PUBLIC LICENSE
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest

your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the

original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Obtaining Technical Assistance

We offer best-effort free support for `ser`. "best-effort" means that we try to solve your problems via email as soon as we can, subject to available manpower. If you need commercial support, contact info@iptel.org.

To receive feedback to your inquiries, we recommend you to subscribe to the *serusers* mailing list and post your queries there. This mailing list is set up for mutual help by the community of `ser` users and developers.

Mailing List Instructions

- Public archives and subscription form: http://mail.iptel.org/mailman/listinfo/serusers [4]
- To post, send an email to serusers@iptel.org
- If you think you encountered an error, please submit the following information to avoid unnecessary round-trip times:
  - Name and version of your operating system -- you can obtain it by calling **uname -a**
  - `ser` distribution: release number and package
  - `ser` build -- you can obtain it by calling **ser -V**
  - Your `ser` configuration file
  - `ser` logs -- with default settings few logs are printed to **syslog** facility which typically dumps them to `/var/log/messages`. To enable detailed logs dumped to `stderr`, apply the following configuration options: **debug=8, log_stderror=yes, fork=no**.
  - Captured SIP messages -- you can obtain them using tools such as `ngrep` or `ethereal`.

If you are concerned about your privacy and do not wish your queries to be posted and archived publicly, you may post to serhelp@iptel.org. E-mails to this address are

only forwarded to iptel.org's `ser` development team. However, as the team is quite busy you should not be surprised to get replies with considerable delay.

More Information

Most up-to-date information including latest and most complete version of this documentation is always available at our website, http://www.iptel.org/ser/. For information on how to install ser, read INSTALL. SGML documentation is available in the 'doc' directory. A SIP tutorial (slide set) is available at http://www.iptel.org/sip/ .

Notes

1. http://www.iptel.org/
2. http://www.iptel.org/user/
3. http://developer.berlios.de/bugs/?group_id=480
4. http://mail.iptel.org/mailman/listinfo/serusers
5. http://www.iptel.org/ser/
6. http://www.iptel.org/sip/

Chapter 2. Introduction to SER

Request Routing and SER Scripts

The most important concept of every SIP server is that of request routing. The request routing logic determines the next hop of a request. It can be for example used to implement user location service or enforce static routing to a gateway. Real-world deployments actually ask for quite complex routing logic, which needs to reflex static routes to PSTN gateways, dynamic routes to registered users, authentication policy, capabilities of SIP devices, etc.

SER's answer to this need for routing flexibility is a routing language, which allows administrators to define the SIP request processing logic in a detailed manner. They can for example easily split SIP traffic by method or destination, perform user location, trigger authentication, verify access permissions, and so on.

The primary building block of the routing language are *actions*. There are built-in actions (like **forward** for stateless forwarding or **strip** for stripping URIs) as well as external actions imported from shared library modules. All actions can be combined in compound actions by enclosing them in braces, e.g. **{a1(); a2();}**. Actions are aggregated in one or more *route blocks*. Initially, only the default routing block denoted by **route[0]** is called. Other routing blocks can be called by the action **route(blocknumber)**, recursion is permitted. The language includes *conditional statements*.

The routing script is executed for every received request in sequential order. Actions may return positive/negative/zero value. Positive values are considered success and evaluated as TRUE in conditional expressions. Negative values are considered FALSE. Zero value means error and leaves execution of currently processed route block. The route block is left too, if **break** is explicitly called from it.

The easiest and still very useful way for `ser` users to affect request routing logic is to determine next hop statically. An example is routing to a PSTN gateway whose static IP address is well known. To configure static routing, simply use the action **forward(IP_address, port_number)**. This action forwards an incoming request "as is" to the destination described in action's parameters.

Example 2-1. Static Forwarding

```
# if requests URI is numerical and starts with
# zero, forward statelessly to a static destination

if (uri=~"^sip:0[0-9]*@iptel.org) {
    forward( 192.168.99.3, 5080 );
}
```

However, static forwarding is not sufficient in many cases. Users desire mobility and change their location frequently. Lowering costs for termination of calls in PSTN requires locating a least-cost gateway. Which next-hop is taken may depend on user's preferences. These and many other scenarios need the routing logic to be more dynamic. We describe in the Section called *Conditional Statements* how to make request processing subject to various conditions and in the Section called *Request URI Rewriting* how to determine next SIP hop.

Conditional Statements

A very useful feature is the ability to make routing logic depend on a condition. A script condition may for example distinguish between request processing for served and foreign domains, IP and PSTN routes, it may split traffic by method or username,

it may determine whether a request should be authenticated or not, etc. ser allows administrators to form conditions based on properties of processed request, such as method or uri, as well as on virtually any piece of data on the Internet.

Example 2-2. Conditional Statement

This example shows how a conditional statement is used to split incoming requests between a PSTN gateway and a user location server based on request URI.

```
# if request URI is numerical, forward the request to PSTN gateway...
if (uri=~"^sip:[0-9]+@foo.bar") { # match using a regular expression
    forward( gateway.foo.bar, 5060 );
} else { # ... forward the request to user location server otherwise
    forward( userloc.foo.bar, 5060 );
};
```

Conditional statements in ser scripts may depend on a variety of expressions. The simplest expressions are action calls. They return true if they completed successfully or false otherwise. An example of an action frequently used in conditional statements is **search** imported from textops module. **search** action leverages textual nature of SIP and compares SIP requests against a regular expression. The action returns true if the expression matched, false otherwise.

Example 2-3. Use of search Action in Conditional Expression

```
# prevent strangers from claiming to belong to our domain;
# if sender claims to be in our domain in From header field,
# better authenticate him
if (search("(f|From): .*@mydomain.com)) {
    if (!(proxy_authorize("mydomain.com" /* realm */,"subscriber" /* ta-
ble name */ ))) {
            proxy_challenge("mydomain.com /* ream */, "1" /* use qop */ );
            break;
    }
}
```

As modules may be created, which export new functions, there is virtually no limitation on what functionality ser conditions are based on. Implementers may introduce new actions whose return status depends on request content or any external data as well. Such actions can query SQL, web, local file systems or any other place which can provide information wanted for request processing.

Furthermore, many request properties may be examined using existing built-in operands and operators. Available left-hand-side operands and legal combination with operators and right-hand-side operands are described in Table 2-1. Expressions may be grouped together using logical operators: negation (**!**), AND (**&&**), OR (**||** and precedence parentheses (**()**).

Operators and Operands

There is a set of predefined operators and operands in ser, which in addition to actions may be evaluated in conditional expressions.

Left hand-side operands, which ser understands are the following:

- *method*, which refers to request method such as REGISTER or INVITE
- *uri*, which refers to current request URI, such as "sip:john.doe@foo.bar"

Note: Note that "uri" always refers to current value of URI, which is subject to change be uri-rewriting actions.

- *scr_ip*, which refers to IP address from which a request came.
- *dst_ip* refers to server's IP address at which a request was received

ser understands the following operators:

- == stands for equity
- =~ stands for regular expression matching
- logical operators: and, or, negation, parentheses (C-notation for the operators may be used too)

Table 2-1. Valid Combinations of Operands and Operators in Expressions

left-hand-side operand	valid operators	valid right-hand side operators	exam-ples/comments
method	== (exact match), =~ (regular expression matching)	string	method=="INVITE" || method=="ACK" || method=="CANCEL"
uri	== (exact match), =~ (regular expression matching)	string	uri=="sip:foo@bar.com" matches only if exactly this uri is in request URI
	== (exact match)	myself	the expression uri==myself is true if the host part in request URI equals a server name or a server alias (set using the alias option in configuration file)
src_ip	== (match)	IP, IP/mask_length, IP/mask, hostname, myself	src_ip==192.168.0.0/16 matches requests coming from a private network
dst_ip by `ser`	== (match)	IP, IP/mask_length, IP/mask, hostname, myself	dst_ip==127.0.0.1 matches if a request was received via loopback interface

Example 2-4. More examples of use of `ser` operators and operands in conditional

statements

```
# using an action as condition input; in this
# case, an actions 'search' looks for Contacts
# with private IP address in requests; the condition
# is processed if such a contact header field is
# found

if (search("^(Contact|m): .*@(192\.168\.|10\.|172\.16)")) {
# ....

# this condition is true if request URI matches
# the regular expression "@bat\.iptel\.org"
    if (uri=~"@bat\.iptel\.org") {
# ...

# and this condition is true if a request came
# from an IP address (useful for example for
# authentication by IP address if digest is not
# supported) AND the request method is INVITE

# if ( (src_ip==192.68.77.110 and method=="INVITE")
# ...
```

URI Matching

URI matching expressions have a broad use in a SIP server and deserve more explanation. Typical uses of URI matching include implementation of numbering plans, domain matching, binding external applications to specific URIs, etc. This section shows examples of typical applications of URI-matching.

Domain Matching

One of most important uses of URI matching is deciding whether a request is targeted to a served or outside domain. Typically, different request processing applies. Requests for outside domains are simply forwarded to them, whereas more complex logic applies to requests for a served domain. The logic may include saving user's contacts when REGISTER requests are received, forwarding requests to current user's location or a PSTN gateways, interaction with external applications, etc.

The easiest way to decide whether a request belongs a served domain is using the **myself** operand. The expression "uri==myself" returns true if domain name in request URI matches name of the host at which `ser` is running. This may be insufficient in cases when server name is not equal to domain name for which the server is responsible. For example, the "uri==myself" condition does not match if a server "sipserver.foo.bar" receives a request for "sip:john.doe@foo.bar". To match other names in URI than server's own, set up the `alias` configuration option. The option may be used multiple times, each its use adds a new item to a list of aliases. The myself condition returns then true also for any hostname on the list of aliases.

Example 2-5. Use of uri==myself Expression

```
# ser powers a domain "foo.bar" and runs at host sipserver.foo.bar;
# Names of served domains need to be stated in the aliases
# option; myself would not match them otherwise and would only
# match requests with "sipserver.foo.bar" in request-URI
alias="foo.bar"
alias="sales.foo.bar"
route[0] {
        if (uri==myself) {
            # the request either has server name or some of the
```

```
            # aliases in its URI
            log(1,"request for served domain")
            # some domain-specific logic follows here ....
        } else {
            # aha -- the server is not responsible for this
            # requests; that happens for example with the following URIs
            #  - sip:a@marketing.foo.bar
            #  - sip:a@otherdomain.bar
            log(1,"request for outbound domain");
            # outbound forwarding
            t_relay();
        };
}
```

It is possible to recognize whether a request belongs to a domain using regular expressions too. Care needs to be paid to construction of regular expressions. URI syntax is rich and an incorrect expression would result in incorrect call processing. The following example shows how an expression for domain matching can be formed.

Example 2-6. Domain Matching Using Regular Expressions

In this example, server named "sip.foo.bar" with IP address 192.168.0.10 is responsible for the "foo.bar" domain. That means, requests with the following hostnames in URI should be matched:

- foo.bar, which is the name of server domain
- sip.foo.bar, since it is server's name and some devices put server's name in request URI
- 192.168.0.10, since it is server's IP address and some devices put server's IP address in request URI

Note how this regular expression is constructed. In particular:

- User name is optional (it is for example never included in REGISTER requests) and there are no restrictions on what characters it contains. That is what *(.+@)?* mandates.
- Hostname must be followed by port number, parameters or headers -- that is what the delimiters *[:;\?]* are good for. If none it these follows, the URI must be ended (*$*). Otherwise, longer hostnames such as 192.168.0.101 or foo.bar.otherdomain.com would mistakenly match.
- Matches are case-insensitive. All hostnames "foo.bar", "FOO.BAR" and "FoO.bAr" match.

```
if (uri=~"^sip:(.+@)?(192\.168\.0\.10|(sip\.)?foo\.bar)([:;\?].*)?$")
      log(1, "yes, it is a request for our domain");
      break;
 };
```

Numbering Plans

Other use of URI matching is implementation of dialing plans. A typical task when designing a dialing plan for SIP networks is to distinguish between "pure-IP" and PSTN destinations. IP users typically have either alphanumerical or numerical usernames. The numerical usernames are convenient for PSTN callers who can only use numeric keypads. Next-hop destination of IP users is looked up dynamically using user location database. On the other hand, PSTN destinations are always indicated by nummerical usernames. Requests to PSTN are statically forwarded to well-known PSTN gateways.

Example 2-7. A simple Numbering Plan

This example shows a simple dialing plan which reserves dialing prefix "8" for IP users, other numbers are used for PSTN destinations and all other non-nummerical usernames are used for IP users.

```
# is it a PSTN destination? (is username nummerical and does not begin with 8?)
if (uri=~"^sip:[0-79][0-9]*@") { # ... forward to gateways then;
      # check first to which PSTN destination the requests goes;
      # if it is US (prefix "1"), use the gateway 192.168.0.1...
      if (uri=~"^sip:1") {
           # strip the leading "1"
           strip(1);
           forward(192.168.0.1, 5060);
      } else {
           # ... use the gateway 10.0.0.1 for all other destinations
           forward(10.0.0.1, 5060);
      }
      break;
} else {
      # it is an IP destination -- try to lookup it up in user location DB
      if (!lookup("location")) {
          # bad luck ... user off-line
          sl_send_reply("404", "Not Found");
          break;
      }
      # user on-line...forward to his current destination
      forward(uri:host,uri:port);
}
```

Request URI Rewriting

The ability to give users and services a unique name using URI is a powerful tool. It allows users to advertise how to reach them, to state to whom they wish to communicate and what services they wish to use. Thus, the ability to change URIs is very important and is used for implementation of many services. "Unconditional forwarding" from user "boss" to user "secretary" is a typical example of application relying on change of URI address.

`ser` has the ability to change request URI in many ways. A script can use any of the following built-in actions to change request URI or a part of it: **rewriteuri**, **rewritehost**, **rewritehostport**, **rewriteuser**, **rewriteuserpass** and **rewriteport**. When later in the script a forwarding action is encountered, the action forwards the request to address in the rewritten URI.

Example 2-8. Rewriting URIs

```
if (uri=~"dan@foo.bar") {
    rewriteuri("sip:bla@somewherelse.com")
    # forward statelessly to the destination in current URI, i.e.,
    # to sip:bla@somewherelese.com:5060
    forward( uri:host, uri:port);
}
```

Two more built-in URI-rewriting commands are of special importance for implementation of dialing plans and manipulation of dialing prefixes. **prefix(s)** , inserts a string "s" in front of SIP address and **strip(n)** takes away the first "n" characters of a SIP address. See Table 2-2 for examples of use of built-in URI-rewriting actions.

Commands exported by external modules can change URI too and many do so. The most important application is changing URI using the user location database. The command **lookup(table)** looks up current user's location and rewrites user's address with it. If there is no registered contact, the command returns a negative value.

Example 2-9. Rewriting URIs Using User Location Database

```
# store user location if a REGISTER appears
if (method=="REGISTER") {
   save("mydomain1");
} else {
# try to use the previously registered contacts to
# determine next hop
   if(lookup("mydomain1")) {
     # if found, forward there...
     t_relay();
   } else {
     # ... if no contact on-line, tell it upstream
     sl_send_reply("404", "Not Found" );
   };
};
```

External applications can be used to rewrite URI too. The "exec" module provides script actions, which start external programs and read new URI value from their output. **exec_uri** and **exec_user** both call an external program, pass current URI or its user part to it respectively, wait until it completes, and eventually rewrite current URI with its output.

It is important to realize that `ser` operates over *current URI* all the time. If an original URI is rewritten by a new one, the original will will be forgotten and the new one will be used in any further processing. In particular, the uri matching operand and the user location action **lookup** always take current URI as input, regardless what the original URI was.

Table 2-2 shows how URI-rewriting actions affect an example URI, sip:12345@foo.bar:6060.

Table 2-2. URI-rewriting Using Built-In Actions

Example Action	Resulting URI
rewritehost("192.168.0.10") rewrites the hostname in URI, other parts (including port number) remain unaffected.	sip:12345@192.168.10:6060

Example Action	Resulting URI
rewriteuri("sip:alice@foo.bar"); rewrites the whole URI completely.	sip:alice@foo.bar
rewritehost-port("192.168.0.10:3040")rewrites both hostname and port number in URI.	sip:12345@192.168.0.10:3040
rewriteuser("alice") rewrites user part of URI.	sip:alice@foo.bar:6060
rewriteuserpass("alice:pw") replaces the pair user:password in URI with a new value.	sip:alice:pw@foo.bar:6060
rewriteport("1234") replaces port number in URI	sip:12345@foo.bar:1234
prefix("9") inserts a string ahead of user part of URI	sip:912345@foo.bar:6060
strip(2) removes leading characters from user part of URI	sip:345@foo.bar:6060

You can verify whether you understood URI processing by looking at the following example. It rewrites URI several times. The question is what is the final URI to which the script fill forward any incoming request.

Example 2-10. URI-rewriting Quiz

```
exec_uri("echo sip:2234@foo.bar; echo > /dev/null");
strip(2);
if (uri=~"^sip:2") {
    prefix("0");
} else {
    prefix("1");
};
forward(uri:host, uri:port);
```

The correct answer is the resulting URI will be "sip:134@foo.bar". **exec_uri** rewrites original URI to "sip:2234@foo.bar", **strip(2)** takes two leading characters from username away resulting in "34@iptel.org", the condition does not match because URI does not begin with "2" any more, so the prefix "1" is inserted.

Destination Set

Whereas needs of many scenarios can by accommodated by maintaining a single request URI, some scenarios are better served by multiple URIs. Consider for example a user with address john.doe@iptel.org. The user wishes to be reachable at his home phone, office phone, cell phone, softphone, etc. However, he still wishes to maintain a single public address on his business card.

To enable such scenarios, `ser` allows translation of a single request URI into multiple outgoing URIs. The ability to forward a request to multiple destinations is known as *forking* in SIP language. All outogoing URIs (in trivial case one of them) are called *destination set*. The destination set always includes one default URI, to which additional URIs can be appended. Maximum size of a destination set is limited by a compile-time constant, MAX_BRANCHES, in `config.h`.

Some actions are designed for use with a single URI whereas other actions work with the whole destination set.

Actions which are currently available for creating the destination set are **lookup** from usrloc module and **exec_uri/exec_user** from exec module. **lookup** fills in the destination set with user contact's registered previously with REGISTER requests. The **exec** actions fill in the destination set with output of an external program. In both cases, current destination set is completely rewritten. New URIs can be appended to destination set by a call to the built-in action **append_branch(uri)**.

Currently supported features which utilize destination sets are *forking* and *redirection*. Action **t_relay** (TM module) for stateful forwarding supports forking. If called with a non-trivial destination set, **t_relay** forks incoming request to all URIs in current destination set. See Example 2-9. If a user previously registered from three locations, the destination set is filled with all of them by **lookup** and the **t_relay** command forwards the incoming request to all these destinations. Eventually, all user's phone will be ringing in parallel.

SIP redirection is another feature which leverages destination sets. It is a very lightweighted method to establish communication between two parties with minimum burden put on the server. In `ser`, the action **sl_send_reply** (SL module) is used for this purpose. This action allows to generate replies to SIP requests without keeping any state. If the status code passed to the action is 3xx, the current destination set is printed in reply's Contact header fields. Such a reply instructs the originating client to retry at these addresses. (See Example 2-16).

Most other `ser` actions ignore destination sets: they either do not relate to URI processing (**log**, for example) or they work only with the default URI. All URI-rewriting functions such as **rewriteuri** belong in this category. URI-comparison operands only refers to the first URI (see the Section called *Operators and Operands*). Also, the built-in action for stateless forwarding, **forward** works only with the default URI and ignores rest of the destination set. The reason is a proxy server willing to fork must guarantee that the burden of processing multiple replies is not put unexpectedly on upstream client. This is only achievable with stateful processing. Forking cannot be used along with stateless **forward**, which thus only processes one URI out of the whole destination set. Also, the uri comparison operand (see the Section called *Operators and Operands*) refers only to current URI and ignores the rest of destination set.

External Modules

`ser` provides the ability to link the server with external third-party shared libraries. Lot of functionality which is included in the `ser` distribution is actually located in modules to keep the server "core" compact and clean. Among others, there are modules for checking max_forwards value in SIP requests (maxfwd), transactional processing (tm), record routing (rr), accounting (acc), authentication (auth), SMS gateway (sms), replying requests (sl), user location (usrloc, registrar) and more.

In order to utilize new actions exported by a module, ser must first load it. To load a module, the directive **loadmodule "filename"** must be included in beginning of a `ser` script file.

Example 2-11. Using Modules

This example shows how a script instructs `ser` to load a module and use actions exported by it. Particularly, the sl module exports an action **sl_send_reply** which makes `ser` act as a stateless user agent and reply all incoming requests with 404.

```
# first of all, load the module!
loadmodule "/usr/lib/ser/modules/sl.so
route{
    # reply all requests with 404
    sl_send_reply("404", "I am so sorry -- user not found");
```

```
}
```

Note: Note that unlike with core commands, all actions exported by modules must have parameters enclosed in quotation marks in current version of `ser`. In the following example, the built-in action **forward** for stateless forwarding takes IP address and port numbers as parameters without quotation marks whereas a module action **t_relay** for stateful forwarding takes parameters enclosed in quotation marks.

Example 2-12. Parameters in built-in and exported actions

```
# built-in action doesn't enclose IP addresses and port numbers
# in quotation marks
forward(192.168.99.100, 5060);
# module-exported functions enclose all parameters in quotation
# marks
t_relay_to("192.168.99.100", "5060");
```

Many modules also allow users to change the way how they work using predefined parameters. For example, the authentication module needs to know location of MySQL database which contains users' security credentials. How module parameters are set using the **modparam** directive is shown in Example 2-13. **modparam** always contains identification of module, parameter name and parameter value. Description of parameters available in modules is available in module documentation.

Yet another thing to notice in this example is module dependency. Modules may depend on each other. For example, the authentication modules leverages the mysql module for accessing mysql databases and sl module for generating authentication challenges. We recommend that modules are loaded in dependency order to avoid ambiguous server behaviour.

Example 2-13. Module Parameters

```
# ------------------ module loading ----------------------------------
-

# load first modules on which 'auth' module depends;
# sl is used for sending challenges, mysql for storage
# of user credentials
loadmodule "modules/sl/sl.so"
loadmodule "modules/mysql/mysql.so"
loadmodule "modules/auth/auth.so"

# ------------------ module parameters -------------------------------
-
# tell the auth module the access data for SQL database:
# username, password, hostname and database name
modparam("auth", "db_url","sql://ser:secret@dbhost/ser")


# --------------------------  request routing logic ------------------
-

# authenticate all requests prior to forwarding them

route{

        if (!proxy_authorize("foo.bar" /* realm */,
                        "subscriber" /* table name */ )) {
```

```
                proxy_challenge("foo.bar", "0");
                break;
        };
        forward(192.168.0.10,5060);
}
```

Writing Scripts

This section demonstrates simple examples how to configure server's behaviour using the `ser` request routing language. All scripts follow the `ser` language syntax, which dictates the following block ordering:

- *global configuration parameters* -- these value affect behaviour of the server such as port number at which it listens, number of spawned children processes, and log-level. See the Section called *Core Options* in Chapter 5 for a list of available options.
- *module loading* -- these statements link external modules, such as transaction management (tm) or stateless UA server (sl) dynamically. See the Section called *Modules* in Chapter 5 for a list of modules included in `ser` distribution.

  Note: If modules depend on each other, than the depending modules must be loaded after modules on which they depend. We recommend to load first modules **tm** and **sl** because many other modules (authentication, user location, accounting, etc.) depend on these.

- *module-specific parameters* -- determine how modules behave; for example, it is possible to configure database to be used by authentication module.
- one or more *route blocks* containing the request processing logic, which includes built-in actions as well as actions exported by modules. See the Section called *Core Commands* in Chapter 5 for a list of built-in actions.
- optionally, if modules supporting reply processing (currently only TM) are loaded, one or more *reply_route blocks* containing logic triggered by received replies. Restrictions on use of actions within **reply_route** blocks apply -- see the Section called *Core Commands* in Chapter 5 for more information.

For more complex examples, see the etc directory in `ser` source distribution. It contains the `iptel.cfg` script which is in production use at iptel.org's public SIP site and exploits most of `ser` features.

Default Configuration Script

The configuration script, `ser.cfg`, is a part of every `ser` distribution and defines default behaviour. It allows users to register with the server and have requests proxied to each other.

After performing routine checks, the script looks whether incoming request is for served domain. If so and the request is "REGISTER", `ser` acts as SIP registrar and updates database of user's contacts. Optionally, it verifies user's identity first to avoid unauthorized contact manipulation.

Non-REGISTER request for served domains are then processed using user location database. If a contact is found for requested URI, script execution proceeds to stateful forwarding, a negative 404 reply is generated otherwise. Requests outside served domain are always statefully forwarded.

Note that this simple script features several limitations:

- By default, authentication is turned off to avoid dependency on mysql. Unless it it turned on, anyone can register using any name and "steal" someone else's calls.
- Even it authentication is turned on, there is no relationship between authentication username and address of record. That means that for example a user authenticating himself correctly with "john.doe" id may register contacts for "gw.bush". Site policy may wish to mandate authentication id to be equal to username claimed in To header field. check_to action from auth module can be used to enforce such a policy.
- There is no dialing plan implemented. All users are supposed to be reachable via user location database. See the Section called *Numbering Plans* for more information.
- The script assumes users will be using server's name as a part of their address of record. If users wish to use another name (domain name for example), this must be set using the `alias` options. See the Section called *Domain Matching* for more information.
- If authentication is turned on by uncommenting related configuration options, clear-text user passwords will by assumed in back-end database.

Example 2-14. Default Configuration Script

```
#
# $Id: ser.cfg,v 1.12 2002/10/21 02:40:06 jiri Exp $
#
# simple quick-start config script
#

# ----------- global configuration parameters ------------------------
-

debug=3          # debug level (cmd line: -ddddddddddd)
fork=no
log_stderror=yes# (cmd line: -E)
check_via=no # (cmd. line: -v)
dns=no           # (cmd. line: -r)
rev_dns=no      # (cmd. line: -R)
port=5060
children=4
fifo="/tmp/ser_fifo"

listen=192.168.0.16

# ------------------ module loading ----------------------------------
-

# Uncomment this if you want to use SQL database
#loadmodule "/usr/lib/ser/modules/mysql.so"

loadmodule "modules/sl/sl.so"
loadmodule "modules/tm/tm.so"
loadmodule "modules/rr/rr.so"
loadmodule "modules/maxfwd/maxfwd.so"
loadmodule "modules/usrloc/usrloc.so"
loadmodule "modules/registrar/registrar.so"
```

```
# Uncomment this if you want digest authentication
# mysql.so must be loaded !
#loadmodule "/usr/lib/ser/modules/auth.so"

# ----------------- setting module-specific parameters --------------
-

# -- usrloc params --

modparam("usrloc", "db_mode",   0)

# Uncomment this if you want to use SQL database
# for persistent storage and comment the previous line
#modparam("usrloc", "db_mode", 2)

# -- auth params --
# Uncomment if you are using auth module
#
#modparam("auth", "secret", "alsdkhglaksdhfkloiwr")
#modparam("auth", "calculate_ha1", yes)
#
# If you set "calculate_ha1" parameter to yes (which true in this con-
fig),
# uncomment also the following parameter)
#
#modparam("auth", "password_column", "password")

# -------------------------  request routing logic ------------------
-

# main routing logic

route{

 # initial sanity checks -- messages with
 # max_forwars==0, or excessively long requests
 if (!mf_process_maxfwd_header("10")) {
  sl_send_reply("483","Too Many Hops");
  break;
 };
 if (len_gt( max_len )) {
  sl_send_reply("513", "Message too big");
  break;
 };

 # Do strict routing if pre-loaded route headers present
 rewriteFromRoute();

 # if the request is for other domain use UsrLoc
 # (in case, it does not work, use the following command
 # with proper names and addresses in it)
 if (uri==myself) {

  if (method=="REGISTER") {

# Uncomment this if you want to use digest authentication
#   if (!www_authorize("iptel.org", "subscriber")) {
#    www_challenge("iptel.org", "0");
#    break;
#   };

   save("location");
   break;
  };
```

```
  # native SIP destinations are handled using our USRLOC DB
  if (!lookup("location")) {
   sl_send_reply("404", "Not Found");
   break;
  };
 };
 # forward to current uri now
 if (!t_relay()) {
  sl_reply_error();
 };

}
```

Stateful User Agent

This examples shows how to make ser act as a stateful user agent (UA). Ability to act as as a stateful UA is essential to many applications which terminate a SIP path. These applications wish to focus on their added value. They do not wish to be involved in all SIP gory details, such as request and reply retransmission, reply formatting, etc. For example, we use the UA functionality to shield SMS gateway and instant message store from SIP transactional processing. The simple example bellow issues a log report on receipt of a new transaction. If we did not use a stateful UA, every single request retransmission would cause the application to be re-executed which would result in duplicated SMS messages, instant message in message store or log reports.

The most important actions are **t_newtran** and **t_reply**. **t_newtran** shields subsequent code from retransmissions. It returns success and continues when a new request arrived. It exits current route block immediately on receipt of a retransmissions. It only returns a negative value when a serious error, such as lack of memory, occurs.

t_reply generates a reply for a request. It generates the reply statefully, i.e., it is kept for future retransmissions in memory.

Example 2-15. Stateful UA

```
   #
# $Id: uas.cfg,v 1.5 2002/10/04 21:37:11 jiri Exp $
#
# this example shows usage of ser as user agent
# server which does some functionality (in this
# example, 'log' is used to print a notification
# on a new transaction) and behaves statefuly
# (e.g., it retransmits replies on request
# retransmissions)

# ------------------ module loading ----------------------------------
-

loadmodule "modules/sl/sl.so"
loadmodule "modules/tm/tm.so"


# -------------------------  request routing logic -------------------
-

# main routing logic
```

```
route{
 # for testing purposes, simply okay all REGISTERs
 if (method=="REGISTER") {
  log("REGISTER");
  sl_send_reply("200", "ok");
  break;
 };

 # create transaction state; abort if error occured
 if ( !t_newtran()) {
  sl_reply_error();
  break;
 };

 # the following log will be only printed on receipt of
 # a new message; retranmissions are absorbed by t_newtran
 log(1, "New Transaction Arrived\n");
        # do what you want to do as a sever...
 if (uri=~"a@") {
  if (!t_reply("409", "Bizzar Error")) {
   sl_reply_error();
  };
 } else {
  if (!t_reply("699", "I don't want to chat with you")) {
   sl_reply_error();
  };
     };
}
```

Redirect Server

The redirect example shows how to redirect a request to multiple destination using 3xx reply. Redirecting requests as opposed to proxying them is essential to various scalability scenarios. Once a message is redirected, `ser` discards all related state and is no more involved in subsequent SIP transactions (unless the redirection addresses point to the same server again).

The key `ser` actions in this example are **append_branch** and **sl_send_reply** (sl module).

append_branch adds a new item to the destination set. The destinations set always includes the current URI and may be enhanced up to `MAX_BRANCHES` items. **sl_send_reply** command, if passed SIP reply code 3xx, takes all values in current destination set and adds them to Contact header field in the reply being sent.

Example 2-16. Redirect Server

```
   #
# $Id: redirect.cfg,v 1.4 2002/10/04 21:37:11 jiri Exp $
#
# this example shows use of ser as stateless redirect server
#

# ------------------ module loading ----------------------------------
-

loadmodule "modules/sl/sl.so"
```

```
# -------------------------  request routing logic ------------------
-

# main routing logic

route{
 # for testing purposes, simply okay all REGISTERs
 if (method=="REGISTER") {
  log("REGISTER");
  sl_send_reply("200", "ok");
  break;
 };
 # rewrite current URI, which is always part of destination ser
 rewriteuri("sip:parallel@iptel.org:9");
 # append one more URI to the destination ser
 append_branch("sip:redirect@iptel.org:9");
 # redirect now
 sl_send_reply("300", "Redirect");
}
```

Executing External Script

Like in the previous example, we show how to make ser act as a redirect server. The difference is that we do not use redirection addresses hardwired in `ser` script but get them from external shell commands. We also use ser's ability to execute shell commands to log source IP address of incoming SIP requests.

The new commands introduced in this example are **exec_msg** and **exec_uri**. **exec_msg** takes current requests, starts an external command, and passes the requests to the command's standard input. It also passes request's source IP address in environment variable named `SRC_IP`.

exec_uri serves for URI rewriting by external applications. The **exec_uri** action passes current URI to the called external program as command-line parameter, and rewrites current destination set with the program's output. An example use would be an implementation of a Least-Cost-Router, software which returns URI of the cheapest PSTN provider for a given destination based on some pricing tables. Example 2-17 is much easier: it prints fixed URIs on its output using shell script **echo** command.

Example 2-17. Executing External Script

```
   #
# $Id: exec.cfg,v 1.4 2002/10/04 21:37:11 jiri Exp $
#
# this example shows use of ser as stateless redirect server
# which rewrites URIs using an exernal utility
#

# ------------------ module loading ----------------------------------
-

loadmodule "modules/exec/exec.so"
loadmodule "modules/sl/sl.so"

# -------------------------  request routing logic ------------------
-

# main routing logic
```

```
route{
 # for testing purposes, simply okay all REGISTERs
 if (method=="REGISTER") {
  log("REGISTER");
  sl_send_reply("200", "ok");
  break;
 };

 # first dump the message to a file using cat command
 exec_msg("printenv SRCIP > /tmp/exectest.txt; cat >> /tmp/exectest.txt");
 # and then rewrite URI using external utility
 # note that the last echo command trashes input parameter
 if (exec_uri("echo sip:mra@iptel.org;echo sip:mrb@iptel.org;echo>/dev/null")) {
  sl_send_reply("300", "Redirect");
 } else {
  sl_reply_error();
  log(1, "alas, rewriting failed\n");
 };
}
```

Reply Processing (Forward on Unavailable)

Many services depend on status of messages relayed downstream: *forward on busy* and *forward on no reply* to name the most well-known ones. To support implementation of such services, `ser` allows to return to request processing when request forwarding failed. When a request is reprocessed, new request branches may be initiated or the transaction can be completed at discretion of script writer.

The primitives used are **t_on_negative(r)** and **reply_route[r]{}.** If **t_on_negative** is called before a request is statefuly forwarded and a forwarding failure occurs, `ser` will return to request processing in a **reply_route** block. Failures include receipt of a SIP error (status code >= 300) from upstream or not receiving any final reply within final response period.

In Example 2-18, **reply_route[1]** is set to be entered on error using the **t_on_negative(1)** action. Within this reply block, `ser` is instructed to initiate a new branch and try to reach called party at another destination (sip:nonsense@iptel.org). To deal with the case when neither the alternate destination succeeds, `t_on_negative` is set again. If the case really occurs, **reply_route[2]** is entered and a last resort destination (sip:foo@iptel.org) is tried.

Example 2-18. Reply Processing

```
   #
# $Id: onr.cfg,v 1.5 2002/10/04 21:37:11 jiri Exp $
#
# example script showing both types of forking;
# incoming message is forked in parallel to
# 'nobody' and 'parallel', if no positive reply
# appears with final_response timer, nonsense
# is retried (serial forking); than, destination
# 'foo' is given last chance

# ------------------ module loading ----------------------------------
-

loadmodule "modules/sl/sl.so"
loadmodule "modules/tm/tm.so"
```

```
# ----------------- setting module-specific parameters --------------
-

# -- tm params --
# set time for which ser will be waiting for a final response;
# fr_inv_timer sets value for INVITE transactions, fr_timer
# for all others
modparam("tm", "fr_inv_timer", 15 )
modparam("tm", "fr_timer", 10 )

# -------------------------  request routing logic ------------------
-

# main routing logic

route{
 # for testing purposes, simply okay all REGISTERs
 if (method=="REGISTER") {
  log("REGISTER");
  sl_send_reply("200", "ok");
  break;
 };
 # try these two destinations first in parallel; the second
 # destination is targeted to sink port -- that will make ser
 # wait until timer hits
 seturi("sip:nobody@iptel.org");
 append_branch("sip:parallel@iptel.org:9");
 # if we do not get a positive reply, continue at reply_route[1]
 t_on_negative("1");
 # forward the request to all destinations in destination set now
 t_relay();
}

reply_route[1] {
 # forwarding failed -- try again at another destination
 append_branch("sip:nonsense@iptel.org");
 log(1,"first redirection\n");
 # if this alternative destination fails too, proceed to reply_route[2]
 t_on_negative("2");
}

reply_route[2] {
 # try out the last resort destination
 append_branch("sip:foo@iptel.org");
 log(1, "second redirection\n");
 # we no more call t_on_negative here; if this destination
 # fails too, transaction will complete
}
```

Chapter 3. Server Operation

Recommended Operational Practices

Operation of a SIP server is not always easy task. Server administrators are challenged by broken or misconfigured user agents, network and host failures, hostile attacks and other stress-makers. All such situations may lead to an operational failure. It is sometimes very difficult to figure out the root reason of a failure, particularly in a distributed environment with many SIP components involved. In this section, we share some of our practices and refer to tools which have proven to make life of administrators easier

1. Keeping track of messages is good

Frequently, operational errors are discovered or reported with a delay. Users frustrated by an error frequently approach administrators and scream "even though my SIP requests were absolutely ok yesterday, they were mistakenly denied by your server". If administrators do not record all SIP traffic at their site, they will be no more able to identify the problem reason. We thus recommend that site operators record all messages passing their site and keep them stored for some period of time. They may use utilities such as `ngrep` or `tcpdump` . There is also a utility `scripts/harv_ser.sh` in `ser` distribution for post-processing of captures messages. It summarizes messages captured by reply status and user-agent header field.

2. Real-time Traffic Watching

Looking at SIP messages in real-time may help to gain understanding of problems. Though there are commercial tools available, using a simple, text-oriented tool such as `ngrep` makes the job very well thanks to SIP's textual nature.

Example 3-1. Using `ngrep`

In this example, all messages at port 5060 which include the string "bkraegelin" are captured and displayed

```
[jiri@fox s]$ ngrep bkraegelin@ port 5060
interface: eth0 (195.37.77.96/255.255.255.240)
filter: ip and ( port 5060 )
match: bkraegelin@
#
U +0.000000 153.96.14.162:50240 -> 195.37.77.101:5060
REGISTER sip:iptel.org SIP/2.0.
Via: SIP/2.0/UDP 153.96.14.162:5060.
From: sip:bkraegelin@iptel.org.
To: sip:bkraegelin@iptel.org.
Call-ID: 0009b7aa-1249b554-6407d246-72d2450a@153.96.14.162.
Date: Thu, 26 Sep 2002 22:03:55 GMT.
CSeq: 101 REGISTER.
Expires: 10.
Content-Length: 0.
.

#
U +0.000406 195.37.77.101:5060 -> 153.96.14.162:5060
SIP/2.0 401 Unauthorized.
Via: SIP/2.0/UDP 153.96.14.162:5060.
From: sip:bkraegelin@iptel.org.
To: sip:bkraegelin@iptel.org.
Call-ID: 0009b7aa-1249b554-6407d246-72d2450a@153.96.14.162.
CSeq: 101 REGISTER.
WWW-Authenticate: Digest realm="iptel.org", nonce="3d9385170000000043acbf6ba9c9741790e0
Server: Sip EXpress router(0.8.8 (i386/linux)).
```

```
Content-Length: 0.
Warning: 392 127.0.0.1:5060 "Noisy feedback tells: pid=31604 req_src_ip=153.96.14.162 i
```

3. Tracing Errors in Server Chains

A request may pass any number of proxy servers on its path to its destination. If an error occurs, it may be quite difficult to learn in which of the servers in the chain it originated and what was its cause. `ser` does its best and displays extensive diagnostics information in SIP replies. This information is part of the warning header field, and contains the following facts:

- Server's IP Address -- good to identify from which server in a chain the reply came
- Incoming and outgoing URIs -- good to learn for which URI the reply was generated, as it may be rewritten many times in the path
- Number of Via header fields in replied request -- that helps in assessment of request path length.

A nice utility for debugging server chains is `sipsak`, Swiss Army Knife, traceroute-like tool for SIP developed at iptel.org. It allows you to send OPTIONS request with low, increasing Max-Forwards header-fields and follow how it propagates in SIP network. See its webpage at http://sipsak.berlios.de/ [1].

Example 3-2. Use of SIPSak for Learning SIP Path

```
[jiri@bat sipsak]$ ./sipsak -T -s sip:7271@iptel.org
warning: IP extract from warning activated to be more informational
0: 127.0.0.1 (0.456 ms) SIP/2.0 483 Too Many Hops
1: ?? (31.657 ms) SIP/2.0 200 OK
 without Contact header
```

Note that in this example, the second hop server does not issue any warning header fields in replies and it is thus impossible to display its IP address in `SIPsak`'s output.

4. Watching Server Health

Watching Server's operation status in real-time may also be a great aid for troubleshooting. `ser` has an excellent facility, a FIFO server, which allows UNIX tools to access server's internals. (It is similar to how Linux tool access Linux kernel via the proc file system.) The FIFO server accepts commands via a FIFO (named pipe) and returns data asked for. Administrators do not need to learn details of the FIFO communication and can serve themselves using a front-end utility `serctl`. Of particular interest for monitoring server's operation are `serctl` commands **ps** and **moni**. The former displays running `ser` processes, whereas the latter shows statistics.

Example 3-3. serctl ps command

This example shows 10 processes running at a host. The process 0, "attendant" watches child processes and terminates all of them if a failure occurs in any of them. Processes 1-4 listen at local interface and processes 5-8 listen at Ethernet interface at port number 5060. Process number 9 runs FIFO server, and process number 10 processes all server timeouts.

```
[jiri@fox jiri]$ serctl ps
```

```
0 31590 attendant
1 31592 receiver child=0 sock=0 @ 127.0.0.1::5060
2 31595 receiver child=1 sock=0 @ 127.0.0.1::5060
3 31596 receiver child=2 sock=0 @ 127.0.0.1::5060
4 31597 receiver child=3 sock=0 @ 127.0.0.1::5060
5 31604 receiver child=0 sock=1 @ 195.37.77.101::5060
6 31605 receiver child=1 sock=1 @ 195.37.77.101::5060
7 31606 receiver child=2 sock=1 @ 195.37.77.101::5060
8 31610 receiver child=3 sock=1 @ 195.37.77.101::5060
9 31611 fifo server
10 31627 timer
```

5. Is Server Alive

It is essential for solid operation to know continuously that server is alive. We've been using two tools for this purpose. `sipsak` does a great job of "pinging" a server, which may be used for alerting on unresponsive servers.

`monit` is a server watching utility which alerts when a server dies.

6. Setting Proper Log Level

If something is going wrong and you are in doubts what causes the error, increase log level. Additional log messages may help you to trace the error reason. Be careful though: `ser` is very talkative in higher debugging levels. Too noisy log files are difficult to read too and server's operation slows down noticeably.

7. Dealing with DNS

SIP standard leverages DNS. Administrators of `ser` should be aware of impact of DNS on server's operation. Server's attempt to resolve an unresolvable address may block a server process in terms of seconds. To be safer that the server doesn't stop responding due to being blocked by DNS resolving, we recommend the following practices:

- Start a sufficient number of children processes. If one is blocked, the other children will keep serving.
- Use DNS caching. For example, in Linux, there is an `nscd` daemon available for this purpose.
- Process transactions statefully if memory allows. That helps to absorb retransmissions without having to resolve DNS for each of them.

8. Labeling Outbound Requests

Without knowing, which pieces of script code a relayed request visited, troubleshooting would be difficult. Scripts typically apply different processing to different routes such as to IP phones and PSTN gateways. We thus recommend to label outgoing requests with a label describing the type of processing applied to the request.

Attaching "routing-history" hints to relayed requests is as easy as using the **append_hf** action exported by textops module. The following example shows how different labels are attached to requests to which different routing logic was applied.

Example 3-4. "Routing-history" labels

```
# is the request for our domain?
# if so, process it using UsrLoc and label it so.
if (uri=~[@:\.]domain.foo") {
   if (!lookup("location")) {
    sl_send_reply("404", "Not Found");
    break;
   };
   # user found -- forward to him and label the request
   append_hf("P-hint: USRLOC\r\n");
} else {
# it is an outbound request to some other domain --
# indicate it in the routing-history label
   append_hf("P-hint: OUTBOUND\r\n");
};
t_relay();
```

This is how such a labeled requests looks like. The last header field includes a label indicating the script processed the request as outbound.

```
#
U 2002/09/26 02:03:09.807288 195.37.77.101:5060 -> 203.122.14.122:5060
SUBSCRIBE sip:rajesh@203.122.14.122 SIP/2.0.
Max-Forwards: 10.
Via: SIP/2.0/UDP 195.37.77.101;branch=53.b44e9693.0.
Via: SIP/2.0/UDP 203.122.14.115:16819.
From: sip:rajeshacl@iptel.org;tag=5c7cecb3-cfa2-491d-a0eb-72195d4054c4.
To: sip:rajesh@203.122.14.122.
Call-ID: bd6c45b7-2777-4e7a-b1ae-11c9ac2c6a58@203.122.14.115.
CSeq: 2 SUBSCRIBE.
Contact: sip:203.122.14.115:16819.
User-Agent: Windows RTC/1.0.
Proxy-Authorization: Digest username="rajeshacl", realm="iptel.org", al-
gorithm="MD5", uri="sip:rajesh@203.122.14.122", nonce="3d924fe900000000fd6227db9e565b73
Expires: 1800.
Content-Length: 0.
P-hint: OUTBOUND.
```

HOWTOs

This section is a "cookbook" for dealing with common tasks, such as user management or controlling access to PSTN gateways.

1. User Management

There are two tasks related to management of SIP users: maintaining user accounts and maintaining user contacts. Both these jobs can be done using the `serctl` command-line tool. Also, the complimentary web interface, `serweb`, can be used for this purpose as well.

If user authentication is turned on, which is a highly advisable practice, user account must be created before a user can log in. To create a new user account, call the **serctl add** utility with username, password and email as parameters. It is important that the environment `SIP_DOMAIN` is set to your realm and matches realm values used in your script. The realm value is used for calculation of credentials stored in subscriber database, which are bound permanently to this value.

```
[jiri@cat gen_ha1]$ export SIP_DOMAIN=foo.bar
[jiri@cat gen_ha1]$ serctl add newuser secret newuser@foo.bar
```

```
MySql Password:
new user added
```

serctl can also change user's password or remove existing accounts from system permanently.

```
[jiri@cat gen_ha1]$ serctl passwd newuser newpassword
MySql Password:
password change succeeded
[jiri@cat gen_ha1]$ serctl rm newuser
MySql Password:
user removed
```

User contacts are typically automatically uploaded by SIP phones to server during registration process and administrators do not need to worry about them. However, users may wish to append permanent contacts to PSTN gateways or to locations in other administrative domains. To manipulate the contacts in such cases, use serctl ul tool. Note that this is the only correct way to update contacts -- direct changes to back-end MySql database do not affect server's memory. Also note, that if persistence is turned off (usrloc "db_mode" parameter set to "0"), all contacts are gone on server reboot. Make sure that persistence is enabled if you add permanent contacts.

To add a new permanent contact for a user, call serctl ul add <username> <contact>. To delete all user's contacts, call serctl ul rm <username>. serctl ul show <username> prints all current user's contacts.

```
[jiri@cat gen_ha1]$ serctl ul add newuser sip:666@gateway.foo.bar
sip:666@gateway.foo.bar
200 Added to table
('newuser','sip:666@gateway.foo.bar') to 'location'
[jiri@cat gen_ha1]$ serctl ul show newuser
<sip:666@gateway.foo.bar>;q=1.00;expires=1073741812
[jiri@cat gen_ha1]$ serctl ul rm newuser
200 user (location, newuser) deleted
[jiri@cat gen_ha1]$ serctl ul show newuser
404 Username newuser in table location not found
```

2. User Aliases

Frequently, it is desirable for a user to have multiple addresses in a domain. For example, a user with username "john.doe" wants to be reachable at a shorter address "john" or at a nummerical address "12335", so that PSTN callers with digits-only keypad can reach him too.

With ser, you can maintain a special user-location table and translate existing aliases to canonical usernames using the **lookup** action from usrloc module. The following script fragment demonstrates use of **lookup** for this purpose.

Example 3-5. Configuration of Use of Aliases

```
if (!uri==myself) { # request not for our domain...
  route(1); # go somewhere else, where outbound requests are processed
  break;
};
# the request is for our domain -- process registrations first
```

```
if (method=="REGISTER") { route(3); break; };

# look now, if there is an alias in the "aliases" table; don't care
# about return value: whether there is some or not, move ahead then
lookup("aliases");

# there may be aliases which translate to other domain and for which
# local processing is not appropriate; check again, if after the
# alias translation, the request is still for us
if (!uri==myself) { route(1); break; };

# continue with processing for our domain...
...
```

The table with aliases is updated using the `serctl` tool. `serctl alias add <alias> <uri>` adds a new alias, `serctl alias show <user>` prints an existing alias, and `serctl alias rm <user>` removes it.

```
[jiri@cat sip_router]$ serctl alias add 1234 sip:john.doe@foo.bar
sip:john.doe@foo.bar
200 Added to table
('1234','sip:john.doe@foo.bar') to 'aliases'
[jiri@cat sip_router]$ serctl alias add john sip:john.doe@foo.bar
sip:john.doe@foo.bar
200 Added to table
('john','sip:john.doe@foo.bar') to 'aliases'
[jiri@cat sip_router]$ serctl alias show john
<sip:john.doe@foo.bar>;q=1.00;expires=1073741811
[jiri@cat sip_router]$ serctl alias rm john
200 user (aliases, john) deleted
```

Note that persitence needs to be turned on in usrloc module. All changes to aliases will be otherwise lost on server reboot. To enable persistence, set the db_mode usrloc parameter to a non-zero value.

```
# ....load module ...
loadmodule "modules/usrloc/usrloc.so"
# ... turn on persistence -- all changes to user tables are immediately
# flushed to mysql
modparam("usrloc", "db_mode",   1)
# the SQL address:
modparam("usrloc", "db_url","sql://ser:secret@dbhost/ser")
```

3. Access Control (PSTN Gateway)

It is sometimes important to exercise some sort of access control. A typical use case is when `ser` is used to guard a PSTN gateway. If a gateway was not well guarded, unauthorized users would be able to use it to terminate calls in PSTN, and cause high charges to its operator.

There are few issues you need to understand when configuring `ser` for this purpose. First, if a gateway is built or configured to accept calls from anywhere, callers may easily bypass your access control server and communicate with the gateway directly. You then need to enforce at transport layer that signaling is only accepted if coming

via ser and deny SIP packets coming from other hosts and port numbers. Your network must be configured not to allow forged IP addresses. Also, you need to turn on record-routing to assure that all session requests will travel via ser. Otherwise, caller's devices would send subsequent SIP requests directly to your gateway, which would fail because of transport filtering.

Authorization (i.e., the process of determining who may call where) is facilitated in ser using *group membership* concept. Scripts make decisions on whether a caller is authorized to make a call to a specific destination based on user's membership in a group. For example a policy may be set up to allow calls to international destinations only to users, who are members of an "int" group. Before user's group membership is checked, his identity must be verified first. Without cryptographic verification of user's identity, it would be impossible to assert that a caller really is who he claims to be.

The following script demonstrates, how to configure ser as an access control server for a PSTN gateway. The script verifies user identity using digest authentication, checks user's privileges, and forces all requests to visit the server.

Example 3-6. Script for Gateway Access Control

```
#
# $Id: iptel.cfg,v 1.38 2002/10/04 21:40:31 jiri Exp $
#
# example: ser configured as PSTN gateway guard; PSTN gateway is located
# at 192.168.0.10
#

# ------------------ module loading ----------------------------------
-

loadmodule "modules/sl/sl.so"
loadmodule "modules/tm/tm.so"
loadmodule "modules/acc/acc.so"
loadmodule "modules/rr/rr.so"
loadmodule "modules/maxfwd/maxfwd.so"
loadmodule "modules/mysql/mysql.so"
loadmodule "modules/auth/auth.so"

# ----------------- setting module-specific parameters --------------
-

modparam("auth", "db_url","sql://ser:heslo@localhost/ser")
modparam("auth", "calculate_ha1", yes)
modparam("auth", "password_column", "password")

# -- acc params --
modparam("acc", "log_level", 1)
# that is the flag for which we will account -- don't forget to
# set the same one :-)
modparam("acc", "acc_flag", 1 )

# -------------------------  request routing logic ------------------
-

# main routing logic

route{

 /* ********* ROUTINE CHECKS  ********************************** */

 # filter too old messages
 if (!mf_process_maxfwd_header("10")) {
  log("LOG: Too many hops\n");
  sl_send_reply("483","Too Many Hops");
```

```
 break;
};
if (len_gt( max_len )) {
 sl_send_reply("513", "Wow -- Message too large");
 break;
};

/* ********* RR ******************************** */

/* Do strict routing if route headers present */
rewriteFromRoute();
/* record-route INVITEs -- all subsequent requests must visit us */
if (method=="INVITE") {
 addRecordRoute();
};

# now check if it really is a PSTN destination which should be handled
# by our gateway; if not, and the request is an invitation, drop it -
_
# we cannot terminate it in PSTN; relay non-INVITE requests -- it may
# be for example BYEs sent by gateway to call originator
if (!uri=~"sip:\+?[0-9]+@.*") {
 if (method=="INVITE") {
  sl_send_reply("403", "Call cannot be served here");
 } else {
  forward(uri:host, uri:port);
 };
 break;
};

# account completed transactions via syslog
setflag(1);

# free call destinations ... no authentication needed
if ( is_user_in("Request-URI", "free-pstn")  /* free destinations */
   |  uri=~"sip:[79][0-9][0-9][0-9]@.*"  /* local PBX */
   | uri=~"sip:98[0-9][0-9][0-9][0-9]") {
 log("free call");
} else if (src_ip==192.168.0.10) {
 # our gateway doesn't support digest authentication;
 # verify that a request is coming from it by source
 # address
 log("gateway-originated request");
} else {
 # in all other cases, we need to check the request against
 # access control lists; first of all, verify request
 # originator's identity

 if (!proxy_authorize( "gateway" /* realm */,
   "subscriber" /* table name */))  {
  proxy_challenge( "gateway" /* realm */, "0" /* no qop */ );
  break;
 };

 # authorize only for INVITEs -- RR/Contact may result in weird
 # things showing up in d-uri that would break our logic; our
 # major concern is INVITE which causes PSTN costs

 if (method=="INVITE") {

  # does the authenticated user have a permission for local
  # calls (destinations beginning with a single zero)?
  # (i.e., is he in the "local" group?)
  if (uri=~"sip:0[1-9][0-9]+@.*") {
   if (!is_in_group("local")) {
    sl_send_reply("403", "No permission for local calls");
```

```
     break;
    };
   # the same for long-distance (destinations begin with two zeros")
   } else if (uri=~"sip:00[1-9][0-9]+@.*") {
    if (!is_in_group("ld")) {
     sl_send_reply("403", " no permission for LD ");
     break;
    };
   # the same for international calls (three zeros)
   } else if (uri=~"sip:000[1-9][0-9]+@.*") {
    if (!is_in_group("int")) {
     sl_send_reply("403", "International permissions needed");
     break;
    };
   # everything else (e.g., interplanetary calls) is denied
   } else {
    sl_send_reply("403", "Forbidden");
    break;
   };

  }; # INVITE to authorized PSTN

 }; # authorized PSTN

 # if you have passed through all the checks, let your call go to GW!

 rewritehostport("192.168.0.10:5060");

 # forward the request now
 if (!t_relay()) {
  sl_reply_error();
  break;
 };

}
```

Use the `serctl` tool to maintain group membership. `serctl acl grant <username> <group>` makes a user member of a group, `serctl acl show <username>` shows groups of which a user is member, and `serctl acl revoke <username> [<group>]` revokes user's membership in one or all groups.

```
[jiri@cat sip_router]$ serctl acl grant john int
MySql Password:
+------+-----+---------------------+
| user | grp | last_modified       |
+------+-----+---------------------+
| john | int | 2002-12-08 02:09:20 |
+------+-----+---------------------+
```

4. Accounting

In some scenarios, like termination of calls in PSTN, SIP administrators may wish to keep track of placed calls. `ser` can be configured to report on completed transactions. Reports are sent by default to `syslog` facility. Experimental support for RADIUS and mysql accounting exists as well.

Note that `ser` is no way call-stateful. It reports on completed transactions, i.e., after a successful call set up is reported, it drops any call-related state. When a call is

terminated, transactional state for BYE request is created and forgotten again after the transaction completes. This is a feature and not a bug -- keeping only transactional state allows for significantly higher scalability. It is then up to the accounting application to correlate call initiation and termination events.

To enable call accounting, tm and acc modules need to be loaded, requests need to be processed statefuly and labeled for accounting.

Example 3-7. Configuration with Enabled Accounting

```
#
# $Id: iptel.cfg,v 1.38 2002/10/04 21:40:31 jiri Exp $
#
# example: accounting calls to nummerical destinations
#

# ------------------ module loading ----------------------------------
-

loadmodule "modules/tm/tm.so"
loadmodule "modules/acc/acc.so"
loadmodule "modules/sl/sl.so"
loadmodule "modules/maxfwd/maxfwd.so"

# ----------------- setting module-specific parameters --------------
-

# -- acc params --
# set the reporting log level
modparam("acc", "log_level", 1)
# number of flag, which will be used for accounting; if a message is
# labeled with this flag, its completion status will be reported
modparam("acc", "acc_flag", 1 )

# -------------------------  request routing logic ------------------
-

# main routing logic

route{

 /* ********* ROUTINE CHECKS  ********************************** */

 # filter too old messages
 if (!mf_process_maxfwd_header("10")) {
  log("LOG: Too many hops\n");
  sl_send_reply("483","Too Many Hops");
  break;
 };
 if (len_gt( max_len )) {
  sl_send_reply("513", "Wow -- Message too large");
  break;
 };

 # nummerical destinations will be labeled for accounting, others not
 if (uri=~"sip:\+?[0-9]+@") {
  setflag(1);
 };

 # forward the request statefuly now; (we need *stateful* forwarding,
 # because the stateful mode correlates requests with replies and
 # drops retranmissions; otherwise, we would have to report on
 # every single message received)
 if (!t_relay()) {
  sl_reply_error();
```

```
  break;
 };

}
```

5. Reliability

It is essential to guarantee continuous service operation even under erroneous conditions, such as host or network failure. The major issue in such situations is transfer of operation to a backup infrastructure and making clients use it.

The SIP standard's use of DNS SRV records has been explicitly constructed to handle with server failures. There may be multiple servers responsible for a domain and referred to by DNS. If it is impossible to communicate with a primary server, a client can proceed to another one. Backup servers may be located in a different geographic area to minimize risk caused by areal operational disasters: lack of power, flooding, earthquake, etc.

> Unless there are redundant DNS servers, fail-over capability cannot be guaranteed.

Unfortunately, at the moment of writing this documentation (end of December 2002) only very few SIP products actually implement the DNS fail-over mechanism. Unless networks with SIP devices supporting this mechanism are built, alternative mechanisms must be used to force clients to use backup servers. Such a mechanism is disconnecting primary server and replacing him with a backup server locally. It unfortunately precludes geographic dispersion and requires network multihoming to avoid dependency on single IP access. Another method is to update DNS when failure of the primary server is detected. The primary drawback of this method is its latency: it may take long time until all clients learn to use the new server.

The easier part of the redundancy story is replication of `ser` data. `ser` relies on replication capabilities of its back-end database. This works with one exception: user location database. User location database is a frequently accessed table, which is thus cached in server's memory to improve performance. Back-end replication does not affect in-memory tables, unless server reboots. To facilitate replication of user location database, server's SIP replication feature must be enabled in parallel with back-end replication.

The design idea of replication of user location database is easy: Replicate any successful REGISTER requests to a peer server. To assure that digest credentials can be properly verified, both servers need to use the same digest generation secret and maintain synchronized time. A known limitation of this method is it does not replicate user contacts entered in another way, for example using web interface through FIFO server. The following script example shows configuration of a server that replicates all REGISTERs.

Example 3-8. Script for Replication of User Contacts

```
#
# $Id: ser.cfg,v 1.12 2002/10/21 02:40:06 jiri Exp $
#
# demo script showing how to set-up usrloc replication
#
```

```
# ----------- global configuration parameters ------------------------
-

debug=3          # debug level (cmd line: -ddddddddddd)
fork=no
log_stderror=yes # (cmd line: -E)

# ------------------ module loading ----------------------------------
-

loadmodule "modules/mysql/mysql.so"
loadmodule "modules/sl/sl.so"
loadmodule "modules/tm/tm.so"
loadmodule "modules/maxfwd/maxfwd.so"
loadmodule "modules/usrloc/usrloc.so"
loadmodule "modules/registrar/registrar.so"
loadmodule "modules/auth/auth.so"

# ----------------- setting module-specific parameters --------------
-

# digest generation secret; use the same in backup server;
# also, make sure that the backup server has sync'ed time
modparam("auth", "secret", "alsdkhglaksdhfkloiwr")

# -------------------------  request routing logic -------------------
-

# main routing logic

route{

 # initial sanity checks -- messages with
 # max_forwars==0, or excessively long requests
 if (!mf_process_maxfwd_header("10")) {
  sl_send_reply("483","Too Many Hops");
  break;
 };
 if (len_gt( max_len )) {
  sl_send_reply("513", "Message too big");
  break;
 };

 # if the request is for other domain use UsrLoc
 # (in case, it does not work, use the following command
 # with proper names and addresses in it)
 if (uri==myself) {

  if (method=="REGISTER") {

   # verify credentials
   if (!www_authorize("foo.bar", "subscriber")) {
    www_challenge("foo.bar", "0");
    break;
   };

   # if ok, update contacts and ...
   save("location");
   # ... if this REGISTER is not a replica from our
   # peer server, replicate to the peer server
   if (!src_ip==backup.foo.bar) {
    t_replicate("backup.foo.bar", "5060");
   };
   break;
  };
  # do whatever else appropriate for your domain
```

```
  log("non-REGISTER\n");
 };
}
```

6. Stateful versus Stateless Forwarding

`ser` allows both stateless and stateful request processing. This memo explains what are pros and cons of using each method. The rule of thumb is "stateless for scalability, stateful for services".

Stateless forwarding with the **forward(uri:host, uri:port)** action guarantees high scalability. It withstands high load and does not run out of memory. A perfect use of stateless forwarding is load distribution.

Stateful forwarding using the **t_relay()** action is known to scale worse. It can quickly run out of memory and consumes more CPU time. Nevertheless, there are scenarios which are not implementable without stateful processing. In particular:

- *Accounting* requires stateful processing to be able to collect transaction status and issue a single report when a transaction completes.
- *Forking* only works with stateful forwarding. Stateless forwarding only forwards to the default URI out of the whole destination set.
- *DNS resolution*. DNS resolution may be better server with stateful processing. If a request is forwarded to a destination whose address takes long time to resolve, a server process is blocked and unresponsive. Subsequent request retransmissions from client will cause other processes to block too if requests are processed statelessly. As a result, `ser` will quickly run out of available processes. With stateful forwarding, retransmissions are absorbed and do not cause blocking of another process.
- *Forwarding Services*. All sort of services with the "forward_on_event" logic, which rely on **t_on_negative** tm action must be processed statefuly.

7. Serving Multiple Domains

`ser` can be configured to serve multiple domains. To do so, you need to take the following steps:

1. Create separate subscriber and location database table for each domain served and name them uniquely.
2. Configure your script to distinguish between multiple served domains. Use regular expressions for domain matching as described in Example 2-6.
3. Update table names in usrloc and auth actions to reflect names you created in 1.

Troubleshooting

This section gathers practices how to deal with errors known to occur frequently.

1. SIP requests are replied by `ser` with "483 Too Many Hops" or "513 Message Too Large"

In both cases, the reason is probably an error in request routing script which caused an infinite loop. You can easily verify whether this happens by watching SIP traffic on loopback interface. A typical reason for misrouting is a failure to match local domain correctly. If a server fails to recognize a request for itself, it will try to forward it to current URI in believe it would forward them to a foreign domain. Alas, it forwards the request to itself again. This continues to happen until value of max_forwards header field reaches zero or the request grows too big. Solutions is easy: make sure that domain matching is correctly configured. See the Section called *Domain Matching* in Chapter 2 for more information how to get it right.

2. Windows Messenger authentication fails.

The most likely reason for this problem is a bug in Windows Messenger. WM only authenticates if server name in request URI equals authentication realm. After a challenge is sent by SIP server, WM does not resubmit the challenged request at all and pops up authentication window again. If you want to authenticate WM, you need to set up your realm value to equal server name. If your server has no name, IP address can be used as realm too.

3. I receive "ERROR: t_newtran: transaction already in process" in my logs.

That looks like an erroneous use of tm module in script. tm can handle only one transaction per request. If you attempt to instantiate a transaction multiple times, `ser` will complain. Anytime any of **t_newtran**, **t_relay** or **t_relay_to** actions is encountered, tm attempts to instantiate a transaction. Doing so twice fails. Make sure that any of this commands is called only once during script execution.

Notes

1. http://sipsak.berlios.de/

Chapter 4. Application Writing

ser offers several methods to leverage and extend its abilities. They primarily differ in how easy-to-use and powerful they are. The easiest way is to couple ser with external applications via the *exec* module. This module allows execution of logic and URI manipulation by external applications on request receipt. While very simple, many useful services can be implemented this way. External applications can be written in any programming language and do not be aware of SIP at all. They only care of URIs. For example, an external shell script may send an email whenever a request for a user arrives.

Example 4-1. Using External Scripts

```
# send email if a request for johndoe arrives
if (uri=~"^sip:johndoe@") {
     exec_msg("echo 'body: call arrived'|mail -s 'call for you' johndoe");
}
```

The other extreme method for extending ser capabilities is to write new modules in C. This method takes deeper understanding of ser internals but gains the highest flexibility. Modules can implement arbitrary brand-new commands upon which ser scripts can rely on. Guidelines on module programming can be found in ser programmer's handbook available from iptel.org website.

There is a middle way too, which allows easy service creation without in-depth knowledge of ser internals. ser allows external applications to link via its built-in *application FIFO server* described in the next section.

Application FIFO Server

Application FIFO server is a very powerful method to program SIP services. The most valuable benefit is it works with SIP-unaware applications written in any programming language. Textual nature of the FIFO interface allows for easy integration with a lot of existing programs. Today, ser's complementary web-interface, serweb, written in PHP, leverages the FIFO interface when displaying and changing user location records stored in server's memory. It uses this interface to send instant messages too, without any knowledge of underlying SIP stack. Another application relying on the FIFO interface is serctl, ser management utility. The command-line utility can browse server's in-memory user-location database, display running processes and operational statistics.

The way the FIFO server works is similar to how /proc filesystem works on some operating systems. It provides a human-readable way to access ser's internals. Applications dump their requests into the FIFO server and receive a status report when request processing completes. ser exports a lot of its functionality located in both the core and external modules through the FIFO server.

FIFO requests are formed easily. They begin with a command enclosed in colons and followed by name of file or pipe (relative to /tmp/ path), to which a reply should be printed. The first request line may be followed by additional lines with command-specific parameters. For example, the **t_uac** FIFO command for initiating a transaction allows to pass additional header fields and message body to a newly created transaction. Each request is terminated by an empty line. Whole requests must be sent by applications atomically in a single batch to avoid mixing with requests from other applications. Requests are sent to pipe at which ser listens (filename configured by the fifo config file option).

An easy way to use the FIFO interface is via the `serctl` command-line tool. When called along with "fifo", FIFO command name, and optional parameters, the tool generates a FIFO request and prints request result. The following example shows use of this tool with the **uptime** and **which** commands. **uptime** returns server's running time, **which** returns list of available FIFO commands. Note that only the built-in FIFO command set is displayed as no modules were loaded in this example.

Example 4-2. Use of `serctl` to Access FIFO Server

```
[jiri@cat test]$ serctl fifo uptime
Now: Fri Dec  6 17:56:10 2002
Up Since: Fri Dec  6 17:56:07 2002
Up time: 3 [sec]

[jiri@cat test]$ serctl fifo which
ps
which
version
uptime
print
```

The request which the `serctl` command-line tool sent to FIFO server looked like this:

Example 4-3. uptime FIFO Request

```
:uptime:ser_receiver_1114
```

This request contains no parameters and consists only of command name enclosed in colons and name of file, to which a reply should be printed. FIFO replies consist of a status line followed by optional parameters. The status line consists, similarly to SIP reply status, of a three-digit status code and a reason phrase. Status codes with leading digit 2 (200..299) are considered positive, any other values indicate an error. For example, FIFO server returns "500" if execution of a non-existing FIFO command is requested.

Example 4-4. FIFO Errors

```
[jiri@cat sip_router]$ sc fifo foobar
500 command 'foobar' not available
```

A powerful feature of the FIFO server is the ability to export new commands from modules for use by external applications. Currently, usrloc module exports FIFO commands for maintaining in-memory user location database and tm module exports FIFO commands for management of SIP transactions. See the example in `examples/web_im/send_im.php` for how to initiate a SIP transaction (instant message) from a PHP script via the FIFO server. This example uses FIFO command **t_uac**. The command is followed by parameters: header fields and message body. The same FIFO command can be used from other environments to send instant messages too. The following example shows how to send instant messages from a shell script.

Example 4-5. Sending IM From Shell Script

```
#!/bin/sh
#
# call this script to send an instant message; script parameters
# will be displayed in message body
#

cat > /tmp/ser_fifo <<EOF
:t_uac_from:hh
NOTIFY
sip:originator@foo.bar
sip:receiver@127.0.0.1
foo: bar_special_header
x: y
p_header: p_value
Contact: <sip:devnull@192.168.0.100:9>
Content-Type: text/plain; charset=UTF-8

Hello world!!!! $@
.
EOF
```

See the Section called *FIFO Commands Reference* in Chapter 5 for a complete listing of FIFO commands available with current `ser` distribution.

Chapter 5. Reference

Core Options

Core options are located in beginning of configuration file and affect behaviour of the server.

- `debug` - Set log level, this is number between 0 and 9. Default is 0.
- `fork` - If set to yes, the server will spawn children. If set to no, the main process will be processing all messages. Default is yes.

  Note: Disabling child spawning is useful mainly for debugging. When `fork` is turned off, some features are unavailable: there is no attendant process, no pid file is generated, and server listens only at one address. Make sure you are debugging the same interface at which `ser` listens. The easiest way to do so is to set the interface using `listen` option explicitly.

- `log_stderror` - If set to yes, the server will print its debugging information to standard error output. If set to no, **syslog** will be used. Default is no (printing to syslog).
- `listen` - list of all IP addresses or hostnames SER should listen on.

  Note: This parameter may repeat several times, then SER will listen on all addresses. For example, the following command-line options (equivalent to "listen" config option) may be used: **ser -l foo -l bar -p 5061 -l x -l y** will listen on foo:5060, bar:5061 & x:5061 & y:5061

- `alias` - Add IP addresses or hostnames to list of name aliases. All requests with hostname matching an alias will satisfy the condition "uri==myself".
- `dns` - Uses dns to check if it is necessary to add a "received=" field to a via. Default is no.
- `rev_dns` - Same as dns but use reverse DNS. Default is no.
- `port` - Listens on the specified port (default 5060). It applies to the last address specified in listen and to all the following that do not have a corresponding "port" option.
- `maxbuffer` - Maximum receive buffer size which will not be exceeded by the auto-probing procedure even if the OS allows. Default value is MAX_RECV_BUFFER_SIZE, which is 256k.
- `children` - Specifies how many children should be created when fork is set to yes. Default value is CHILD_NO, which is 8. Running multiple children allows a server to server multiple requests in parallel when request processing block (e.g., on DNS lookup). Note that `ser` typically spawns additional processes, such as timer process or FIFO server. If FIFO server is turned on, you can watch running processes using the `serctl` utility.
- `check_via` - Turn on or off Via host checking when forwarding replies. Default is no.
- `syn_branch` - Shall the server use stateful synonym branches? It is faster but not reboot-safe. Default is yes.
- `mem_log` - Debugging level for memory statistics. Default is L_DBG -- memory statistics are dumped only if `debug` is set high.

- `sip_warning` - Should replies include extensive warnings? By default yes, it is good for trouble-shooting.
- `fifo` - FIFO special file pathname, for example "/tmp/ser_fifo". Default is no filename -- no FIFO server is started then. We recommend to set it so that accompanying applications such as `serweb` or `serctl` can communicate with `ser`.
- `fifo_mode` - Permissions of the FIFO special file.
- `server_signature` - Should locally-generated messages include server's signature? By default yes, it is good for trouble-shooting.
- `reply_to_via` - A hint to reply modules whether they should send reply to IP advertised in Via or IP from which a request came.
- `user` | `uid` - uid to be used by the server.
- `group` | `gid` - gid to be used by the server.
- `loadmodule` - Specifies a module to be loaded (for example "/usr/lib/ser/modules/tm.so")
- `modparam` - Module parameter configuration. The commands takes three parameters:
  - *module* - Module in which the parameter resides.
  - *parameter* - Name of the parameter to be configured.
  - *value* - New value of the parameter.

Core Commands

Route Blocks and Process Control

Route block and process control keywords determine the order in which SIP requests are processed.

- **route[number]{...}** - This marks a "route block" in configuration files. route blocks are basic building blocks of `ser` scripts. Each route block contains a sequence of `SER` actions enclosed in braces. Multiple route blocks can be included in a configuration file. When script execution begins on request receipt, route block number 0 is entered. Other route blocks serve as a kind of sub-routines and may be entered by calling the action **route(n)**, where n is number of the block. The action **break** exits currently executed route block. It stops script execution for route block number 0 or returns to calling route block otherwise.

  Example 5-1. route

```
route[0] {
        # call routing block number 2
 route(2);
}

route[2] {
    forward("host.foo.bar", 5060);
}
```

- **reply_route** is used to restart request processing when a negative reply for a previously relayed request is received. It is only used along with tm module, which stores the original requests and can return to their processing later. To activate processing of a **reply_route** block, call the TM action **t_on_negative(route_number)** before calling **t_relay**. When a negative reply comes back, the desired **reply_route** will be entered and processing of the original request may continue.

The set of actions applicable from within **reply_route** blocks is limited. Permitted actions are URI-manipulation actions, logging and sending stateful replies using **t_reply_unsafe**. Use of other actions may lead to unpredictable results. (We plan to add syntactical checks in the future so that improper action use will be detected during server start-up.)

Example 5-2. reply_route

```
reply_route[1] {
    # for some reason, the original forwarding attempt
    # failed, try at another URI
    append_branch("sip:nonsense@iptel.org");
    # if this new attempt fails too, try another reply_route
    t_on_negative("2");
}
```

- The action **break** exits currently executed route block. It stops script execution for route block number 0 or returns to calling route block otherwise.

  Note: We recommend to use **break** after any request forwarding or replying. This practice helps to avoid erroneous scripts that continue execution and mistakenly send another reply or forward a request to another place, resulting in protocol confusion.

  Example: break;

- **route(n)** - call routing block route[n]{...}; when the routing block n finishes processing, control is passed back to current block and processing continues.
- **if (condition) statement** - Conditional statement.

  Example 5-3. Use of if

```
if (method=="REGISTER) {
    log("register received\n");
};
```

- **if - else** - If-Else Conditional statement.

  Example 5-4. Use of if-else

```
if (method=="REGISTER) {
    log("register received\n");
} else {
    log("non-register received\n");
};
```

Flag Manipulation

- **setflag** - Set flag in the message.

  Example: setflag(1);

- **resetflag** - Reset flag in the message.

  Example: resetflag(1);

- **isflagset** - Test whether a particular flag is set.

  Example 5-5. isflagset

```
if (isflagset(1)) {
    ....
};
```

Manipulation of URI and Destination Set

- **rewritehost | sethost | seth** - Rewrite host part of the Request URI.

  Example: sethost("foo.bar.com");

- **rewritehostport | sethostport | sethp** - Rewrite host and port part of the Request URI.

  Example: sethostport("foo.bar.com:5060");

- **rewriteuser | setuser | setu** - Rewrite or set username part of the Request URI.

  Example: setuser("joe");

- **rewriteuserpass | setuserpass | setup** - Rewrite or set username and password part of the Request URI.

  Example: setuserpass("joe:mypass");

- **rewriteport | setport | setp** - Rewrite or set port of the Request URI.

  Example: setport("5060");

- **rewriteuri | seturi** - Rewrite or set the whole Request URI.

  Example: seturi("sip:joe@foo.bar.com:5060");

- **revert_uri** - Revert changes made to the Request URI and use original Request URI.

  Example: revert_uri();

- **prefix** - Add prefix to username in Request URI.

  Example: prefix("123");

- **strip** - Remove first n characters of username in Request URI.

  Example: strip(3);

- **append_branch** - Append a new destination to destination set of the message.

Example 5-6. Use of append_branch

```
# redirect to these two destinations: a@foo.bar and b@foo.bar
# 1) rewrite the current URI
rewriteuri("sip:a@foo.bar");
# 2) append another entry to the destination ser
append_branch("sip:b@foo.bar");
# redirect now
sl_send_reply("300", "redirection");
```

Message Forwarding

- **forward(uri, port)** - Forward the request to given destination statelessly. The uri and port parameters may take special values 'uri:host' and 'uri:port' respectively, in which case SER forwards to destination set in current URI. All other elements in a destination set are ignored by stateless forwarding.

  Example: forward("foo.bar.com"); # port defaults to 5060

- **send** - Send the message as is to a third party

  Example: send("foo.bar.com");

Logging

- **log([level], message)** - Log a message.

  Example: log(1, "This is a message with high log-level set to 1\n");

  Logging is very useful for troubleshooting or attracting administrator's attention to unusual situations. `ser` reports log messages to `syslog` facility unless it is configured to print them to `stderr` with the `log_stderr` configuration option. Log messages are only issued if their log level exceeds threshold set with the `debug` configuration option. If log level is omitted, messages are issued at log level 4.

Miscellaneous

- **len_gt** - If length of the message is greater than value given as parameter, the command will return 1 (indicating true). Otherwise -1 (indicating false) will be returned. It may take 'max_len' as parameter, in which case message size is limited to internal buffer size BUF_SIZE (3040 by default).

  Example 5-7. Use of len_gt

```
# deny all requests larger in size than 1 kilobyte
if (len_gt(1024)) {
    sl_send_reply("513", "Too big");
    break;
};
```

Command Line Parameters

Note: Command-Line parameters may be overridden by configuration file options which take precedence over them.

- *-h* - Displays a short usage description, including all available options.
- *-c* - Performs loop checks and computes branches.
- *-r* - Uses dns to check if it is necessary to add a "received=" field to a via.
- *-R* - Same as -r but uses reverse dns.
- *-v* - Turns on via host checking when forwarding replies.
- *-d* - Turns on debugging, multiple -d increase debugging level.
- *-D* - Runs ser in the foreground (it doesn't fork into daemon mode).
- *-E* - Sends all the log messages to stderr.
- *-V* - Displays the version number.
- *-f config-file* - Reads the configuration from "config-file" (default ./ser.cfg).
- *-l address* - Listens on the specified address. Multiple -l mean listening on multiple addresses. The default behaviour is to listen on all the ipv4 interfaces.
- *-p port* - Listens on the specified port (default 5060). It applies to the last address specified with -l and to all the following that do not have a corresponding -p.
- *-n processes-no* - Specifies the number of children processes forked per interface (default 8).
- *-b max_rcv_buf_size* - Maximum receive buffer size which will not be exceeded by the auto-probing procedure even if the OS allows.
- *-m shared_mem_size* - Size of the shared memory which will be allocated (in Megabytes).
- *-w working-dir* - Specifies the working directory. In the very improbable event that will crash, the core file will be generated here.
- *-t chroot-dir* - Forces ser to chroot after reading the config file.
- *-u uid* - Changes the user id under which ser runs.
- *-g gid* - Changes the group id under which ser runs.
- *-P pid-file* - Creates a file containing the pid of the main ser process.
- *-i fifo-path* - Creates a fifo, useful for monitoring ser status.

`serctl` command

`serctl` is a command-line utility which allows to perform most of management tasks needed to operate a server: adding users, changing their passwords, watching server status, etc. Usage of utility is as follows:

Example 5-8. `serctl` usage

```
usage:
           * subscribers *
 serctl add <username> <password> <email> .. add a new subscriber (*)
 serctl passwd <username> <passwd> ......... change user's password (*)
 serctl rm <username> ...................... delete a user (*)
```

```
 serctl mail <username> ..................... send an email to a user
 serctl alias show [<alias>] ................ show aliases
 serctl alias rm <alias> .................... remove an alias
 serctl alias add <alias> <uri> ............. add an aliases

          * access control lists *
 serctl acl show [<username>] ............... show user membership
 serctl acl grant <username> <group> ....... grant user membership (*)
 serctl acl revoke <username> [<group>] .... grant user membership(s) (*)

          * usrloc *
 serctl ul show [<username>]................. show in-RAM online users
 serctl ul rm <username> .................... delete user's UsrLoc entries
 serctl ul add <username> <uri> ............. introduce a permanent Ur-
Loc entry
 serctl showdb [<username>] ................. show online users flushed in DB

          * server health *
 serctl monitor ............................. show internal status
 serctl ps .................................. show runnig processes
 serctl fifo ................................ send raw commands to FIFO

   Commands labeled with (*) will prompt for a MySQL password.
   If the variable PW is set, the password will not be prompted.
```

Note: Prior to using the utility, you have to first set the environment variable `SIP_DOMAIN` to locally appropriate value (e.g., "foo.com"). It is needed for calculation of user credentials, which depend on SIP digest realm.

Example 5-9. Example Output of Server Watching Command sc monitor

```
[cycle #: 2; if constant make sure server lives and fifo is on]
Server: Sip EXpress router(0.8.8 (i386/linux))
Now: Thu Sep 26 23:16:48 2002
Up Since: Thu Sep 26 12:35:27 2002
Up time: 38481 [sec]

Transaction Statistics
Current: 0 (0 waiting) Total: 606 (0 local)
Replied localy: 34
Completion status 6xx: 0, 5xx: 1, 4xx: 86, 3xx: 0,2xx: 519

Stateless Server Statistics
200: 6218 202: 0 2xx: 0
300: 0 301: 0 302: 0 3xx: 0
400: 0 401: 7412 403: 2 404: 1258 407: 116 408: 0 483: 0 4xx: 25      500: 0 5xx: 0
6xx: 0
xxx: 0
failures: 0

UsrLoc Stats
Domain Registered Expired
'aliases' 9 0
'location' 29 17
```

Modules

Module description is currently located in READMEs of respective module directories. In the current `ser` distribution, there are the following modules:

- *acc* -- call accounting using `syslog` facility. Depends on tm.
- *auth* -- digest authentication. Depends on sl and mysql.
- *exec* -- execution of shell programs.
- *jabber* -- gateway between SIMPLE and Jabber instant messaging. Depends on tm and mysql.
- *maxfwd* -- checking max-forwards header field.
- *msilo* -- message silo. Store for undelivered instant messages. Depends on tm and mysql.
- *mysql* -- mysql database back-end.
- *registrar, usrloc* -- User Location database. Works in in-memory mode or with mysql persistence support. Depends on sl, and on mysql if configured for use with mysql.
- *rr* -- Record Routing (strict and loose)
- *sl* -- stateless User Agent server.
- *sms* -- SIMPLE/SMS gateway. Depends on tm. Takes special hardware.
- *textops* -- textual request operations.
- *tm* -- transaction manager (stateful processing).

The most frequently used actions exported by modules are summarized in Table 5-1. For a full explanation of module actions, see documentation in respective module directories in source distribution of `ser`.

Table 5-1. Frequently Used Module Actions

Command	Modules	Parameters	Comments
addRecordRoute	rr	none	record-route request
append_hf	textops	header field	append a header field to the end of request's header
check_from	auth	none	check if username in from header field matches authentication id
check_to	auth	none	check if username in To header field matched authentication id
exec_uri	exec	command name	execute an external command and replace destination set with its output

Command	Modules	Parameters	Comments
is_in_group	auth	group name	check if a user, as identified by digest credentials, is a member of a group‘
is_user	auth	user id	returns true if request credentials belong to a user
is_user_in	auth	user, group	check if user is member of a group; user can be gained from request URI ("Request-URI"), To header field ("To"), From header field ("From") or digest credentials ("Credentials")
lookup	usrloc	table name	attempt to translate request URI using user location database; returns false if no contact for user found;
mf_process_maxfwd	maxfwd_header	default max_forwards value	return true, if request’s max_forwards value has not reached zero yet; if none is included in the request, set it to value in parameter
proxy_authorize	auth	realm, subscriber table	returns true if requests contains proper credentials, false otherwise
proxy_challenge	auth	realm, qop	challenge user to submit digest credentials; qop may be turned off for backwards compatibility with elderly implementations
rewriteFromRoute	rr	none	strict routing: use Route header field if present in request
save	usrloc	table name	for use in registrar: save content of Contact header fields in user location database and reply with 200

Command	Modules	Parameters	Comments
search	textops	regular expression	search for a regular expression match in request
sl_send_reply	sl	status code, reason phrase	reply a request statelessly
t_relay	tm	none	stateful forwarding to locations in current destination set
t_on_negative	tm	reply_route number	set reply_route block which shall be entered if stateful forwarding fails
t_replicate	tm	host, port number	replicate a request to a destination

FIFO Commands Reference

This section lists currently supported FIFO commands. Some of them are built-in in `ser` core, whereas others are exported by modules. The most important exporters are now tm and usrloc module. tm FIFO commands allow users to initiate transactions without knowledge of underlying SIP stack. usrloc FIFO commands allow users to access in-memory user-location database. Note that that is the only way how to affect content of the data-base in real-time. Changes to MySql database do not affect in-memory table unless `ser` is restarted.

Table 5-2. FIFO Commands

Command	Module	Parameters	Comments
ps	core	none	prints running `ser` processes
which	core	none	prints list of available FIFO commands
version	core	none	prints version of `ser`
uptime	core	none	prints `ser`'s running time
sl_stats	sl	none	prints statistics for sl module
t_stats	tm	none	print statistics for tm module
t_hash	tm	none	print occupation of transaction table (mainly for debugging)

Command	Module	Parameters	Comments
t_uac_from	tm	method, sender's URI, destination URI, optional header fields terminated by empty line, message body terminated by a line with a single dot	initiate a transaction
ul_stats	usrloc	none	print usrloc statistics
ul_rm	usrloc	table name, user name	remove all user's contacts from user-location database
ul_rm_contact	usrloc	table name, user name, contact	remove a user's contact from user-location database
ul_dump	usrloc	none	print content of in-memory user-location database
ul_flush	usrloc	none	flush content of in-memory user-location cache in persistent database (MySQL)
ul_add	usrloc	table name, user name, contact, expiration, priority (q)	insert a contact address in user-location database
ul_show_contact	usrloc	table, user name	show user's contact addresses in user-location database

Used Database Tables

`ser` includes MySQL support to guarantee data persistence across server reboots and storage of users' web environment. The data stored in the database include user profiles, access control lists, user location, etc. Note that users are not supposed to alter the data directly, as it could introduce inconsistency between data on persistence storage and in server's memory. The following list enumerates used tables and explains their purpose.

- subscriber -- table of users. It includes user names and security credentials, as well as additional user information.
- reserved -- reserved user names. `serweb` does not permit creation of accounts with name on this list.
- phonebook -- user's personal phonebooks. Accessible via `serweb`.

- pending -- table of unconfirmed subscription requests. Used by `serweb`.
- missed_calls -- table of missed calls. Can be fed by acc modules if mysql support is turned on. Displayed by `serweb`.
- location -- user contacts. Typically updated through `ser`'r registrar functionality.
- grp -- group membership. Used by auth module to determine whether a user belongs to a group.
- event -- allows users to subscribe to additional services. Currently unused.
- aliases -- keeps track of alternative user names.
- active_sessions -- keeps track of currently active web sessions. For use by `serweb`.
- acc -- keeps track of accounted calls. Can be fed by acc module if mysql support is turned on. Displayed by `serweb`.
- config -- maintains attribute-value pairs for keeping various information. Currently not used.
- silo -- message store for instant messages which could not have been delivered.
- version -- keeps version number of each table definition.