

Lire User's Manual

Joost van Baal

Wessel Dankers

Francis J. Lacoste

Wolfgang Sourdeau

Egon L. Willighagen

Lire User's Manual

by Joost van Baal, Wessel Dankers, Francis J. Lacoste, Wolfgang Sourdeau, and Egon L. Willighagen

Copyright © 2000, 2001, 2002, 2003, 2004 Stichting LogReport Foundation

This manual 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 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 manual (see COPYING); if not, check with <http://www.gnu.org/copyleft/gpl.html> (<http://www.gnu.org/copyleft/gpl.html>) or write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111, USA.

Revision History

Revision 1.5 \$Date: 2004/04/06 21:26:22 \$

\$Id: user-manual.dbx,v 1.77 2004/04/06 21:26:22 wsourdeau Exp \$

Table of Contents

Preface.....	i
What This Book Contains	i
How Is This Book Organized?	i
If You Don't Find Something In This Manual	i
I. Lire Overview	i
1. Introducing Lire.....	1
What Is Lire?	1
Supported Systems.....	1
Supported Applications	1
Supported Output Format	4
What Lire Can't Do	5
2. Installing Lire	6
Client Installation.....	6
Requirements	6
Installing.....	6
Standalone Installation.....	6
Requirements	6
Minimum Requirements	6
Requirements for Other Output Formats	7
Other Optional Requirements.....	8
Installing.....	8
Anonymized Client Installation.....	9
Requirements	10
Installing.....	10
Responder Installation	10
Requirements	10
Installation.....	10
Installing Under MTA's using procmail as their MDA	11
Installing Under Exim.....	12
Installing Under qmail	13
3. Running Lire	14
Lire's configuration system.....	14
Getting Started	14
Using A Responder.....	14
Generating A Report From A Log File.....	15
Selecting Output Format	15
Including Charts in the Report	16
Merging Reports	16
Manual Merging.....	16
Automating Merging using cron: some ideas	17
Gotchas.....	17
Holes in your reporting period	17
Take care when changing report configuration file parameters	17
Sending Anonymized Log Files To A Responder	18
Processing The Responder's Results.....	18
Running Lire In A Server Cluster.....	19
Using Mail.....	19
4. Automating Lire	20
Automatically Processing Log Files Using Cron	20

Configuring lr_cron	20
Using lr_cron within Cron	20
5. Customizing Lire's Reports.....	21
The Report's Configuration File	21
Selecting Subreports	22
Reordering The Subreports.....	22
Changing Parameters	22
Using Subreports On Filtered Input.....	23
II. Reports Reference.....	24
6. Database Reports.....	25
Supported Log Format	25
MySQL's Log	25
Report Descriptions and Configuration	25
Actions By Period Database Report.....	25
Most Active Users Database Report	26
Most Accessed Databases Database Report.....	26
Queries By Type.....	26
Filter Descriptions and Configuration	26
7. Dialup Reports.....	28
Supported Log Format	28
Report Descriptions and Configuration.....	28
Connections By Connection Time Dialup Report	28
Connections By Period Dialup Report	28
Connections By Connection Type Dialup Report	28
Connections By Connect Status Dialup Report.....	28
Connections By Hangup Status Dialup Report	29
Connection Time By Type By Period Dialup Report	29
Most Connections From Telephone Number Dialup Report.....	29
Most Connections To Telephone Number Dialup Report	29
Cost By Period Dialup Report.....	30
Cost By Telephone Number Dialup Report.....	30
8. DNS Reports	31
Supported Log Format	31
Bind8 Query Log	31
Bind9 Query Log	31
Report Descriptions and Configuration	32
Top Requesting Hosts Report.....	32
Top Requesting Hosts Report.....	33
Top Requested Names Report.....	33
Distribution of Request Types by Method DNS Report	33
Distribution of Request Types Report.....	34
Requests By Period DNS Report	34
Requests By Timeslot DNS Report.....	34
Requests by Period by Method DNS Report	34
Requests by Timeslot by Method DNS Report.....	35
Filter Descriptions and Configuration	35
Select Resolver Filter	35
9. DNS Zone Reports	37
Supported Log Format	37
Report Descriptions and Configuration	37
Top Requesting AXFR Hosts Report.....	37

Top Requesting AXFR Hosts Report	37
Top Requesting AXFR Hosts Report	37
Top Denied AXFR Requests	38
Top Denied Dynamic DNS Updates Report	38
Top Newly Loaded Zones Report	39
Filter Descriptions and Configuration	39
AXFR denied only	39
AXFR denied only	39
AXFR denied only	39
Loaded zones only	39
Loaded zones only	40
10. Email Reports	41
Supported Log Format	41
ArGoSoft Mail Server	41
Exim	41
Netscape Messaging Server	42
Postfix	42
Qmail	43
Sendmail	44
Reports' Descriptions and Configuration	44
Deliveries Attempts By Period By Status Email Report	44
Deliveries Attempts By Period Email Report	45
Deliveries Attempts By Delay Email Report	45
Deliveries Attempts By Size Email Report	45
Failed Deliveries By Relay Email Report	46
Different From Domain by Period Email Report	46
Different From Email for Selected Domains by Period Email Report	46
Different To Domain by Period Email Report	47
Different To Email for Selected Domains by Period Email Report	47
Highest Average Delay By To Relay And To Domain Email Report	48
Most Deliveries Between Relays Email Report	48
Most Deliveries From Domain Email Report	48
Most Deliveries From User By Domain Email Report	48
Most Deliveries From Relay Email Report	49
Largest Email Exchange Email Report	49
Largest Email Exchange Per Relay Pair Email Report	50
Status Summary For Most Common Domains	50
Most Deliveries To Domain Email Report	50
Most Deliveries To User By Domain Email Report	51
Most Deliveries From Relay Email Report	51
Largest Volume Received From Domain Email Report	51
Largest Volume Sent To Domain Email Report	52
Tracked Recipients Email Report	52
Tracked Senders Email Report	52
Volume Delivered By Period Email Report	53
Email Summary	53
Filters' Descriptions and Configuration	53
Select Client IP Filter	53
11. Firewall Reports	55
Supported Log Format	55
Cisco ACL	55
IPChains	55

IP Filter	56
IPTables.....	56
WebTrends Enhanced Log Format.....	57
Report Descriptions and Configuration	58
Bytes by Period Firewall Report	58
Traffic's Volume by Rule Firewall Report	58
Bytes by Timeslot Firewall Report	58
Top Bytes per From-IP Report.....	59
Top Bytes per From-IP per Port Report	59
Top Bytes per To-ip Report	60
Top Bytes per destination IP per Port Report.....	60
Top blocked tcp packets per source IP per destination port Report	60
Packets by Period Firewall Report	61
Packets by Rule Firewall Report.....	61
Packets by Timeslot Firewall Report	61
Top Volume to Destination by Source Firewall Report	62
Top Volume to Destination by Source Firewall Report	62
Top Messages Firewall Report	63
Top Messages Firewall Report	63
Top Messages Firewall Report	63
Top Packets by Source IP Report.....	64
Top Packets by Destination IP Report.....	64
Top Packets by Destination by Port Firewall Report	64
Top Packets Destination by Source Firewall Report.....	65
Top Packets Source by Destination Firewall Report.....	65
Filter Descriptions and Configuration	66
Exclude Source IP Filter	66
Exclude Destination IP Filter	66
Select Action Filter	66
Select Source IP Filter.....	67
Select Destination IP Filter	67
Select Destination Port.....	67
12. FTP Reports.....	69
Supported Log Format	69
Microsoft Internet Information Server	69
Xferlog	69
Report Descriptions and Configuration	70
Top Remote Host FTP Report.....	70
Bytes by Period FTP Report	70
Bytes by User by Period FTP Report.....	70
Bytes by Period by User FTP Report.....	71
Bytes by Direction by User with count by Period FTP Report.....	71
Top Files FTP Report	72
Top Uploaded Files FTP Report	72
Top Downloaded Files FTP Report.....	72
Top Users FTP Report.....	73
Top by User (Bytes Transferred) FTP Report	73
Tracked Users FTP Report.....	73
Tracked Files FTP Report	74
Number of Transfers by Direction FTP Report	74
Number of Transfers by Transfer Type FTP Report	74
Each Transfer by Filename Report.....	74

Filter Descriptions and Configuration	75
13. Message Store Reports	76
Supported Log Formats	76
Report Descriptions and Configuration	76
Failed Logins by Period Message Store Report.....	76
Closed Event by Period Message Store Report.....	76
Events by Protocol Message Store Report	76
Successful Login by Period Message Store Report	77
Top User Logins Message Store Report.....	77
Top User Most Message Leftover in Store Report.....	77
Top User Most Message Leftover Store Report.....	78
Top User Most Message Store Report	78
Top User Most Message Store Report	78
Top Users doing Select Message Store Report	79
Unique Users by Period Message Store Report	79
Filter Descriptions and Configuration	79
Select Client Host Filter	79
14. Print Reports.....	81
Supported Log Format	81
CUPS page_log	81
LPRng Account Log File	81
Report Descriptions and Configuration	81
Jobs per Printer Print Report.....	82
Top Users Print Report.....	82
Number Of Jobs For Each Number Of Sheets	82
Jobs per Period Print Report	82
Jobs per User per Period Print Report.....	83
Jobs per Printer per Period Print Report	83
Sheets per User Print Report.....	83
Sheets per Period Print Report	83
Sheets per User per Period Print Report	84
Billing Report.....	84
Billing per Printer Report.....	84
Filter Descriptions and Configuration	84
15. Proxy Reports	86
Supported Log Formats	86
Microsoft Internet Security and Acceleration Server.....	86
Squid	86
WebTrends Enhanced Format	87
Report Descriptions and Configuration	88
Bytes by Cache Result	88
Bytes by Object's Source	88
Bytes Transferred By Period Proxy Report.....	88
Bytes Transferred By Timeslot Proxy Report.....	88
Requests by Cache Result	89
Requests By Period Proxy Report.....	89
Requests By Size Proxy Report	89
Number of Requests By Timeslot Proxy Report.....	89
Requests By Request's Time Proxy Report	90
Top Clients by Destinations Proxy Report.....	90
Top Destinations by Number of Requests.....	90
Top Destinations by Bytes Downloaded	91

Top Destinations by Clients	91
Top Destinations by Users Proxy Report	92
Top Users by Destinations Proxy Report	92
Top MIME types by Transferred Size	92
Top Users by Bytes Proxy Report	93
Top URLs by Users Proxy Report	93
Filter Descriptions and Configuration	93
Select Cache Result Filter	94
16. Syslog Reports	95
Supported Log Formats	95
Report Descriptions and Configuration	95
Messages by Facility Syslog Report	95
Messages by Level Syslog Report	95
Messages by Period Syslog Report	95
Top Hosts Syslog Report	95
Top Messages by Period Syslog Report	96
Top Messages by Process Syslog Report	96
Top Messages Syslog Report	97
Top Processes by Period Syslog Report	97
Top Processes Syslog Report	97
Filter Descriptions and Configuration	98
Exclude Message Filter	98
Exclude Priority Filter	98
Exclude Process Filter	98
Select Host Filter	99
Select Message Filter	99
Select Priority Filter	99
Select process Filter	100
17. WWW Reports	101
Supported Log Format	101
Common Log Format	101
Combined Log Format	102
CLF With mod_gzip Extensions	102
Referer Log Format	102
Logs With Virtual Host Information	103
W3C Extended Log Format	103
Report Descriptions and Configuration	104
Bytes By Period WWW Report	104
Bytes Per Directory WWW Report	104
Bytes By HTTP Result By Period WWW Report	104
Bytes By HTTP Result WWW Report	105
Bytes Per Request WWW Report	105
Client Hosts By Period WWW Report	105
Search Engines with Keywords Report	105
Requests By Browser WWW Report	106
Number of Requests By Period WWW Report	106
Requests By Browser Language WWW Report	107
Requests By HTTP Method WWW Report	107
Requests By OS WWW Report	107
Requests By Result By Period WWW Report	107
Requests By HTTP Result WWW Report	108
Requests By Gzip Result WWW Report	108

Requests By Robot Report	108
Requests By Top Level Domain Report.....	108
Requests By Attack Report	109
Requests By Keywords Report	109
Requests By User Agent WWW Report	109
Requests By Search Engines Report.....	109
Number of Requests By Size WWW Report	109
Number of Requests By Timeslot WWW Report.....	110
Requests By HTTP Protocol Version WWW Report.....	110
Average Compression By File Type WWW Report	110
Most Averaged Compressed Requested File WWW Report	111
Top Client By HTTP Result WWW Report.....	111
Top Client by Size WWW Report.....	111
Top Client WWW Report.....	112
Last Pages By Session WWW Report	112
First Pages By Session WWW Report	112
Most Travelled Referer -> Page Connections WWW Report.....	113
Top Referring Pages WWW Report.....	113
Top Referring Pages By Requested Page WWW Report.....	114
Top Referring Sites WWW Report	114
Most Requested Pages WWW Report	115
Top Traversals WWW Report.....	115
Top URLs By HTTP Result WWW Report.....	116
Most Requested URLs By Client Host WWW Report	116
User Sessions By Period WWW Report	116
Recurring Visitors WWW Report	117
Visit times User Session WWW Report.....	117
Page Counts User Session WWW Report.....	117
Filter Descriptions and Configuration	117
Select URL Filter	118
Select Sessions by Page Filter.....	118
Select Client Host Filter	118
Exclude URL Filter	118
Exclude Sessions by Page Filter	119
Exclude Client Host Filter.....	119
Exclude Referer Filter	119
III. Lire Reference.....	121
18. Installation Parameters	122
./configure parameters	122
Installation Environment Variables.....	123
19. Lire Logging and Error Messages.....	125
Logging.....	125
Log Messages	125
20. Lire Installation Layout.....	127

List of Examples

3-1. Sending a Log File For Processing To A Responder.....	15
3-2. Generating a Report With Ir_log2report.....	15
3-3. Generating A HTML Report	15
3-4. Generating A HTML Report With Charts	16
3-5. Merging Reports	16
3-6. Sending An Anonymized Postfix Log File To A Responder.....	18
3-7. Deanonymizing and Generating A HTML Report	18
5-1. Commented Report Configuration File	21
5-2. FTP Report Configuration File	22
6-1. Sample MySQL Log File	25
8-1. Enabling Query Log In Bind	31
8-2. Sample Bind 8 Query Log	31
8-3. Sample Bind 9 Query Log	31
10-1. ArGoSoft Mail Server Log Sample	41
10-2. Exim Log Sample	41
10-3. Netscape Messaging Server Log Sample	42
10-4. Postfix Log Sample.....	42
10-5. Qmail Log Sample.....	43
10-6. Sendmail Log Sample.....	44
11-1. IOS Log Sample	55
11-2. IPChains Log Sample	55
11-3. IP Filter Log Sample.....	56
11-4. IPTables Log Sample.....	56
11-5. WELF Log Sample.....	57
11-6. SonicWall Log Sample	58
12-1. Microsoft Internet Information Server FTP Log Sample	69
12-2. Xferlog Log Sample	69
14-1. CUPS page_log Log Sample	81
14-2. LPRng Log Sample	81
15-1. Microsoft Internet Security and Acceleration Server Log Sample.....	86
15-2. Squid Log Sample	86
15-3. WELF Log Sample.....	87

Preface

Log file analysis is both an essential and tedious part of system administration. It is essential because it's the best way of profiling the usage of the service installed on the network. It's tedious because programs generate a lot of data and tools to report on this data are unavailable or incomplete and when such tools exists, they are specific to one product, which means that you can't compare your qmail and Exim mail servers.

Lire is a software package developed by the Stichting LogReport Foundation to generate useful reports from raw log files of various network programs. Multiple programs are supported for various types of network services. Lire also supports various output formats for the generated reports.

What This Book Contains

This book is the *Lire User's Manual*. It describes how to install, configure and use Lire. The intended audience is system administrators who want to install and use Lire to gather informations about the services operating on their network.

There is another book, the *Lire Developer's Manual* that is intended for system administrators or programmers that want to extend Lire or want to understand its architecture and design.

How Is This Book Organized?

This book is divided into three parts. Part I gives an overview of what Lire can achieve for you. It explains how to install Lire and gives simple usage patterns for various kinds of environments.

Part II contains comprehensive informations on all the reports that can be generated by Lire. It describes all the supported log files and gives the descriptions and customizable parameters for each report.

Finally, you will find in Part III reference material on all installation options and on all the runtime parameters of Lire.

If You Don't Find Something In This Manual

You can report typos, incorrect grammar or any other editorial problem to <bugs@logreport.org>. We welcome reader's feedback. If you feel that certain parts of this manual aren't clear, are missing information or lacking in any other aspect, please tell us. Of course, if you feel like writing the missing information yourself, we'll very happily accept your patch. We will make our best effort to improve this manual.

Remember, that there is another manual, the *Lire Developer's Manual* which contains comprehensive information on how to extend Lire and describes in detail its internal architecture and design.

There are various mailing lists for Lire's users. There is a general users' discussion list where you can find help on how to install and use Lire. You can subscribe to this mailing list by sending an empty email with a subject of *subscribe* to <questions-request@logreport.org>. Email for the list should be sent to <questions@logreport.org>.

You can keep track of Lire's new release by subscribing to the announcement mailing list. You can subscribe yourself by sending an empty email with a subject of *subscribe* to <announcement-request@logreport.org>.

Finally, if you're interested in Lire's development, there is a development mailing list to which you can subscribe by sending an empty email with a subject of *subscribe* to <development-request@logreport.org>. Email to the list should be sent to <development@logreport.org>.

I. Lire Overview

Chapter 1. Introducing Lire

What Is Lire?

The Lire package is targeted at automatically generating useful reports from raw log files from various services. Currently, Lire can generate reports for a variety of email, web, dns, ftp, print servers and firewalls, and supports multiple output formats. Lire is developed by the Stichting LogReport Foundation, more information about the project can be found on <http://www.logreport.org/>.

Lire is built around the concept of a *superservice*. A superservice is a class of applications which share the same reports. Lire supports 6 superservices: dns, email, firewall, ftp, print and www. This means that log files for all supported email servers (*service* in Lire's parlance) will get similar reports. This is important for heterogeneous environments where you could have e.g. Sendmail and Postfix mail servers running. You will get similar reports which you can compare.

Lire can run in an online responder setup, as a client, as a cron driven system, or as a command line driven system. In an online responder setup, the Lire system receives emails containing log files from other hosts and sends generated reports back by email. In a client setup, the system sends log files by email to another Lire system which runs an online responder and receives reports back. Optionally, the log files can be anonymized before being sent. A cron driven setup reads and processes log files after they're rotated, on the local host. In a command line driven system, users run the Lire scripts on an ad-hoc basis.

Supported Systems

The package is reported to be useable on

- GNU/Linux (Debian GNU/Linux ("potato" and "woody"), Red Hat Linux (7.0, 7.1, 7.2, 7.3), Mandrake Linux (7.0 Air, 7.2 Odyssey))
- BSD (FreeBSD (4.1-STABLE, 4.5-PRERELEASE, 4.4-STABLE), OpenBSD (2.7, 2.8, 2.9, 3.2), Mac OS X v 10.1)
- Solaris (SunOS 5.6 and 5.7)
- HP-UX (11.11)
- AIX (Thanks Raymond Page)
- And yes, it even runs on GNU/Hurd!

The LogReport team tests Lire on various GNU/Linux distributions, as well as on OpenBSD before shipping. Don't worry if your system isn't listed here: it means we haven't had the opportunity to test Lire on your system, it does *not* mean Lire won't run on your system. If Perl runs on your system (which very likely is the case), Lire very likely will run on it too. However, please send us a note on your experiences. We're interested in Lire's portability.

Supported Applications

Lire can generate reports for a variety of dns, email, print, proxy, database, ftp and web servers as well as some firewalls. Here are the applications (services) supported in each superservice.

Database

Lire can generate reports from the log files of database servers:

- MySQL. <http://www.mysql.org/>

For these applications, you will get reports about the number of queries, the top users, the most used databases and more.

Dialup

Lire can generate reports from the log files of Linux kernel 2.4.x isdnlog log files:

- Linux kernel 2.4.x isdnlog <http://www.isdn4linux.de/>

DNS

Lire can generate reports from the query log files of two DNS servers:

- Bind 8. <http://www.isc.org/products/BIND/bind8.html>
- Bind 9. <http://www.isc.org/products/BIND/bind9.html>

For these applications, you will get reports about the number of DNS requests by hour, the top DNS clients, the most requested names and more.

DNS Zone

Lire can generated reports from DNS server logs about DNS Zone transfers: AXFR's and the loading of zones, as logged by e.g BIND 8's named log.

Email

Six email servers are supported by Lire:

- ArGoSoft Mail Server. <http://www.argosoft.com/applications/mailserver/> (<http://www.argosoft.org/>)
- Exim. <http://www.exim.org/>
- Postfix. <http://www.postfix.org/>
- Netscape Messaging Server.
- Qmail. <http://www.qmail.org/>
- Sendmail. <http://www.sendmail.org/>

The email servers' reports will show you the number of deliveries and the volume of email delivered by day, the domains from which you receive or send the most emails, the relays most used, etc.

Firewall

Several packet filtering firewalls are supported by Lire:

- Log files from Cisco IOS <http://www.cisco.com/univercd/cc/td/doc/product/software/> (<http://www.cisco.com/univercd/cc/td/doc/product/software/>).
- IPfilter log files <http://coombs.anu.edu.au/~avalon/ip-filter.html> (<http://coombs.anu.edu.au/~avalon/ip-filter.html>).
- Linux 2.2.X ipchains log files. <http://netfilter.samba.org/ipchains/> (<http://netfilter.samba.org/ipchains/>).
- Linux 2.4.X iptables log files. <http://netfilter.samba.org/> (<http://netfilter.samba.org/>).
- All log files using the WebTrends Enhanced Log Format (<http://www.webtrends.com/partners/welfOverview.htm>). This makes Lire support a potentially large number of firewall products. Consult <http://www.webtrends.com/partners/firewall.htm> for a list. Note that we didn't test Lire with all of those products. We appreciate all feedback regarding how Lire behaves with those products.

The reports generated will include informations about the IP address with the largest volume of data denied, the denied TCP ports, etc.

FTP

Lire can generate reports for FTP servers that use the xferlog log format. Some of the FTP servers known to support that log format:

- BSD ftpd. (As found on OpenBSD, FreeBSD and most UNIXes).
- ProFTPD. <http://www.proftpd.org/>
- Wu-Ftpd. <http://www.wu-ftp.org/>

It also supports log files from Microsoft Internet Information Server, which uses a variant of the W3C Extended Log Format.

The ftp superservice reports will include information such as the clients with the most transfers, the most requested files, the most active users, the amount of bytes transferred by day, etc.

Message Store

Lire can generate reports from log files from two message stores:

- Netscape Messaging Server.
- Netscape Messaging Server Mail Multi Plexor

Print

Lire can generate reports for two print servers:

- CUPS <http://www.cups.org/>
- LPRng <http://www.lprng.com/>

The reports generated will include information about the usage of the printers, statistics on the jobs and users.

Proxy

Lire supports three types of log files for proxy servers:

- Squid. <http://www.squid-cache.org/>
- Microsoft Internet Security and Acceleration Server. <http://www.microsoft.com/isaserver/>
- All log files using the WebTrends Enhanced Log Format (<http://www.webtrends.com/partners/welfOverview.htm>). This makes Lire support a potentially large number of proxy products. Consult <http://www.webtrends.com/partners/firewall.htm> for a list. Note that we didn't test Lire with all of those products. We appreciate all feedback regarding how Lire behaves with those products.

Syslog

Lire can generate overview reports about your syslog log files. It supports more than 8 different syslog log file formats.

WWW

Lire supports the three most common log formats for web servers: common log format (CLF), combined log format and the W3C extended log format (<http://www.w3.org/TR/WD-logfile.html>). Most web servers are able to log in one of those formats. It has been verified that Lire is able to generate reports for the following web servers:

- Apache. <http://httpd.apache.org/>
- Boa. <http://www.boa.org/>
- Microsoft Internet Information Server (3.X, 4.X, 5.X).
- iPlanet Web Server. <http://www.iplanet.com/>

Reports for the www superservice will include information like the number of requests by day, requests by browser, attack detection, top referers, etc. It is Lire's most complete report.

You will find the definite lists of reports available for each superservice in Part II: Reports Reference.

Supported Output Format

Lire supports multiple report output formats. All reports are generated in a native XML format which can be transformed into different other output formats. The following formats are supported:

ASCII

The default output format is ASCII. Simple text reports are best used for daily email reports.

HTML

Lire can generate HTML reports that can be viewed in any web browser. Those reports can include charts for easy overview.

PDF

To print the reports, Lire can generate Adobe PDF output. Like the HTML reports, those can include charts for easy overview.

Excel 95

FIXME

XHTML

Alternatively to the HTML output format, some may prefer reports in the new XML based XHTML. Those can be viewed in recent web browsers like Mozilla. Like most of the other reports, XHTML can include charts for easy overview.

DocBook XML

Lire can generate reports in the standard XML DTD DocBook. This intermediary format can be interesting for those who might want to customize the layouts of the reports. For example, one could generate reports in DocBook XML and use the company's stylesheets to print reports with the company's logo and standard report appearance.

What Lire Can't Do

Even with all the reports available, all those applications supported and all the possible output formats, there are still a number of things that Lire can't do by design. Lire is a *batch report generator*, it isn't a *real-time log analyzer*. There are a lot of real-time alerting tools out there. Lire is designed to generate reports from log files periodically (usually after the log files are rotated).

In case you find something you would like to see Lire do and it is reasonable that Lire should be able to do it, please let us know. In the Section called *If You Don't Find Something In This Manual* in *Preface* you can find how to get in contact with us.

Chapter 2. Installing Lire

Lire supports various installation environments. This chapter contains all there is to know about the installation of Lire in various setup scenarios: from the simple client setup to the installation of an online responder. You can find some quick installation instructions in the `INSTALL` file. Installation notes about specific platforms (Debian GNU/Linux, Solaris, OpenBSD, etc.) can be found in the *Lire FAQ*.

Client Installation

The simplest setup to install Lire in a client-server scenario is where the log files are sent by email to an online responder for processing.

Tip: You can test Lire by using Stichting LogReport Foundation's online responder available at `<log@service.logreport.org>`. (To process sendmail log files, send them to `<log@sendmail.logreport.org>`).

Requirements

To use Lire in such a setup, you only need a mailer (any will do) and an email address where the generated report can get sent to.

Installing

No special installation is necessary. You can generate reports by sending the log files to the responder right away. Consult the the Section called *Using A Responder* in Chapter 3 for the complete story.

Standalone Installation

The most common installation scenario will be where you install Lire on one system to generate daily or weekly reports from cron or by using the command line tools. This setup will install the complete software.

Requirements

Minimum Requirements

To install Lire on a system, you need the following:

- GNU gzip.
- Perl 5.6.1 or later (5.8.3 strongly recommended).
- The XML::Parser perl module. (This one needs the expat library.)

XML::Parser is available from any CPAN mirror. (<http://www.cpan.org/modules/by-module/XML/>).

The expat library is available from <http://expat.sourceforge.net/>.

- The DBD::SQLite perl module, available from <http://www.cpan.org/modules/by-module/DBD/> (which in turn requires DBI from <http://www.cpan.org/modules/by-module/DBI/>).
- The libintl-perl perl module, available from <http://www.cpan.org/modules/by-module/Locale/>.
- The Curses::UI perl module, available from <ftp://ftp.cpan.org/pub/CPAN/modules/by-module/Curses/>. Curses is required as well and is available from the same location.
- Standard UNIX utilities like **sh**, **ls**, **grep**, **bc**, **cut**, **head**, **sort**, **tar**, etc.

Those are the minimal requirements. With those, you will be able to generate ASCII reports only.

Requirements for Other Output Formats

To generate reports in other output formats than ASCII, there are additional requirements as follows:

- An XSLT processor. Currently the only supported XSLT processor is **xsltproc**, included with the XSLT C Library for Gnome (libxslt). You need version 1.0.4 or later. This XSLT C library can be used standalone and does not need the Gnome desktop environment to operate.

You can download this library along with the libxml2 library which it requires from <http://xmlsoft.org/XSLT/>.

When Lire calls **xsltproc**, you'll likely see something like "Attempt to load network entity <http://www.oasis-open.org/docbook/xml/4.1.2/docbookx.dtd>" in the Lire debug messages. This message is given by **xsltproc**, and is somewhat misleading. A more correct message would be: "A network URL should have been loaded but wasn't because you used xsltproc's --nonet option." Lire uses --nonet by default so that Lire doesn't hang behind firewalls. So: don't be scared by this message.

The method to render charts is through the use of **ploticus**. The **ploticus** generates nice looking graphs, especially in combination with PNG, PostScript and PDF output. As a standalone program it is quite easy to install (depending on the operating system you use).

The requirements to generate charts with **ploticus** are:

- The **ploticus** program, available from <http://ploticus.sourceforge.net/>). This package contains everything necessary to render GIF, SVG and PostScript images.

For additional output options, see below. The PNG format is recommended especially for web pages.

- To generate PNG (Portable Network Graphics) you will need the libpng and zlib libraries.

The libpng library is available from <http://www.libpng.org/> (<http://www.libpng.org/pub/png/libpng.html>) and <http://www.gzip.org/zlib> respectively.

- To generate JPEG images, you will also need the libgd and libjpeg libraries.

The GD Graphics Library is available from <http://www.boutell.com/gd/>.

The JPEG Library is available from <http://www.ijg.org>.

- To generate PDF vector image files (for embedding in PDF documents) and high quality PNG images, you will also need the GhostScript PostScript interpreter.

The GhostScript PostScript interpreter is available from <http://www.cs.wisc.edu/~ghost/>.

To generate HTML or XHTML reports, there are in addition to the XSLT processor, the following requirements:

- The DTD for DocBook XML 4.1.2. This is available from <http://www.docbook.org/xml/4.1.2/index.html>.
- Norman Walsh's XSL stylesheets for DocBook. You can download these stylesheets from <http://docbook.sourceforge.net/projects/xsl/index.html>.

To generate PDF reports, there are the following additional requirements:

- Jade, James Clarks' engine for the DSSSL style language. You can also use OpenJade, the name under which Jade is currently being maintained and extended.

You can download Jade from <http://www.jclark.com/jade/>. OpenJade is available from <http://openjade.sourceforge.net/>.

- The DTD for DocBook XML 4.1.2. This is available from <http://www.docbook.org/xml/4.1.2/index.html>.
- Norman Walsh's DSSSL stylesheets for DocBook. You can download those stylesheets from <http://docbook.sourceforge.net/projects/dsssl/index.html>
- (*For PDF output only*). JadeTeX and recent TeX installation.
JadeTeX is available from <http://jadetex.sourceforge.net/>.

Other Optional Requirements

Other optional things you may want to install:

- When available, the **logger** utility can be used to send Lire output to syslog.
- The Time-modules perl module (available from any CPAN mirror, <http://www.cpan.org/modules/by-module/Time/>. If it isn't present in the system, the required files included with Lire will be installed.
- The MIME-Tools perl module (available from any CPAN mirror, <http://www.cpan.org/modules/by-module/MIME/>.

This module is necessary to conveniently send reports by email or to operate a responder.

Installing

Installation of Lire is pretty straightforward:

1. Make sure that you have the requirements installed.
2. Extract the source code:

```
$ gzip -dc lire-version.tar.gz | tar xf -
```

3. Configure the software. You may use the `--prefix` option to specify where you want to install Lire. By default, it will be installed under `/usr/local`.

```
$ cd lire-version
$ ./configure [--prefix=path]
```

Make sure not to use `~` in the *path*. This is known to fail.

It find all requirements you had installed.

Note: As for SGML/XML components (DocBook DTD, Norman Walsh's DSSSL and XSL stylesheets), **configure** should find them if they were installed in "standard" places. (This is somewhere in an `sgml` tree as specified in the FHS and as you will find on most recent GNU/Linux distributions.) If they aren't found, you may hint **configure** by specifying their location through the use of environment variables:

```
$ DBK_XML_DTD=path_to_docbook_dir/docbookx.dtd \
  DBK_DSSSL_STYLESHEETS=path_to_dbk_dsssl_dir \
  DBK_XSL_STYLESHEETS=path_to_dbk_xsl_dir \
  ./configure [--prefix=path]
```

Similarly, you can use other environment variables to hint for other things that Lire can't find. See Chapter 18 for the complete list.

4. Compile the software (if you have XML::Parser installed, this will consist only of generating man pages).

```
$ make
```

5. You may have to become root if you are installing in a directory where only root has write permissions.
6. Install Lire.

```
# make install
```

That's it! You have a complete Lire installation and are ready to generate some reports. See Chapter 3 for information on using Lire.

Anonymized Client Installation

Although the client-only setup is the easiest to install and use, some people might understandably be worried about sending log files that may contain sensitive data to a public online responder. That is why Lire supports

anonymizing of log files. In an anonymized client setup, hostnames, emails and IP addresses in the log files are anonymized before being sent to the responder. The responder replies with a report in the Lire XML report format which is then de-anonymized by the client and transformed into the appropriate output format.

Requirements

The anonymized client installation has the same requirements as a standalone installation (see the Section called *Standalone Installation*). Like in the Standalone Installation, those will vary according to the output format you want to support.

Additionally, to support the anonymizing process, you will need Berkeley DB and the DB_File perl module. This module is part of the standard perl installation, but on proprietary UNIX systems you may have to install it separately.

Installing

There is no difference between the anonymized client installation and the Standalone Installation procedure. Consult the Section called *Standalone Installation*.

Responder Installation

When you want to generate reports for several servers, it is best to install Lire as a responder on one system to which to other systems can send their log files. This section describes how to setup Lire as a responder.

Requirements

Responder installation has the same requirements as the standalone installation (see the Section called *Standalone Installation*).

There is the following additional requirement:

- The MIME-Tools perl module (available from any CPAN mirror, <http://www.cpan.org/modules/by-module/MIME/>).

Installation

Basic installation procedure is the same as a standalone installation (see the Section called *Standalone Installation*). You might want to change the `--with-spooldir` option to **configure** (the default is `prefix/var/spool/lire`):

```
$ ./configure [--prefix=path --with-spooldir=path_to_spooldir]
```

Lire in a responder setup runs the **lr_spoold** daemon which scans maildirs where requests are delivered. Consequently, to finish the responder installation you have to create a maildir for each service you want to support and setup delivery to those maildirs.

Note: A *maildir* is a mailbox format first developed as part of Qmail where messages are stored in a directory hierarchy instead of a single file. You can find more informations about the maildir format at <http://www.courier-mta.org/maildirmake.html>.

As far as Lire is concerned, a maildir is a subdirectory *service/Maildir/new* which contains email messages in separate files.

The *sysconfdir/lire/address.cf* contains the name of the maildirs that are to be scanned and the type of log files that the emails should contain.

Refer to your MTA's documentation for notes on how to setup delivery to maildir. We give some notes on how to do this in the following sections.

The **lr_setup_responder** script can be used to setup some required infrastructure for the responder. Alternatively, one can execute the setup manually: One can create the maildirs by doing e.g.

```
$ cd ~/lire
$ mkdir -p var/spool/lire/common
$ maildirmake var/spool/lire/common/Maildir
$ cd ~/lire/var/spool/lire
$ mkdir bind8_query postfix qmail sendmail
$ maildirmake bind8_query/Maildir
$ maildirmake postfix/Maildir
$ maildirmake qmail/Maildir
$ maildirmake sendmail/Maildir
```

maildirmake gets distributed with qmail and with the Courier Mail Server <http://www.courier-mta.org>. If you haven't set up delivery to maildirs yet, doing a

```
$ maildirmake foo
```

is about the same as doing

```
$ mkdir foo
$ mkdir foo/cur foo/new foo/tmp
$ chmod og-rwx foo foo/*
```

Installing Under MTA's using procmail as their MDA

On many systems, procmail is used as the default Mail Delivery Agent. For instance, sendmail very often is configured to use procmail. If your MTA is configured like this, you can use procmail to take care of delivering to the right Maildir. We give some hints on how to get this done.

In Lire's *\$HOME/.procmailrc* you can put

```

:0:
* ^To:.*combined-log@
<LR_SPOOLDIR>/combined/Maildir/new

:0:
* ^To:.*sendmail-log@
<LR_SPOOLDIR>/sendmail/Maildir/new

```

etc. Make sure to replace `<LR_SPOOLDIR>` by the appropriate path. After that, you'll only have to make sure that the addresses `combined-log`, `sendmail-log`, etc. are aliases for the Lire user. You can then run `lr_spoold` to monitor the spool archives.

Installing Under Exim

There is more than one way to setup maildir delivery on a system running exim <http://www.exim.org/>. We show only one.

Be sure to have "maildir_format" enabled in the `address_directory:` section, e.g.

```

address_directory:
  driver = appendfile
  no_from_hack
  prefix = ""
  suffix = ""
  maildir_format

```

in your `exim.conf`'s transport configuration. Furthermore, have "directory_transport" transport in the `userforward` driver set to "address_directory", e.g.

```

userforward:
  driver = forwardfile
  file_transport = address_file
  pipe_transport = address_pipe
  reply_transport = address_reply
  directory_transport = address_directory
  no_verify
  check_ancestor
  check_local_user
  file = .forward
  modemask = 002
  filter

```

in your `exim.conf`'s directors configuration. Create a maildir, e.g. `~/.lire/var/spool/combined/`. (See the `qmail` section for how to do this.) Finally, do e.g.

```

$ cat <<EOT > .forward
> # Exim filter
> save \${home}/.lire/var/spool/combined/
> EOT

```

One could create more than one maildir, and configure the useraccount to store email messages for different services in different maildirs. We wont go into this such detail here though.

Installing Under qmail

Suppose your configure-time prefix was \$HOME/lire.

```
$ cd ~/lire/var/spool/lire/postfix
$ maildirmake Maildir
$ echo './lire/var/spool/lire/postfix/Maildir/' > .qmail-postfix
```

Get mail to postfix@yourhost delivered to hibou-postfix@yourhost, and controlled by ~hibou/.qmail-postfix:

```
$ su
# cd /var/qmail/control/users
# vi assign
=postfix:hibou:1028:1028:/home/hibou::-:postfix:
```

Get mail to anybody@postfix.yourdomain delivered to the local postfix mailbox:

```
# vi virtualdomains
postfix.yourdomain:postfix
```

Now send your qmail-send process a SIGHUP.

Chapter 3. Running Lire

This chapter describes the various ways that you can use Lire to process log files to generate reports. The next chapter (Chapter 4) explains how you can set up your system to process your log files automatically at regular interval.

Lire's configuration system

Lire's configuration is spread over several files. These files are pulled from different locations, each with a different purpose.

`prefix/share/lire/defaults/`

This configuration directory contains configuration files of Lire and any various third-party components you may have installed. You should not edit these files, modifications will be overwritten during an upgrade of Lire or those third-party components.

`~/.lire/config.xml`

This is the configuration file where your customized configuration is stored.

The configuration of a lire session is build up first from the files found in the defaults directory and the user configuration file in the end.

Getting Started

To start configuring Lire, simply fire up **lire** from a terminal and select the Lire—>Preferences menu. Initially all values will be unset so that only defaults are used. If you want to override a default all you need to do is set the value.

If a value is invalid for a given variable, a message will be displayed next to the widget holding the erroneous configuration variable. Such variables will be ignored when Lire runs and the defaults will be used instead.

Your customized configuration file will only be saved if you press the OK button. Otherwise the changes will be ignored, as you would normally expect.

Using A Responder

The easiest way to generate a report from your log file is to send your log file to a responder. The report will be sent to you by email to the address specified in the *Reply-To:* or *From:* header. To use a responder, you only need your standard mailer.

To save bandwidth, responders accept log files compressed using **gzip**, **compress** or **zip**. The log file can be sent in the email body or in a MIME attachment.

Note: Although any mailer will do, you should take care of the following when sending your log file:

- Make sure that your mailer won't insert new lines to wrap long log lines.
- Make sure that your mailer sets the standard MIME headers when using transfer encoding.

- When sending the log file as a MIME attachment, make sure that there are no other attachments (such as a signature) after the log file.

As a public service Stichting LogReport Foundation offers an online responder. To use it, you just send your log file to the appropriate responder for the log format you are using. The email addresses available can be found at <http://logreport.org/lire/or/> (<http://logreport.org/lire/or/>).

Example 3-1. Sending a Log File For Processing To A Responder

In this example, a bind8 query log file is sent to the LogReport responder for processing. The report will be sent back to the user who ran the **mail** command.

```
$ mail -s "Bind8 Log" log@bind8-query.logreport.org < \
/var/log/query.log
```

To save bandwidth, please send big log files in compressed format only. E.g., do:

```
$ mutt -s "`hostname` `date`" -a \
/var/log/apache/common.log.1.gz log@common.logreport.org < \
/dev/null
```

For more privacy, it is possible to send an anonymized log to the responder. Consult the Section called *Sending Anonymized Log Files To A Responder* for more information.

Generating A Report From A Log File

To generate a report from a log file, you use the **lr_log2report** command. This command expects a log file on standard input and will output the report on standard output. It takes the Lire service of the log file as its argument. The various supported services can be found by running **lr_check_service -l**, see the **lr_log2report(1)** man page for details. The **lr_log2report** command will log a lot of information on standard error. You can use the **lr_run** wrapper to filter those messages according to your preferences.

Example 3-2. Generating a Report With lr_log2report

This is the way to generate a report in the default output format for a log file taken from an Apache log server.

```
$ lr_run lr_log2report combined < \
/var/log/apache/access_log > ~/report.txt
```

Selecting Output Format

Another output format than the default one (usually text) can be selected by using the **-o** switch with the **lr_log2report** command.

Example 3-3. Generating A HTML Report

To generate a HTML report from the same log file as above, you would use the following command:

```
$ lr_run lr_log2report -o html_page combined < \
  /var/log/apache/access_log > ~/apache.tar
```

This will create a **tar** archive containing the report. Unpack it using the **tar** command, or unpack it on the fly, doing:

```
$ lr_run lr_log2report -o html_page combined < \
  /var/log/apache/access_log | tar xf -
```

This will create a `report` directory in the current working directory, holding the HTML report file.

Including Charts in the Report

Lire can generate charts (pie chart, bar chart, line graph or histogram) for some reports. If you have the necessary programs installed, you can tell **lr_log2report** to make the charts by adding the `-i` option to your command line (Only the PDF, HTML, XHTML and DocBook XML output formats support charts).

Example 3-4. Generating A HTML Report With Charts

To include charts with the HTML report, you would use the following command:

```
$ lr_run lr_log2report -o html -i combined < \
  /var/log/apache/access_log | tar xfc - /tmp
```

In this case, we piped the output to **tar** because what is outputted is a tar file. This tar file contains a directory called `report` which contains the charts in PNG format and the HTML report in the `index.html` file. The command creates a directory `/tmp/report/`.

Merging Reports

Since June 2002, Lire supports the merging of reports: one can combine two reports into one bigger report. This can be used to generate e.g. a weekly report from 7 daily reports, or generate a site-wide report from reports about the behaviour of each server on a site.

We describe some of the issues involved here. More in-depth information can be found in the **lr_xml_merge** manpage.

Manual Merging

We give an example.

Example 3-5. Merging Reports

To process two BIND v9 logfiles, and merge the reports, one would run:

```
$ zcat /var/log/named.2.gz | lr_run lr_log2report \
  -o xml bind9_query > $XMLDIR/20020622.xml
$ zcat /var/log/named.1.gz | lr_run lr_log2report \
  -o xml bind9_query > $XMLDIR/20020623.xml
$ lr_run lr_xml2report -U dns -S bind9_query \
  $XMLDIR/20020622.xml $XMLDIR/20020623.xml > \
  $ASCIIDIR/20020622-20020623.txt
```

Of course, refer to the **lr_log2report** and **lr_xml2report** manpages for more detailed information.

Automating Merging using cron: some ideas

Suppose you're running a daily cronjob which feeds your log to **lr_log2report** (possibly via **lr_cron**). If you'd like to generate weekly reports, you should first make sure `LR_ARCHIVE` is set to '1'. This way, each daily XML report will get saved in files like e.g. `$XMLDIR/hostname/20020602085048-20020603221309.xml.gz` where `XMLDIR` is `/var/lib/lire/data/report/xml/dns/bind9` (the raw log files, as well as the intermediate DLF files, will get saved too). Now you can add a weekly cronjob, which unzips the last 7 files from this directory to 7 temporary files, and runs:

```
lr_run lr_xml2mail -U yoursuperservice -S yourservice \
  uncompressed-xml-file-nr-1 uncompressed-xml-file-nr-2 ... \
  uncompressed-xml-file-nr-7 joe@example.com
```

Be sure to clean up the 7 uncompressed files afterwards. Refer to the **lr_xml2mail** manpage for more details.

Gotchas

The merging functionality is very powerful, and allows you to shoot yourself in the foot. We document some pitfalls.

Holes in your reporting period

When merging XML report files `xml.3` (2002-06-02 08:50:48 CEST - 2002-06-09 08:05:06 CEST) and `xml.1` (2002-06-16 08:18:40 CEST - 2002-06-21 22:13:09 CEST), the generated report will gladly display "Reporting on period: 2002-06-02 08:50:48 CEST - 2002-06-21 22:13:09 CEST": There is *no* safeguard against forgetting in-between report files.

Take care when changing report configuration file parameters

In some cases, changing the report configuration file (`superservice.cfg`) just before merging might lead to bogus data in your report.

Consider this case: Our `firewall.cfg` file looks like:

```
=section Denied Packets Reports
|select-action action_match="denied"
top-pkt-by-src          ips_to_show=10
```

We process some firewall logs, and archive the XML reports. Now we set `ips_to_show` to 1000, and merge the XML reports. This could incorrectly omit some IPs! You've got no guarantee the exact top 1000 IPs are shown. This is due to the fact the XML reports do *not* contain all information from the log: they're *reports*, after all.

Due to these issues, the merging is implemented with some heuristics: we keep more data than what's requested by the user in the XML report, to be able to handle most after-the-fact merging requests. We've tested the algorithm with a pretty broad range of real-life log files, and found out generally, the merged reports do give a good reflection of what actually has happened on the network: the heuristic is pretty well chosen. However, if you really need guaranteed 100% accurate data, generate your report directly from the raw logs. If you just want a quick overview, the merging is more suitable. Just make sure you're not cranking the limit parameters up too high in this case.

See also the Report Generating chapter in the Lire Architecture part of the *Lire Developer's Manual*.

Sending Anonymized Log Files To A Responder

For more privacy, you can anonymize your log somewhat before sending it to a responder. Lire includes a command called `lr_anonymize` which will transform everything that looks like an IP address, an email or a domain name into an anonymized form (10.0.0.1, 2.0.0.10.in-addr.arpa, 11.example.com, <john.doe@2.example.com>, etc.) The mapping between the real value and its anonymized form is saved in a disk database so that you can reverse the process when you receive the report from the responder.

The procedure is quite simple, you just have to filter your log file through `lr_anonymize` and make sure that the subject of your email starts with `anon`.

Example 3-6. Sending An Anonymized Postfix Log File To A Responder

To send an anonymized postfix log file to the Stichting LogReport Foundation responder, you would use a command like:

```
$ grep ' postfix/' /var/log/mail.log | \
  lr_run lr_anonymize /tmp/anon | \
  mail -s "anon Daily Report" log@postfix.logreport.org
```

The `/tmp/anon` is the database that is used to save the mapping between the real and anonymized values.

Warning

lr_anonymize will overwrite the content of that database, so if you reuse the database, make sure that you don't have two concurrent requests to a responder because you will lose the first mappings!

Processing The Responder's Results

The responder will generate a report in an XML format specific to Lire. To obtain a "normal" report from this, you first deanonymize it, then run the appropriate converter on the deanonymized report. The converter for a specific output format is called `lr_xml2format`. For example, you would use the `lr_xml2pdf` command to generate a PDF report.

Example 3-7. Deanonymizing and Generating A HTML Report

To generate a HTML report from the XML report you received from the responder, you would use the following command:

```
$ lr_run lr_deanonymize /tmp/anon < /tmp/anon-report.xml > /tmp/report.xml
$ lr_run lr_xml2report -o html /tmp/report.xml > /tmp/report.html
```

You could also generate charts by adding the `-i` to the **lr_xml2report** command.

Running Lire In A Server Cluster

Using Mail

You can monitor a set of maildirs which receive email messages containing log files for the services as listed in `address.cf` by doing something like:

```
$ lr_run lr_spoold
```

This enables you to configure one host as a reporting host (or "online responder"), while other machines send their log files to it by email for processing. (If remote syslogging is used, a cron-driven setup is sufficient.)

A publicly available online responder is running at `log@<servicename>.logreport.org`; see <http://logreport.org/lire/or/> for more information.

Chapter 4. Automating Lire

This chapter discusses various ways to configure Lire for generating periodical reports from your system logs.

Automatically Processing Log Files Using Cron

The easiest way to have Lire generate reports from the various log files available on your system is through a **cron** job. Lire includes a script called **lr_cron** which takes care of calling the appropriate batch of commands on the appropriate log files. The reports can be generated either in the ASCII format or in any other format, as long as they are supported by Lire and the required software packages are available on your system. During the **lr_cron** process, they are afterward sent to an email address of your choice or stored in a location of your filesystem. To use **lr_cron**, you have to configure you DLF stores using **lire**.

Configuring lr_cron

An **lr_cron** process has to be configured in two steps:

- configuring the jobs related to a store
- configuring cronjobs to run **lr_cron** on the specified stores at the specified frequency

The first step is achieved by executing the **lire** command. Open or create a DLF store using the **Store→Open...** or the **Store→New...** menu. You will then see a list of the *Import jobs* and the *Report jobs* configured in this store. Edit them suiting your needs and remember the period parameters you have indicated for each of those jobs.

The purpose of the *Import jobs* is to schedule imports of log files within the DLF store while *Report jobs* are the processes through which the gathered data are computed and rendered into a report, which in turn can parametrized and scheduled for sending.

Each time **lr_cron** runs on a store, it executes the Import jobs and the Report jobs in order, depending of the period of their schedule.

Using lr_cron within Cron

Installing cron jobs is really easy since the only parameters given to **lr_cron** are a *period* and a *store*. The lines to add to your crontab should look similar to:

```
0 0 * * * /usr/bin/lr_cron daily /var/lib/lire/www_store
0 0 * * 0 /usr/bin/lr_cron weekly /var/lib/lire/email_store
```

Once activated like this, report(s) will be sent on a weekly and/or daily basis.

Chapter 5. Customizing Lire's Reports

This chapter explains how to customize the reports generated by Lire. Each superservice comes with a number of subreports which select various information from the log file for inclusion in the final report.

The report's configuration for a specific superservice is in a file named *superservice.cfg*. This file is looked for in `$sysconfdir/lire` and `$HOME/.lire/etc/`. If you're not happy with the default configuration as shipped with Lire, you're advised to copy `$sysconfdir/lire/superservice.cfg` to `$HOME/.lire/etc/` and edit this copy. This copy will of course not get touched on Lire upgrades.

The Report's Configuration File

The configuration file is *line oriented*. Empty lines are ignored as well as line starting with `#`. The configuration file is divided in sections introduced by lines starting with the directive `=section`. The section's title follows the directive.

Each section can optionally contain a list of filters that will be used to filter the records used to generate the subreports. A filter is set up by starting a line with the pipe (`|`) character followed by the filter's id. The rest of the line will usually contain filter's parameter assignments. If any filters is used in the section, they must come after the section's title and before the subreports.

Other lines are interpreted as subreport id's.

Example 5-1. Commented Report Configuration File

Here is an example configuration for the DNS superservice.

```
=section All Requests
top-requesting-hosts             hosts_to_show=10
top-requested-names              names_to_show=10
requesttype-distribution
requests-by-period              period=1d

=section Recursive Requests
|select-resolver    method="recurs"
top-requesting-hosts             hosts_to_show=10
top-requested-names              names_to_show=10
requesttype-distribution
requests-by-period              period=1d

=section Non Recursive Requests
|select-resolver    method="nonrec"
top-requesting-hosts             hosts_to_show=10
top-requested-names              names_to_show=10
requesttype-distribution
requests-by-period              period=1d
```

This DNS report will contain three sections. The first section "All Requests" doesn't use any filters and thus the four configured subreports (`top-requesting-hosts`, `top-requested-names`, `requesttype-distribution` and `requests-by-period`) will be computed using all the requests. The second section ("Recursive Requests") has one filter set up (`select-resolver`) which will select only recursive requests. This section will contain the same subreports as the "All Requests" section, but calculated on a different input set. The third section ("Non Recursive Requests") is similar to the second section with the exception that only non-recursive requests will be used to compute the subreports.

Selecting Subreports

Example 5-2. FTP Report Configuration File

Here is an example configuration file for the FTP superservice.

```
# Report configuration for the FTP super service

# Top X reports
top-remote-host hosts_to_show=10
#top-files files_to_show=10
top-files-in files_to_show=10
top-files-out files_to_show=10
top-users users_to_show=10

# By day reports
bytes-by-period period="1d"

# Transfers by X reports
transfers-by-direction
transfers-by-type
```

The FTP superservice will thus contain seven subreports. Because the line with the `top-files` subreport starts with `#`, this subreport won't be included in the report. To include that subreport in the report, you would have to remove the `#` character.

In Part II, you will find all the subreports and filters available for all superservices.

Reordering The Subreports

Ordering is very simple. The order in which subreport lines appear in the config files is the order in which the subreports will be given in the output. Rearranging the lines in these configuration files reorders them in the output. For example, in the above example, `transfers-by-type` will be the last report given in the output. You can reorder sections in a similar manner.

Changing Parameters

Many subreports (and filters) can be customized through the configuration files. For example, consider this excerpt from the DNS configuration file previously listed.

```
top-requesting-hosts             hosts_to_show=10
top-requested-names              names_to_show=10
requesttype-distribution
requests-by-period              period=1d
```

All reports are selected, but furthermore, for the subreports giving a "Top X" output the number X can be defined. With the above configuration the report `top-requesting-hosts` will give a Top 10. Specifying '0' as the value for a top-report will cause all data to get displayed.

Caution

All variable settings must be placed on the same line as the subreport's id or filter's id!

A more exotic example is taken from the WWW superservice configuration file:

```
top-referers-by-page referer_to_show=5 page_to_show=10 referer_exclusion='^-$'
```

In this example a Perl regular expression is used as content for the `referer_exclusion` variable. This expression matches all referers -. Such referers are found in the log file in cases when e.g. the URL of your web page was typed by the client user. (When users visit your page by clicking on a link in a page, referring to your page, the page containing the link will be mentioned in the referer field.) All referers that match - will be excluded from the analysis.

In Part II, you will find the description of all the available subreports along with their parameters.

Using Subreports On Filtered Input

Some subreport contains support for filtering their input. One such subreport is the `top-referers-by-page` subreport. But sometimes a subreport doesn't support filtering its input, or you want to have exactly the same information as other subreports already included in the report but on a subset of the records. The solution to this problem is to set up filters in a specific section. You can find an example of this in the DNS report's configuration file that was commented in the Section called *The Report's Configuration File*.

We give two more examples of filtering usage here. A `firewall.cfg` file could read

```
=section Permitted Traffic Reports, 100.56.0.x
|select-action action_match="permitted"
|select-to-ip ip-range="^100\\.56\\.0\\."
top-pkt-by-src          ips_to_show=30
top-pkt-by-dst          ips_to_show=30
```

A `www.cfg` file could read

```
=section Visitors Reports, excluding 10.0.0.2
|exclude-client_host client_match="^10\\.0\\.0\\.2$"
top-client_host          client_to_show=10
top-client_host-by-size  client_to_show=10
```

In Part II, you will find the description of all the available filters along with their parameters.

II. Reports Reference

Chapter 6. Database Reports

Supported Log Format

Lire currently only supports the query log of MySQL. This log file contains all the connectinons and queries sent to your database server.

MySQL's Log

The MySQL's log file will contain information about each start and shutdown of your database server, as well as all connections and queries processed by the database server during its session.

Example 6-1. Sample MySQL Log File

```
/usr/sbin/mysqld, Version: 3.23.43-debug-log, started with:
Tcp port: 3306  Unix socket: /var/run/mysqld/mysqld.sock
Time           Id Command      Argument
011226 21:32:57      1 Connect    root@localhost on
011226 21:33:01      1 Query      show tables
011226 21:33:08      1 Query      show databases
011226 21:33:46      1 Quit
011226 21:34:32      2 Connect    Access denied for user: \
'jdoe@localhost' (Using password: YES)
011226 21:34:42      3 Connect    Access denied for user: \
'jdoe@localhost' (Using password: YES)
011226 21:35:59      6 Connect    jdoe@localhost on
                6 Init DB    nmrshiftdb
                6 Query      SHOW VARIABLES
011226 21:36:00      6 Query      CREATE TABLE molecules \
(molid INT, CMLcode TEXT)
                6 Query      CREATE TABLE chemnames \
(molid INT, autonom TEXT, name TEXT)
```

Report Descriptions and Configuration

Actions By Period Database Report

ID: actions-by-period

Chart: None

This report shows the number of actions in configurable time periods.

Parameters

period

This parameter controls the time period over which the deliveries are aggregated.

Defaults to 1d.

Most Active Users Database Report

ID: top-users

Chart: bars

This report lists the users that do the most actions.

Parameters

users_to_show

This parameter controls the number of users to display in the report.

Defaults to 10.

Most Accessed Databases Database Report

ID: top-databases

Chart: bars

This report lists the databases that were the most accessed.

Parameters

databases_to_show

This parameter controls the number of databases to display in the report.

Defaults to 10.

Queries By Type

ID: top-querytypes

Chart: bars

This report shows the number of queries by query type.

This report doesn't have any parameters.

Filter Descriptions and Configuration

None.

Chapter 7. Dialup Reports

Supported Log Format

Lire supports logs of one dialup connections: Linux kernel 2.4.x isdnlog.

Report Descriptions and Configuration

Connections By Connection Time Dialup Report

ID: `connections-by-connection-time`

Chart: None

This report shows the number of connections aggregated in configurable connection time range.

Parameters

size

This parameter controls the size of the connection time ranges over which the connections are aggregated.

Defaults to 1m.

Connections By Period Dialup Report

ID: `connections-by-period`

Chart: histogram

This report shows the number of connections aggregated in configurable time periods. It does only include succesful connections.

Parameters

period

This parameter controls the time period over which the connections are aggregated.

Defaults to 1d.

Connections By Connection Type Dialup Report

ID: `connections-by-type`

Chart: histogram

This report shows the number of connections aggregated by connection type (data or speech).

This report doesn't have any parameters.

Connections By Connect Status Dialup Report

ID: `connections-by-connect-status`

Chart: histogram

This report shows the number of connections aggregated by connect status (ring, busy, connected or failed).

This report doesn't have any parameters.

Connections By Hangup Status Dialup Report

ID: `connections-by-hangup-status`

Chart: histogram

This report shows the number of connections aggregated by hangup status (no answer, normal, unallocated).

This report doesn't have any parameters.

Connection Time By Type By Period Dialup Report

ID: `connection-time-by-period`

Chart: bars

This report lists the total connection time by period.

Parameters

period

This parameter controls the time period over which the connection time is aggregated.

Defaults to 1d.

Most Connections From Telephone Number Dialup Report

ID: `top-telephone-number-from`

Chart: bars

This report lists the telephone numbers from which the most connections were made.

Parameters

numbers_to_show

This parameter controls the number of telephone numbers to display in the report.

Defaults to 10.

Most Connections To Telephone Number Dialup Report

ID: top-telephone-number

Chart: bars

This report lists the telephone numbers to which to most connections were made.

Parameters

numbers_to_show

This parameter controls the number of telephone numbers the display in the report.

Defaults to 10.

Cost By Period Dialup Report

ID: cost-by-period

Chart: bars

This report lists the total cost by period.

Parameters

period

This parameter controls the time period over which the total cost is aggregated.

Defaults to 1d.

Cost By Telephone Number Dialup Report

ID: cost-by-telephone-number

Chart: bars

This report lists the total cost for each telephone number which has been called to.

This report doesn't have any parameters.

Chapter 8. DNS Reports

Supported Log Format

Lire supports query logs of two DNS servers: Bind 8 and Bind 9.

Note: You have to enable query logging in bind, something which is not turned on by default.

Example 8-1. Enabling Query Log In Bind

To enable query logging in Bind 8 or Bind 9, you should add the following to your `named.conf` configuration file:

```
logging {
    channel query_logging {
        file "/var/log/named_querylog"
        versions 3 size 100M;
        print-time yes;                // timestamp log entries
    };

    category queries {
        query_logging;
    };
};
```

Bind8 Query Log

Bind 8's query logs contain one entry for each DNS query made to the name server. It logs the time of the query (you have to set `print-time` to `yes` for this), the IP of the requesting client, the name queried, the type of the query and the protocol. Recursive queries will have a + after the XX which appears in all query entries.

Example 8-2. Sample Bind 8 Query Log

```
10-Apr-2000 00:01:20.307 XX /10.2.3.4/1.2.3.in-addr.arpa/SOA/IN
10-Apr-2000 00:01:20.308 XX+/10.4.3.2/host.foo.com/A/IN
```

Bind9 Query Log

Bind 9 logs the same information as Bind 8 (except whether the request was recursive or not) but in a different format.

Note: We also support the new date format introduced in Bind9 9.3 which also contains the year (15-Jul-2002).

Example 8-3. Sample Bind 9 Query Log

print-severity and *print-category* were set to *yes* to obtain that log. Lire also accepts logs where those are turned off.

```
Feb 25 11:09:43.651 queries: info: client 10.0.0.3#1035: \
    query: 3.example.com.nl IN A
Feb 25 11:09:48.739 queries: info: client 10.0.0.3#1035: \
    query: 3.example.com.nl IN A
Feb 25 12:50:32.476 queries: info: client 10.0.0.3#1035: \
    query: 21.example.com.co.uk IN A
Feb 25 12:50:34.110 queries: info: client 10.0.0.3#1035: \
    query: 22.example.com IN A
```

Tip: If you miss the recursive flag from Bind 8, it is possible to add back that feature by patching Bind 9. The following patch by Wytze van der Raay will add a + or - after the query type to indicate whether the query was recursive or not. Lire will detect that the log file was made by a patched Bind 9.

```
# patch bin/named/query.c to log recursive/non-recursive query indication
SRC=bin/named/query.c
if [ -f ${SRC}.org ]
then
    echo "Patched ${SRC} already in place"
else
    echo "Patch ${SRC} for recursive/non-recursive query indication"
    cp -p ${SRC} ${SRC}.org
    patch -p0 ${SRC} <<\!
--- bin/named/query.c.org      Mon Sep 24 22:57:48 2001
+++ bin/named/query.c        Tue Sep 25 09:55:21 2001
@@ -3272,7 +3272,8 @@
    dns_rdatatype_format(rdataset->type, typename, sizeof(typename));

    ns_client_log(client, NS_LOGCATEGORY_QUERIES, NS_LOGMODULE_QUERY,
-               level, "query: %s %s %s", namebuf, classname, typename);
+               level, "query: %s %s %s%s", namebuf, classname, typename,
+               WANTRECURSION(client) ? "+" : "-");
}

void
!
fi
```

Report Descriptions and Configuration

Top Requesting Hosts Report

ID: top-clients-detailed-report

Chart: None

This report lists the requesting hosts with the most requests and it displays more detailed information on those requests: it shows the breakdown of requested names and request types.

Parameters

hosts_to_show

This parameter controls the number of hosts to display in the report.

Defaults to 10.

names_to_show

This parameter controls the number of requested names to display by host in the report.

Defaults to 5.

Top Requesting Hosts Report

ID: top-requesting-hosts

Chart: bars

This report lists the requesting hosts with the most requests.

Parameters

hosts_to_show

This parameter controls the number of hosts to display in the report.

Defaults to 10.

Top Requested Names Report

ID: top-requested-names

Chart: bars

This report lists the most requested names.

Parameters

names_to_show

This parameter controls the number of names to display in the report.

Defaults to 10.

Distribution of Request Types by Method DNS Report

ID: requesttype-by-method

Chart: None

This report shows the distribution of the type of requests splitted by method.

This report doesn't have any parameters.

Distribution of Request Types Report

ID: requesttype-distribution

Chart: bars

This reports on the distribution of the request types.

This report doesn't have any parameters.

Requests By Period DNS Report

ID: requests-by-period

Chart: histogram

This report shows the number of requests aggregated in configurable time periods.

Parameters

period

This parameter controls the time period over which the requests are aggregated.

Defaults to 1d.

Requests By Timeslot DNS Report

ID: requests-by-timeslot

Chart: histogram

This report shows the number of requests aggregated in configurable time slots (by hours of the day, by days of the weeks, etc.).

Parameters

timeslot

This parameter controls the time slot over which the requests are aggregated.

Defaults to 1h.

Requests by Period by Method DNS Report

ID: req-by-period-by-method

Chart: None

This report shows the number of requests aggregated in configurable time periods and splitted by the resolver's type.

Parameters

period

This parameter controls the time period over which the requests are aggregated.

Defaults to 1d.

Requests by Timeslot by Method DNS Report

ID: req-by-timeslot-by-method

Chart: None

This report shows the number of requests aggregated in configurable time slots (by hours of the day, by days of the weeks, etc.) and splitted by the resolver's type.

Parameters

timeslot

This parameter controls the time slot over which the requests are aggregated.

Defaults to 1h.

Filter Descriptions and Configuration

Select Resolver Filter

ID: select-resolver

This filter specification can be used to select only DLF records having a particular resolver type.

Parameters

method

This parameter sets the type of resolver that should be selected. Possible methods are:

nonrec

Non-recursive requests.

recurs

Recursive requests.

Defaults to ^recurs\$.

Chapter 9. DNS Zone Reports

Supported Log Format

Lire supports named log files from BIND 8.

Report Descriptions and Configuration

Top Requesting AXFR Hosts Report

ID: `top-axfr-host`

Chart: bars

This report lists the requesting hosts with the most requests.

Parameters

hosts_to_show

This parameter controls the number of hosts to display in the report.

Defaults to 10.

Top Requesting AXFR Hosts Report

ID: `top-axfr-zone`

Chart: bars

This report lists the requesting hosts with the most requests.

Parameters

zones_to_show

This parameter controls the number of hosts to display in the report.

Defaults to 10.

Top Requesting AXFR Hosts Report

ID: `top-axfr`

Chart: bars

This report lists the requesting hosts with the most requests.

Parameters

hosts_to_show

This parameter controls the number of hosts to display in the report.

Defaults to 10.

zones_to_show

This parameter controls the number of zones to display in the report.

Defaults to 10.

Top Denied AXFR Requests

ID: top-denied-axfr

Chart: bars

This report lists the top 10 of denied AXFR requests.

Parameters

hosts_to_show

This parameter controls the number of hosts to display in the report.

Defaults to 10.

zones_to_show

This parameter controls the number of zones to display in the report.

Defaults to 10.

Top Denied Dynamic DNS Updates Report

ID: top-dynamicdenied

Chart: bars

This report lists the top of the denied dynamic DNS packets.

Parameters

zones_to_show

This parameter controls the number of zones to display in the report.

Defaults to 10.

`hosts_to_show`

This parameter controls the number of hosts to display in the report.

Defaults to 10.

Top Newly Loaded Zones Report

ID: `top-loadedzone`

Chart: `bars`

This report lists the top of the newly loaded zones.

Parameters

`zones_to_show`

This parameter controls the number of zones to display in the report.

Defaults to 10.

Filter Descriptions and Configuration

AXFR denied only

ID: `axfr-approved-only`

This filter specification filters out the non-approved AXFR lines.

This filter doesn't have any parameters.

AXFR denied only

ID: `axfr-axfr-only`

This filter specification filters out the transfered AXFR lines.

This filter doesn't have any parameters.

AXFR denied only

ID: `axfr-denied-only`

This filter specification filters out the non-approved AXFR lines.

This filter doesn't have any parameters.

Loaded zones only

ID: `dyndns-only`

This filter specification filters out the non-dynamic DNS entries.

This filter doesn't have any parameters.

Loaded zones only

ID: `loaded-only`

This filter specification filters out the non-loaded zones.

This filter doesn't have any parameters.

Chapter 10. Email Reports

Supported Log Format

Lire supports log files from six different email servers.

ArGoSoft Mail Server

The log files generated by the ArGoSoft Mail Server are supported. For proper operation, you'll need to turn on the following components' logging:

- Log SMTP commands.
- Log SMTP conversations.
- Log to File.

Example 10-1. ArGoSoft Mail Server Log Sample

```
3/17/2002 12:00:03 AM - SMTP connection with 10.0.0.1 [1.example.com] \
ended. ID=3342
3/17/2002 12:00:22 AM - Requested SMTP connection from 10.0.0.2 \
[2.example.com]
3/17/2002 12:00:22 AM - ( 3345) 220 ArGoSoft Mail Server Pro \
for WinNT/2000/XP, Version 1.8 (10.0.0.3)
3/17/2002 12:00:23 AM - ( 3345) HELO greed
3/17/2002 12:00:23 AM - ( 3345) 250 Welcome, 2.example.com \
[10.0.0.2], pleased to meet you
3/17/2002 12:00:23 AM - ( 3345) RSET
3/17/2002 12:00:23 AM - ( 3345) 250 Reset state
3/17/2002 12:00:23 AM - ( 3345) MAIL FROM:<john.doe.1@1.mail.example.com>
3/17/2002 12:00:23 AM - ( 3345) Checking address \
john.doe.1@1.mail.example.com
3/17/2002 12:00:23 AM - ( 3345) Address john.doe.1@1.mail.example.com \
is local
```

Exim

The standard log file from Exim is supported.

Example 10-2. Exim Log Sample

```
2001-03-27 10:00:11 exim 3.16 daemon started: pid=215, -q30m, \
listening for SMTP on port 25
2001-03-27 10:00:11 Start queue run: pid=218
```

```

2001-03-27 10:00:11 End queue run: pid=218
2001-03-27 10:08:01 Start queue run: pid=736
2001-03-27 10:08:01 End queue run: pid=736
2001-03-27 11:29:10 14hpmo-00002f-00 <= john.doe.25@1.mail.example.com \
    U=root P=local S=757
2001-03-27 11:29:11 14hpmo-00002f-00 => egonw \
    <john.doe.21@1.mail.example.com> D=localuser T=local_delivery
2001-03-27 11:29:11 14hpmo-00002f-00 Completed

```

Netscape Messaging Server

Netscape Messaging Server logs its information with **syslog**. No special configuration is necessary.

Example 10-3. Netscape Messaging Server Log Sample

```

[08/Jan/2002:11:30:00 +0100] rodolf smtpd[29296]: \
    General Information: Log created (1010485800)
[08/Jan/2002:11:30:00 +0100] rodolf smtpd[29296]: \
    General Notice: SMTP-Accept:GPM7U000.J7C:\
    <john.doe.1@1.mail.example.com>:[10.0.0.1]:1.example.com.fr:\
    <john.doe.2@1.mail.example.com>:4111:1:<john.doe.3@2.mail.example.com>
[08/Jan/2002:11:30:39 +0100] rodolf smtpd[29296]: \
    General Notice: SMTP-Accept:GPM7V300.A7C:\
    <john.doe.4@1.mail.example.com>:[10.0.0.1]:1.example.com.fr:\
    <john.doe.5@1.mail.example.com>:59347:1:<john.doe.6@2.mail.example.com>
[08/Jan/2002:11:31:09 +0100] rodolf smtpd[29296]: \
    General Notice: SMTP-Accept:GPM7VX00.67E:\
    <john.doe.7@3.mail.example.com>:[10.0.0.1]:1.example.com.fr:\
    <john.doe.8@4.mail.example.com>:4117:1:<john.doe.9@2.mail.example.com>
[08/Jan/2002:11:31:26 +0100] rodolf smtpd[29296]: \
    General Notice: SMTP-Accept:GPM7WE00.D7U:\
    <john.doe.10@5.mail.example.com> (added by 2.example.com.fr):\
    [10.0.0.1]:1.example.com.fr:<john.doe.11@6.mail.example.com>:3278:1:\
    <john.doe.12@2.mail.example.com>
[08/Jan/2002:11:31:33 +0100] rodolf smtpd[29296]: \
    General Notice: SMTP-Accept:GPM7WL00.F86:
    <john.doe.13@7.mail.example.com>:[10.0.0.1]:1.example.com.fr:\
    <john.doe.14@1.mail.example.com>:998:1:<john.doe.15@2.mail.example.com>

```

Postfix

Postfix logs its information with **syslog**. No special configuration is necessary.

Example 10-4. Postfix Log Sample

```
Dec 1 04:02:56 internetsrv postfix/pickup[20919]: 693A3578E: uid=0 from=<root>
Dec 1 04:02:56 internetsrv postfix/cleanup[20921]: 693A3578E: \
    message-id=<john.doe.1@example.com>
Dec 1 04:02:57 internetsrv postfix/qmgr[20164]: 693A3578E: \
    from=<john.doe.2@example.com>, size=617 (queue active)
Dec 1 04:02:57 internetsrv postfix/cleanup[20921]: E325C578D: \
    message-id=<john.doe.1@example.com>
Dec 1 04:02:58 internetsrv postfix/local[20924]: 693A3578E: \
    to=<john.doe.2@example.com>, relay=local, delay=3, \
    status=sent (forwarded as E325C578D)
Dec 1 04:02:58 internetsrv postfix/qmgr[20164]: E325C578D: \
    from=<john.doe.2@example.com>, size=769 (queue active)
```

Qmail

Lire accepts qmail-send Qmail log files where each line starts with the timestamp in numerical (with fraction) format: 982584201.511524. qmail-smtpd logfiles are not (yet) supported.

Tip: If you use **multilog**, you will have to filter your log file through **tai64nfrac**.

Tip: If you redirect your Qmail logs to **syslog**, you can run **lr_desyslog** (included in Lire) to remove the extra **syslog** timestamp:

```
$ lr_desyslog qmail < qmail-syslog.log > qmail.log
```

Example 10-5. Qmail Log Sample

```
998545829.342079 new msg 6416
998545829.342350 info msg 6416: bytes 2657 from \
    <bounce-debian-hurd=john.doe-debian-hurd=john.doe.1@1.mail.example.com> \
    qp 22423 uid 71
998545829.356889 starting delivery 1808: msg 6416 to local \
    john.doe.2@2.mail.example.com
998545829.357096 status: local 1/10 remote 0/20
998545829.445754 delivery 1808: success: did_0+0+1/
998545829.445976 status: local 0/10 remote 0/20
998545829.446056 end msg 6416
998545832.186954 new msg 6416
998545832.187213 info msg 6416: bytes 1957 from \
    <dns-return-13543-john-dns=john.doe.3@3.mail.example.com> qp 22431 uid 71
998545832.196806 starting delivery 1809: msg 6416 to local \
    john.doe.4@2.mail.example.com
```

Sendmail

Sendmail logs its activity through **syslog**. You need to set your *LogLevel* to 9 or higher. Versions 8.10.x and 8.11.x of Sendmail are supported.

Example 10-6. Sendmail Log Sample

```
Oct 29 14:46:13 mailhost sendmail[19504]: alias database /etc/aliases \
rebuilt by root
Oct 29 14:46:13 mailhost sendmail[19504]: /etc/aliases: 40 aliases, \
longest 10 bytes, 395 bytes total
Oct 29 14:52:33 mailhost sendmail[19584]: alias database /etc/aliases \
rebuilt by root
Oct 29 14:52:33 mailhost sendmail[19584]: /etc/aliases: 40 aliases, \
longest 10 bytes, 395 bytes total
Oct 29 15:00:00 mailhost sendmail[19633]: f9U000Y19633: from=root, \
size=257, class=0, nrcpts=1, msgid=<john.doe.1@1.mail.example.com>, \
relay=john.doe.2@2.mail.example.com
Oct 29 15:00:00 mailhost sendmail[19633]: f9U000Y19633: to=root, \
ctladdr=root (0/0), delay=00:00:00, xdelay=00:00:00, mailer=local, \
pri=30257, dsn=2.0.0, stat=Sent
Oct 29 16:00:00 mailhost sendmail[19672]: f9U100619672: from=root, size=257, \
class=0, nrcpts=1, msgid=<john.doe.3@1.mail.example.com>, \
relay=john.doe.2@2.mail.example.com
Oct 29 16:00:00 mailhost sendmail[19672]: f9U100619672: to=root, \
ctladdr=root (0/0), delay=00:00:00, xdelay=00:00:00, mailer=local, \
pri=30257, dsn=2.0.0, stat=Sent
Oct 29 17:00:00 mailhost sendmail[19696]: f9U200V19696: from=root, \
size=257, class=0, nrcpts=1, msgid=<john.doe.4@1.mail.example.com>, \
relay=john.doe.2@2.mail.example.com
Oct 29 17:00:00 mailhost sendmail[19696]: f9U200V19696: to=root, \
ctladdr=root (0/0), delay=00:00:00, xdelay=00:00:00, mailer=local, \
pri=30257, dsn=2.0.0, stat=Sent
```

Reports' Descriptions and Configuration

Deliveries Attempts By Period By Status Email Report

ID: deliveries-by-period-by-status

Chart: None

This report shows the number of delivery attempts that resulted in a specific status, aggregated in configurable time periods.

Parameters

period

This parameter controls the time period over which the deliveries are aggregated.

Defaults to 1d.

Deliveries Attempts By Period Email Report

ID: deliveries-by-period

Chart: histogram

This report shows the number of delivery attempts aggregated in configurable time periods.

Parameters

period

This parameter controls the time period over which the deliveries are aggregated.

Defaults to 1d.

Deliveries Attempts By Delay Email Report

ID: deliveries-by-delay

Chart: None

This report shows the number of deliveries attempts aggregated in configurable delay range.

Parameters

delay-size

This parameter controls the size of the delay ranges over which the deliveries are aggregated.

Defaults to 1s.

Deliveries Attempts By Size Email Report

ID: deliveries-by-size

Chart: None

This report shows the number of deliveries attempts aggregated in configurable size range.

Parameters

size

This parameter controls the size of the size ranges over which the deliveries are aggregated.

Defaults to 1k.

Failed Deliveries By Relay Email Report

ID: errors-by-to-relay

Chart: None

This report shows the errors that happened for each relay that was contacted.

This report doesn't have any parameters.

Different From Domain by Period Email Report

ID: from_domain-by-period

Chart: None

This report shows the number of different email domains that used your email server in an given period.

Parameters

period

This parameter controls the time period which will be used to compute different domains. For example, if you use 1d here, you'll get the number of different source email addresses' domain that sended email through your server in one day.

Defaults to 1d.

Different From Email for Selected Domains by Period Email Report

ID: from_user-from-domain-by-period

Chart: None

This report shows the number of different local parts from some domains that sended at least an email through the mail system in a given period.

Parameters

period

This parameter controls the time period which will be used to compute different user. For example, if you use 1d here, you'll get the number of different email addresses' from the selected domains that sended email through your server in one day.

Defaults to 1d.

select_domain

This parameter contains a regexp that will be used to select the domains for which different from addresses will be computed.

Defaults to .*.

Different To Domain by Period Email Report

ID: to_domain-by-period

Chart: None

This report shows the number of different email domains to which the server sent an email in an given period.

Parameters

period

This parameter controls the time period which will be used to compute different domains. For example, if you use 1d here, you'll get the number of different destination email addresses' domain that received an email through your server in one day.

Defaults to 1d.

Different To Email for Selected Domains by Period Email Report

ID: to_user-from-domain-by-period

Chart: None

This report shows the number of different local parts from some domains that received at least an email through the mail system in a given period.

Parameters

period

This parameter controls the time period which will be used to compute different user. For example, if you use 1d here, you'll get the number of different email addresses' from the selected domains that received an email through your server in one day.

Defaults to 1d.

select_domain

This parameter contains a regexp that will be used to select the domains for which different to addresses will be computed.

Defaults to . *.

Highest Average Delay By To Relay And To Domain Email Report

ID: top-avg-delay-by-to-relay-and-to-domain

Chart: None

This report shows the relay and domain with the highest average delay.

Parameters

delay_to_show

This parameter controls the number of delay to display in the report.

Defaults to 10.

Most Deliveries Between Relays Email Report

ID: top-deliveries-btw-relays

Chart: None

This report lists the connections between two relays with the most deliveries.

Parameters

connection_to_show

This parameter controls the number of connections to display in the report.

Defaults to 10.

Most Deliveries From Domain Email Report

ID: top-from-domain

Chart: bars

This report lists the domain from which we received the most emails.

Parameters

domain_to_show

This parameter controls the number of domains to display in the report.

Defaults to 10.

Most Deliveries From User By Domain Email Report

ID: top-from-email-by-domain

Chart: None

This report lists the user by domain from which we received the most emails.

Parameters

domain_to_show

This parameter controls the number of domains to display in the report.

Defaults to 30.

user_to_show

This parameter controls the number of user by domain to display in the report.

Defaults to 5.

Most Deliveries From Relay Email Report

ID: top-from-relay

Chart: bars

This report lists the relay from which we received the most emails.

Parameters

relay_to_show

This parameter controls the number of relays to display in the report.

Defaults to 10.

Largest Email Exchange Email Report

ID: top-largest-email-exchange

Chart: None

This report the sender and recipient that exchange the largest volume of email.

Parameters

exchange_to_show

This parameter controls the number of sender, recipient to display in the report.

Defaults to 10.

msg_to_show

This parameter controls the number of messages to display in the report.

Defaults to 5.

Largest Email Exchange Per Relay Pair Email Report

ID: top-largest-email-exchange-per-relay

Chart: None

This report lists the pairs of from- and to-relays which exchanged the largest volume of email.

Parameters

exchange_to_show

This parameter controls the number of from-relays to display in the report.

Defaults to 10.

msg_to_show

This parameter controls the number of to-relays per from-relay to display in the report.

Defaults to 5.

Status Summary For Most Common Domains

ID: top-status-by-domain

Chart: histogram

This report is a summary for the most common domains.

Parameters

domain_to_show

This parameter controls the number of domains to display in the report.

Defaults to 10.

Most Deliveries To Domain Email Report

ID: top-to-domain

Chart: bars

This report lists the domain to which we sent the most emails.

Parameters

domain_to_show

This parameter controls the number of domains to display in the report.

Defaults to 10.

Most Deliveries To User By Domain Email Report

ID: top-to-email-by-domain

Chart: None

This report lists the user by domain to which we sent the most emails.

Parameters

domain_to_show

This parameter controls the number of domains to display in the report.

Defaults to 30.

user_to_show

This parameter controls the number of user by domain to display in the report.

Defaults to 5.

Most Deliveries From Relay Email Report

ID: top-to-relay

Chart: bars

This report lists the relay to which we sent the most emails.

Parameters

relay_to_show

This parameter controls the number of relays to display in the report.

Defaults to 10.

Largest Volume Received From Domain Email Report

ID: top-volume-from-domain

Chart: bars

This report lists the domain from which we received the largest volume of mail.

Parameters

domain_to_show

This parameter controls the number of domains to display in the report.

Defaults to 10.

Largest Volume Sent To Domain Email Report

ID: top-volume-to-domain

Chart: bars

This report lists the domains to which the largest volume of mail was sent.

Parameters

domain_to_show

This parameter controls the number of domains to display in the report.

Defaults to 10.

Tracked Recipients Email Report

ID: tracked-recipients

Chart: None

This report shows the number of messages received from each sender for selected recipient emails.

Parameters

tracked_email_re

This parameter controls which recipient emails will be included in the report. E.g. "joe\\@example\\.com" or "(joe|ann|mailer-daemon)\\@example\\.org" or "root\\@example\\.com|john\\.doe\\@example\\.net":
regexps should be Perl style, with suitable escaping.

Defaults to .*.

Tracked Senders Email Report

ID: tracked-senders

Chart: None

This report shows the number of messages sent to each recipients for selected sender emails.

Parameters

tracked_email_re

This parameter controls which sender emails will be included in the report. E.g. "joe\\@example\\.com" or "(joe|ann|mailer-daemon)\\@example\\.org" or "root\\@example\\.com|john\\.doe\\@example\\.net":
 regexp's should be Perl style, with suitable escaping.

Defaults to `.*`.

Volume Delivered By Period Email Report

ID: `volume-by-period`

Chart: `histogram`

This report shows the volume of delivered emails in configurable time period.

Parameters

period

This parameter controls the time period over which the deliveries are aggregated.

Defaults to `1d`.

Email Summary

ID: `summary`

Chart: `histogram`

This report shows a summary of email deliveries.

This report doesn't have any parameters.

Filters' Descriptions and Configuration

Select Client IP Filter

ID: `select-client-ip`

This filter specification can be used to select the deliveries coming from a particular client host or relay.

Parameters

ip-range

This parameter contains the regular expression that will be used to select the deliveries coming from particular hosts. Only deliveries made by a client host matching that regexp will be included in the subreports. The match is done on the client IP address (not its hostname).

Defaults to `.*`.

Chapter 11. Firewall Reports

Supported Log Format

Lire supports logs from many packet filter firewalls.

Cisco ACL

Cisco routers that use IOS can log activity via **syslog**. Lire is able to process the logs entries corresponding to the packet filters.

Example 11-1. IOS Log Sample

```
Aug 19 04:02:34 1.example.com.nl 218963: Aug 19 04:02:32.977: \
%LINEPROTO-5-UPDOWN: Line protocol on Interface BRI0:1, changed \
state to down
Aug 19 04:02:34 1.example.com.nl 218964: Aug 19 04:02:33.262: \
%ISDN-6-DISCONNECT: Interface BRI0:1 disconnected from \
172605440 teraar, call lasted 42 seconds
Aug 19 04:02:35 1.example.com.nl 218965: Aug 19 04:02:33.266: \
%LINK-3-UPDOWN: Interface BRI0:1, changed state to down
Aug 19 04:02:38 1.example.com.nl 218966: Aug 19 04:02:36.103: \
%SEC-6-IPACCESSLOGP: list 102 denied tcp 10.0.0.1(4652) -> \
10.0.0.2(80), 1 packet
Aug 19 04:02:45 1.example.com.nl 218967: Aug 19 04:02:43.543: \
%ISDN-6-LAYER2DOWN: Layer 2 for Interface BR0, TEI 86 changed to down
Aug 19 04:02:53 1.example.com.nl 218968: Aug 19 04:02:51.471: \
%SEC-6-IPACCESSLOGP: list 102 denied tcp 10.0.0.3(2162) -> \
10.0.0.4(80), 1 packet
Aug 19 04:03:06 1.example.com.nl 218969: Aug 19 04:03:04.585: \
%ISDN-6-LAYER2DOWN: Layer 2 for Interface BRI0, TEI 86 changed to down
Aug 19 04:03:10 1.example.com.nl 218970: Aug 19 04:03:08.867: \
%SEC-6-IPACCESSLOGP: list 102 denied tcp 10.0.0.5(2342) -> \
10.0.0.6(80), 1 packet
Aug 19 04:03:12 1.example.com.nl 218971: Aug 19 04:03:10.771: \
%SEC-6-IPACCESSLOGP: list 102 denied tcp 10.0.0.7(1093) -> \
10.0.0.8(80), 1 packet
Aug 19 04:03:36 1.example.com.nl 218972: Aug 19 04:03:34.373: \
%SEC-6-IPACCESSLOGP: list 102 denied tcp 10.0.0.9(3173) -> \
10.0.0.10(80), 1 packet
```

IPChains

IPChains will log packets marked for logging through **syslog** (actually the kernel log buffer which is usually sent to **syslog**). Lire expects the logs in the form of a syslog log file.

Example 11-2. IPChains Log Sample

```

Oct 28 04:02:30 firewall kernel: Packet log: output DENY eth0 PROTO=17 \
    10.0.0.1:137 10.0.0.2:137 L=78 S=0x00 I=36930 F=0x0000 T=64 (#7)
Oct 28 04:07:30 firewall kernel: Packet log: output DENY eth0 PROTO=17 \
    10.0.0.1:137 10.0.0.2:137 L=78 S=0x00 I=37211 F=0x0000 T=64 (#7)
Oct 28 04:07:40 firewall kernel: Packet log: input DENY eth1 PROTO=17 \
    10.0.0.3:138 10.0.0.4:138 L=256 S=0x00 I=37213 F=0x0000 T=64 (#7)
Oct 28 04:07:40 firewall kernel: Packet log: input DENY eth1 PROTO=17 \
    10.0.0.3:138 10.0.0.4:138 L=236 S=0x00 I=37214 F=0x0000 T=64 (#7)
Oct 28 04:08:20 firewall kernel: Packet log: output DENY lo PROTO=17 \
    10.0.0.5:138 10.0.0.2:138 L=256 S=0x00 I=37216 F=0x0000 T=64 (#7)
Oct 28 04:12:30 firewall kernel: Packet log: output DENY eth0 PROTO=17 \
    10.0.0.1:137 10.0.0.2:137 L=78 S=0x00 I=37255 F=0x0000 T=64 (#7)
Oct 28 04:17:30 firewall kernel: Packet log: output DENY eth0 PROTO=17 \
    10.0.0.1:137 10.0.0.2:137 L=78 S=0x00 I=37364 F=0x0000 T=64 (#7)
Oct 28 04:19:40 firewall kernel: Packet log: input DENY eth1 PROTO=17 \
    10.0.0.3:138 10.0.0.4:138 L=256 S=0x00 I=37440 F=0x0000 T=64 (#7)
Oct 28 04:19:40 firewall kernel: Packet log: input DENY eth1 PROTO=17 \
    10.0.0.3:138 10.0.0.4:138 L=236 S=0x00 I=37441 F=0x0000 T=64 (#7)
Oct 28 04:20:20 firewall kernel: Packet log: output DENY lo PROTO=17 \
    10.0.0.5:138 10.0.0.2:138 L=256 S=0x00 I=37453 F=0x0000 T=64 (#7)

```

IP Filter

IP Filter logs selected packets through **syslog**.

Example 11-3. IP Filter Log Sample

```

Oct 30 07:42:29 firewall ipmon[16747]: 07:42:28.585962          ie0 @0:9 \
    b 192.168.48.1,45085 -> 192.168.48.2,22 PR tcp len 20 64 -S OUT
Oct 30 07:40:24 firewall ipmon[16747]: 07:40:23.631307          ep1 @0:6 \
    b 192.168.26.5,113 -> 192.168.26.1,3717 PR tcp len 20 40 -AR OUT
Oct 30 07:42:29 firewall ipmon[16747]: 07:42:28.585962          ie0 @0:9 \
    b 192.168.48.1,45085 -> 192.168.48.2,22 PR tcp len 20 64 -S OUT
Oct 30 07:44:11 firewall ipmon[16747]: 07:44:10.605416 2x          ep1 @0:15 \
    b 192.168.26.1,138 -> 192.168.26.255,138 PR udp len 20 257 IN
Oct 30 07:44:34 firewall ipmon[16747]: 07:44:33.891869          ie0 @0:10 \
    b 192.168.48.1,23406 -> 192.168.48.2,22 PR tcp len 20 64 -S OUT

```

IPTables

IPTables will log packets marked for logging through **syslog** (actually the kernel log buffer which is usually sent to **syslog**). Lire expects the logs in the form of a syslog log file.

A problem with logs from IPTables is that we have no real idea of what happened with the packet (was it denied or permitted). The logging module of IPTables permit to tag each logged packet with a prefix. Lire will interpret packets having a prefix which contains the strings denied, drop, deny or reject as denied packets. All other packets will have an unknown action value (-).

Example 11-4. IPTables Log Sample

```

Sep 21 11:45:17 lire kernel: Packet-drop IN=eth0 OUT=eth0 SRC=10.0.0.1 \
DST=10.0.0.2 LEN=48 TOS=0x00 PREC=0x00 TTL=113 ID=38365 DF \
PROTO=TCP SPT=3117 DPT=80 WINDOW=16384 RES=0x00 SYN URGP=0
Sep 21 11:45:20 lire kernel: Packet-drop IN=eth0 OUT=eth0 SRC=10.0.0.1 \
DST=10.0.0.2 LEN=48 TOS=0x00 PREC=0x00 TTL=113 ID=38478 DF \
PROTO=TCP SPT=3117 DPT=80 WINDOW=16384 RES=0x00 SYN URGP=0
Sep 21 11:45:26 lire kernel: Packet-drop IN=eth0 OUT=eth0 SRC=10.0.0.1 \
DST=10.0.0.2 LEN=48 TOS=0x00 PREC=0x00 TTL=113 ID=38680 DF \
PROTO=TCP SPT=3117 DPT=80 WINDOW=16384 RES=0x00 SYN URGP=0
Sep 21 11:52:46 lire kernel: Packet-drop IN=eth0 OUT=eth0 SRC=10.0.0.1 \
DST=10.0.0.3 LEN=48 TOS=0x00 PREC=0x00 TTL=113 ID=54122 DF \
PROTO=TCP SPT=4532 DPT=80 WINDOW=16384 RES=0x00 SYN URGP=0
Sep 21 11:52:49 lire kernel: Packet-drop IN=eth0 OUT=eth0 SRC=10.0.0.1 \
DST=10.0.0.3 LEN=48 TOS=0x00 PREC=0x00 TTL=113 ID=54222 DF \
PROTO=TCP SPT=4532 DPT=80 WINDOW=16384 RES=0x00 SYN URGP=0
Sep 21 11:52:55 lire kernel: Packet-drop IN=eth0 OUT=eth0 SRC=10.0.0.1 \
DST=10.0.0.3 LEN=48 TOS=0x00 PREC=0x00 TTL=113 ID=54443 DF \
PROTO=TCP SPT=4532 DPT=80 WINDOW=16384 RES=0x00 SYN URGP=0

```

WebTrends Enhanced Log Format

The WELF format is a format developed by WebTrends and supported by many firewall vendors. Products can save log files in that format directly or can log through **syslog**. Either native WELF log files or **syslog**'s log files contain WELF information. Although the log format isn't designed for packet filter firewalls (it can contain information from devices that do network intrusion or proxy services), Lire does its best to map this information to something that can be meaningful.

Example 11-5. WELF Log Sample

```

WTsyslog[1998-08-01 14:05:46 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 04:10:23" fw=WebTrendsSample pri=5 \
msg="ICMP packet dropped" src=10.0.0.2 dst=10.0.0.3 rule=3
WTsyslog[1998-08-01 16:31:00 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 10:35:38" fw=WebTrendsSample pri=6 \
proto=tcp/443 src=10.0.0.4 dst=10.0.0.5 rcvd=4844
WTsyslog[1998-08-01 16:31:01 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 10:35:38" fw=WebTrendsSample pri=6 proto=tcp/443 \
src=10.0.0.4 dst=10.0.0.5 rcvd=6601
WTsyslog[1998-08-01 16:43:59 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 10:48:36" fw=WebTrendsSample pri=5 \
msg="UDP packet dropped" src=10.0.0.6 dst=10.0.0.3 rule=3
WTsyslog[1998-08-01 16:46:13 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 10:50:50" fw=WebTrendsSample pri=5 \
msg="UDP packet dropped" src=10.0.0.7 dst=10.0.0.3 rule=3
WTsyslog[1998-08-01 16:46:13 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 10:50:50" fw=WebTrendsSample pri=6 proto=telnet \
src=10.0.0.4 dst=10.0.0.8 sent=1194

```

Lire also supports some extension uses by SonicWall.

Example 11-6. SonicWall Log Sample

```
Jan  7 15:01:10 lire id=firewall sn=asdlFFFXSD \
    time="2002-01-06 22:42:13" fw=10.0.0.1 pri=6 c=1 m=30 \
    msg="Administrator login failed - incorrect password" n=1 \
    src=10.0.0.2:LAN dst=10.0.0.1
Jan  7 15:01:16 lire id=firewall sn=asdlFFFXSD \
    time="2002-01-06 22:42:19" fw=10.0.0.1 pri=6 c=1 m=29 \
    msg="Successful administrator login" n=1 src=10.0.0.2:LAN dst=10.0.0.1
Jan  7 15:02:32 lire id=firewall sn=asdlFFFXSD \
    time="2002-01-06 22:43:34" fw=10.0.0.1 pri=5 c=128 m=37 \
    msg="UDP packet dropped" n=1 src=10.0.0.3:68 dst=10.0.0.4:67 dstname=DHCP
Jan  7 15:31:43 lire id=firewall time="2002-01-07 15:20:21" \
    fw=10.0.0.5 pri=6 proto=dns src=10.0.0.6 dst=10.0.0.8 rcvd=130 \
    sn=asdlFFFXSD 54 c=1024 m=98 n=31
Jan  7 15:31:43 10.0.0.5 id=firewall time="2002-01-07 15:20:21" \
    fw=10.0.0.5 pri=6 proto=dns src=10.0.0.6 dst=10.0.0.9 rcvd=130 \
    sn=asdlFFFXSD 54 c=1024 m=98 n=32
```

Report Descriptions and Configuration

Bytes by Period Firewall Report

ID: bytes-by-period

Chart: histogram

This report shows the number of bytes aggregated in configurable time periods.

Parameters

period

This parameter controls the time period over which the bytes are aggregated.

Defaults to 1d.

Traffic's Volume by Rule Firewall Report

ID: bytes-by-rule

Chart: bars

This report shows the volume of data logged by each rule.

This report doesn't have any parameters.

Bytes by Timeslot Firewall Report

ID: bytes-by-timeslot

Chart: histogram

This report shows the volume of traffic distributed by timeslots (hours of the day, days of the week, etc.) that passed (or were denied) by your firewall.

Parameters

timeslot

This parameter controls the length of the timeslot over which the packets are aggregated. Use **1h** for “hours of the day” or **1d** for “days of the week”.

Defaults to 1h.

Top Bytes per From-IP Report

ID: bytesperfrom

Chart: bars

This report lists the IP addresses sending the highest data volume.

Parameters

ips_to_show

This parameter controls the number of sending IP addresses to display in the report.

Defaults to 10.

Top Bytes per From-IP per Port Report

ID: bytesperfromperport

Chart: bars

This report lists the volume we were asked to receive per source IP per source port.

Parameters

ips_to_show

This parameter controls the number of sending IP addresses to display in the report.

Defaults to 10.

ports_to_show

This parameter controls the number of source ports to display in the report.

Defaults to 10.

Top Bytes per To-ip Report

ID: bytesperto

Chart: bars

This report lists the IP addresses for which we were asked to sent the highest data volume to.

Parameters

ips_to_show

This parameter controls the number of receiving IP addresses to display in the report.

Defaults to 10.

Top Bytes per destination IP per Port Report

ID: bytesperport

Chart: bars

This report lists the volume were asked to receive per destination IP per port.

Parameters

ips_to_show

This parameter controls the number of receiving IP addresses to display in the report.

Defaults to 10.

ports_to_show

This parameter controls the number of ports to display in the report.

Defaults to 10.

Top blocked tcp packets per source IP per destination port Report

ID: deniedtcpport

Chart: bars

This report lists the destination ports for which we blocked the highest tcp data volume, along with the sending ip addresses

Parameters

ips_to_show

This parameter controls the number of sending IP addresses to display in the report.

Defaults to 10.

ports_to_show

This parameter controls the number of destination ports to display in the report.

Defaults to 10.

Packets by Period Firewall Report

ID: pkt-by-period

Chart: histogram

This report shows the number of packets logged by the firewall aggregated in configurable time period.

Parameters

period

This parameter controls the time period over which the packets are aggregated.

Defaults to 1d.

Packets by Rule Firewall Report

ID: pkt-by-rule

Chart: bars

This report shows the number of packets logged by the firewall for each rules.

This report doesn't have any parameters.

Packets by Timeslot Firewall Report

ID: pkt-by-timeslot

Chart: histogram

This report shows the number of packets distributed by timeslots (hours of the day, days of the week, etc.).

Parameters

timeslot

This parameter controls the length of the timeslot over which the packets are aggregated. Use **1h** for “hours of the day” or **1d** for “days of the week”.

Defaults to 1h.

Top Volume to Destination by Source Firewall Report

ID: top-bytes-dst-by-src

Chart: None

This report will show for a number of source IP addresses that sent the most volume of traffic, the list of destination (destination IP and destination port).

Parameters

src_to_show

This parameter controls the number of source IP addresses to display in the report.

Defaults to 15.

dst_to_show

This parameter controls the number of destination (IP address and port) to display for each source IP.

Defaults to 20.

Top Volume to Destination by Source Firewall Report

ID: top-bytes-src-by-dst

Chart: None

This report will show for each destination (destination IP and port) the list of source IPs that sent the most volume.

Parameters

dst_to_show

This parameter controls the number of destination (IP address and port) to display in the report.

Defaults to 15.

src_to_show

This parameter controls the number of source IP addresses that will be displayed for each destination.

Defaults to 20.

Top Messages Firewall Report

ID: top-msg

Chart: bars

This report shows the top messages (IDS alerts or others) generated by the firewall.

Parameters

msgs_to_show

This parameter controls the number of messages to show in the report.

Top Messages Firewall Report

ID: top-dst-by-msg

Chart: None

This report shows the top destination IPs that are the target of the messages (IDS alerts or others) generated by the firewall.

Parameters

msgs_to_show

This parameter controls the number of messages to show in the report.

ips_to_show

This parameter controls the number of destination IPS to list with each message.

Top Messages Firewall Report

ID: top-src-by-msg

Chart: None

This report shows the top source IPs that are at the origin of the messages (IDS alerts or others) generated by the firewall.

Parameters

msgs_to_show

This parameter controls the number of messages to show in the report.

ips_to_show

This parameter controls the number of source IPS to list with each message.

Top Packets by Source IP Report

ID: top-pkt-by-src

Chart: bars

This report lists the IP addresses that were listed as source in the most packets.

Parameters

ips_to_show

This parameter controls the number of source IP addresses to display in the report.

Defaults to 10.

Top Packets by Destination IP Report

ID: top-pkt-by-dst

Chart: bars

This report lists the IP addresses that were listed as destination in the most packets.

Parameters

ips_to_show

This parameter controls the number of destination IP addresses to display in the report.

Defaults to 10.

Top Packets by Destination by Port Firewall Report

ID: top-pkt-by-dst-by-port

Chart: None

This report will show for each destination IP the ports that received the most packets.

Parameters

dst_to_show

This parameter controls the number of destination IPs to display in the report.

Defaults to 15.

ports_to_show

This parameter controls the number of destination ports that will be displayed for each destination.

Defaults to 20.

Top Packets Destination by Source Firewall Report

ID: top-pkt-dst-by-src

Chart: None

This report will show for a number of source IP addresses that sent the most packets, the list of destination (destination IP and destination port).

Parameters

src_to_show

This parameter controls the number of source IP addresses to display in the report.

Defaults to 15.

dst_to_show

This parameter controls the number of destination (IP address and port) to display for each source IP.

Defaults to 20.

Top Packets Source by Destination Firewall Report

ID: top-pkt-src-by-dst

Chart: None

This report will show for each destination (destination IP and port) the list of source IPs that sent the most packets.

Parameters

dst_to_show

This parameter controls the number of destination (IP address and port) to display in the report.

Defaults to 15.

src_to_show

This parameter controls the number of source IP addresses that will be displayed for each destination.

Defaults to 20.

Filter Descriptions and Configuration

Exclude Source IP Filter

ID: `exclude-from-ip`

This filter specification can be used to select packets sent from any host excluding the specified one.

Parameters

ip-range

This parameter contains the regular expression that will be used to select the packets sent from any host, excluding the matching one(s). Only packets sent from a host not matching that regexp will be included in the subreports. The match is done on the host's IP address (not its hostname).

Defaults to `. *`.

Exclude Destination IP Filter

ID: `exclude-to-ip`

This filter specification can be used to select packets sent to any host excluding the specified one.

Parameters

ip-range

This parameter contains the regular expression that will be used to select the packets sent to any host, excluding the matching one(s). Only packets sent to a host not matching that regexp will be included in the subreports. The match is done on the host's IP address (not its hostname).

Defaults to `. *`.

Select Action Filter

ID: `select-action`

This filter specification can be used to select only the firewall events that were permitted or denied.

Parameters

action_match

This parameter contains the action that should be selected:

`denied`

Select only denied events.

permitted

Select only permitted events.

-

This is also a possible action when we can't determine from the log information if this event was denied or permitted.

Defaults to denied.

Select Source IP Filter

ID: `select-from-ip`

This filter specification can be used to select the packets coming from a particular host.

Parameters

ip-range

This parameter contains the regular expression that will be used to select packets coming from particular hosts. Only packets sent from a host matching that regexp will be included in the subreports. The match is done on the client IP address (not its hostname).

Defaults to `.*`.

Select Destination IP Filter

ID: `select-to-ip`

This filter specification can be used to select packets sent to a particular host.

Parameters

ip-range

This parameter contains the regular expression that will be used to select the packets sent to a particular host. Only packets sent to a host matching that regexp will be included in the subreports. The match is done on the host's IP address (not its hostname).

Defaults to `.*`.

Select Destination Port

ID: `select-to-port`

This filter specification can be used to select packets sent to a particular port on the destination host.

Parameters

port-range

This parameter contains the regular expression that will be used to select the packets sent to a particular port. Only packets sent to a port matching that regexp will be included in the subreports. The match is done on the ports symbolic name, as found by `Libre::Firewall`'s `firewall_number2names` routine.

Defaults to `.*`.

Chapter 12. FTP Reports

Supported Log Format

Lire supports the widely used `xferlog` FTP file transfer log files and logs from the FTP service of Microsoft Internet Information Server.

Microsoft Internet Information Server

The FTP log file from Microsoft Internet Information Server is a variant of the W3C Extended Log Format defined at <http://www.w3.org/TR/WD-logfile.html>.

Lire can use the following fields of the format: *date*, *time*, *c-ip*, *c-dns*, *cs-bytes*, *time-taken*, *cs-uri-stem* and *cs-method*. The other fields will be ignored.

Example 12-1. Microsoft Internet Information Server FTP Log Sample

```
#Software: Microsoft Internet Information Server 4.0
#Version: 1.0
#Date: 2001-11-29 00:01:32
#Fields: time c-ip cs-method cs-uri-stem sc-status
00:01:32 10.0.0.1 [56]created spacedat/091001092951LGW_Data.zip 226
00:01:32 10.0.0.1 [56]created spacedat/html/bx01g01.gif 226
00:01:32 10.0.0.1 [56]created spacedat/html/catlogo.gif 226
00:01:32 10.0.0.1 [56]QUIT - 226
00:03:32 10.0.0.1 [58]USER badm 331
00:03:32 10.0.0.1 [58]PASS - 230
```

Xferlog

The `xferlog` format is supported by a wide range of FTP servers like Wu-Ftpd, ProFTPD or standard BSD `ftpd`.

Example 12-2. Xferlog Log Sample

```
Mon Feb 26 09:48:18 2001 1 1.example.com 147456 \
/var/ftp/pubinfo/sm2/esc/s82e5937.jpg b _ o a \
john.doe.1@mail.example.com ftp 0 * i
Mon Feb 26 10:26:31 2001 1 2.example.com 10593 \
/var/html/public/htdocs/pubinfo/pr/1999/28/extra-photos.html \
a _ i r kellys ftp 0 * c
Mon Feb 26 10:27:50 2001 1 2.example.com 14 \
/var/html/public/htdocs/pubinfo/pr/1999/28/extra-photos.html.LCK \
a _ i r kellys ftp 0 * c
Mon Feb 26 10:28:17 2001 1 2.example.com 14 \
/var/html/public/htdocs/pubinfo/pr/1999/28/extra-photos.html.LCK \
a _ o r kellys ftp 0 * c
Mon Feb 26 10:28:18 2001 1 2.example.com 10591 \
/var/html/public/htdocs/pubinfo/pr/1999/28/extra-photos.html \
```

```

a _ i r kellys ftp 0 * c
Mon Feb 26 12:51:02 2001 2 3.example.com 43063 \
/var/ftp/pubinfo/jpeg/EtaCar3d.jpg b _ o a mozilla@ ftp 0 * c
Mon Feb 26 12:51:17 2001 2 3.example.com 37332 \
/var/ftp/pubinfo/jpeg/EtaCarC.jpg b _ o a mozilla@ ftp 0 * c
Mon Feb 26 12:51:52 2001 6 3.example.com 62823 \
/var/ftp/pubinfo/jpeg/EtaCarD.jpg b _ o a mozilla@ ftp 0 * c
Mon Feb 26 12:52:31 2001 2 3.example.com 33660 \
/var/ftp/pubinfo/jpeg/Neptune.jpg b _ o a mozilla@ ftp 0 * c
Mon Feb 26 12:52:43 2001 2 3.example.com 26295 \
/var/ftp/pubinfo/jpeg/NeptDS.jpg b _ o a mozilla@ ftp 0 * c

```

Report Descriptions and Configuration

Top Remote Host FTP Report

ID: top-remote-host

Chart: bars

This report lists the remote hosts with the most requests.

Parameters

hosts_to_show

This parameter controls the number of remote hosts to display in the report.

Defaults to 10.

Bytes by Period FTP Report

ID: bytes-by-period

Chart: histogram

This report calculates the sum of all transfers by a user-configurable time period.

Parameters

period

This parameter controls the time period which is use to aggregate the records.

Defaults to 1d.

Bytes by User by Period FTP Report

ID: bytes-by-user-by-period

Chart: None

This report shows the bytes transferred by user by period.

Parameters

period

This parameter controls the time period which is used to aggregate the records.

Defaults to 1d.

users_to_show

This parameter controls the number of users to display during each users.

Defaults to 10.

Bytes by Period by User FTP Report

ID: bytes-by-period-by-user

Chart: None

This report shows for each user, the amount of bytes they transferred by period.

Parameters

period

This parameter controls the time period which is used to aggregate the records for each users.

Defaults to 1d.

users_to_show

This parameter controls the number of users to display in the report.

Defaults to 10.

Bytes by Direction by User with count by Period FTP Report

ID: bytes-by-dir-by-user-by-period

Chart: None

This report shows the bytes transferred by direction (in or out) by user with the number of transfers by period.

Parameters

period

This parameter controls the time period which is used to aggregate the records.

Defaults to 1d.

users_to_show

This parameter controls the number of users to display during each period.

Defaults to 10.

Top Files FTP Report

ID: top-files

Chart: bars

This report lists the most requested files.

Parameters

files_to_show

This parameter controls the number of files to display in the report.

Defaults to 10.

Top Uploaded Files FTP Report

ID: top-files-in

Chart: bars

This report lists the most uploaded files.

Parameters

files_to_show

This parameter controls the number of files to display in the report.

Defaults to 10.

Top Downloaded Files FTP Report

ID: top-files-out

Chart: bars

This report lists the most downloaded files.

Parameters

files_to_show

This parameter controls the number of files to display in the report.

Defaults to 10.

Top Users FTP Report

ID: top-users

Chart: bars

This report lists the most active users.

Parameters

users_to_show

This parameter controls the number of users to display in the report.

Defaults to 10.

Top by User (Bytes Transferred) FTP Report

ID: top-users-bytes

Chart: bars

This report lists the users with the highest amount of bytes transferred.

Parameters

users_to_show

This parameter controls the number of users to display in the report.

Defaults to 10.

Tracked Users FTP Report

ID: tracked-users

Chart: None

This report shows each download for tracked users.

Parameters

tracked_user_re

This parameter controls which user will be included in the report.

Defaults to `.*`.

Tracked Files FTP Report

ID: `tracked-files`

Chart: `None`

This report shows each download of tracked files.

Parameters

tracked_file_re

This parameter controls which files will be included in the report.

Defaults to `.*`.

Number of Transfers by Direction FTP Report

ID: `transfers-by-direction`

Chart: `pie`

This report lists the number of transfers by direction.

This report doesn't have any parameters.

Number of Transfers by Transfer Type FTP Report

ID: `transfers-by-type`

Chart: `pie`

This report lists the number of transfers by transfer type.

This report doesn't have any parameters.

Each Transfer by Filename Report

ID: `transfers-by-file`

Chart: `histogram`

This report lists the transfers for each file.

This report doesn't have any parameters.

Filter Descriptions and Configuration

None.

Chapter 13. Message Store Reports

Supported Log Formats

Lire supports log files from Netscape Messaging Server and Netsape Messaging Server Mail Multi Plexor.

Report Descriptions and Configuration

Failed Logins by Period Message Store Report

ID: badlogin-by-period

Chart: None

This report specification will generate a report showing the the number of failed login aggregated by some period of time. The period is configurable.

Parameters

period

This parameter controls the time period over which the failed logins are counted.

Defaults to 1d.

Closed Event by Period Message Store Report

ID: close-by-period

Chart: None

This report specification will generate a report showing the the number of `close` commands aggregated by some period of time. The period is configurable.

Parameters

period

This parameter controls the time period over which the `close` events are counted.

Defaults to 1d.

Events by Protocol Message Store Report

ID: events-by-protocol

Chart: None

This report specification will generate a report showing the the number of events for each protocol supported by the message store server.

This report doesn't have any parameters.

Successful Login by Period Message Store Report

ID: login-by-period

Chart: None

This report specification will generate a report showing the the number of successful logins aggregated by some period of time. The period is configurable.

Parameters

period

This parameter controls the time period over which the `close` events are counted.

Defaults to 1d.

Top User Logins Message Store Report

ID: top-user-login

Chart: None

This report specification will generate a report showing the the users that did the most logins on the message store. The number of users to include in the report is configurable.

Parameters

users_to_show

This parameter controls the number of users to include in the report.

Defaults to 10.

Top User Most Message Leftover in Store Report

ID: top-user-leftover-bytes

Chart: None

This report specification will generate a report showing the the users that transported the most bytes from the message store. The number of users to include in the report is configurable.

Parameters

users_to_show

This parameter controls the number of users to include in the report.

Defaults to 10.

Top User Most Message Leftover Store Report

ID: top-user-leftover-messages

Chart: None

This report specification will generate a report showing the the users that transported the most message from the message store. The number of users to include in the report is configurable.

Parameters

users_to_show

This parameter controls the number of users to include in the report.

Defaults to 10.

Top User Most Message Store Report

ID: top-user-most-bytes

Chart: None

This report specification will generate a report showing the the users that transported the most bytes from the message store. The number of users to include in the report is configurable.

Parameters

users_to_show

This parameter controls the number of users to include in the report.

Defaults to 10.

Top User Most Message Store Report

ID: top-user-most-messages

Chart: None

This report specification will generate a report showing the the users that transported the most message from the message store. The number of users to include in the report is configurable.

Parameters

`users_to_show`

This parameter controls the number of users to include in the report.

Defaults to 10.

Top Users doing Select Message Store Report

ID: `top-user-select`

Chart: None

This report specification will generate a report showing the the users that did the most `select` commands on the server. The number of users to include in the report is configurable.

Parameters

`users_to_show`

This parameter controls the number of users to include in the report.

Defaults to 10.

Unique Users by Period Message Store Report

ID: `unique-user`

Chart: None

This report specification will generate a report showing the the number of different users by some period of time. The period is configurable.

Parameters

`period`

This parameter controls the time period over which the different users are counted.

Defaults to 1d.

Filter Descriptions and Configuration

Select Client Host Filter

ID: `select-client-host`

This filter specification can be used to select only the message store events made by a particular client.

Parameters

client_match

This parameter contains the regular expression that will be used to select the client. Only events made by a client host matching that regexp will be included in the subreports.

Defaults to `.*`.

Chapter 14. Print Reports

Supported Log Format

The print superservice supports printer logs from two print daemons.

CUPS page_log

Information about this format can be found in the CUPS Software Administrators Manual (<http://www.cups.org/sam.html>).

Example 14-1. CUPS page_log Log Sample

```
DANKA_infotec_P450 kurt 137 [19/Aug/2001:16:58:58 +0100] 1 1
P4501 kurt 138 [19/Aug/2001:17:05:06 +0100] 1 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 2 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 3 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 4 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 5 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 6 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 7 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 8 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 9 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 10 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 11 1
P4501 kurt 138 [19/Aug/2001:17:05:08 +0100] 12 1
```

LPRng Account Log File

Example 14-2. LPRng Log Sample

Lire can process the accounting file associated with a LPRng print queue. The format of the file is described at <http://www.lprng.org/LPRng-HOWTO-Multipart/x9481.htm>

```
jobstart '-Hh4.private' '-nroot' '-Pps' '-kcfA938h4.private' \
'-b1093' '-tNov 5 19:39:25'
start '-p12942' '-kcfA938h4.private' '-nroot' '-hh4.private' '-Pps' \
'-c0' '-Fo' '-tSun Nov 5 19:39:25 1995'
filestart '-p12944' '-kcfA938h4.private' '-nroot' '-hh4.private' '-Pps' \
'-c0' '-Ff' '-tSun Nov 5 19:39:27 1995'
fileend '-p12944' '-kcfA938h4.private' '-nroot' '-hh4.private' '-Pps' \
'-b3' '-c0' '-Ff' '-tSun Nov 5 19:39:58 1995'
end '-p12942' '-kcfA938h4.private' '-nroot' '-hh4.private' '-Pps' \
'-b2' '-c0' '-Fo' '-tSun Nov 5 19:39:59 1995'
jobend '-Hh4.private' '-nroot' '-Pps' '-kcfA938h4.private' \
'-b1093' '-tNov 5 19:39:59'
```

Report Descriptions and Configuration

Jobs per Printer Print Report

ID: `jobs-per-printer`

Chart: bars

This report shows the number of jobs for each printer.

This report doesn't have any parameters.

Top Users Print Report

ID: `top-users`

Chart: bars

This report lists the users with the most print jobs.

Parameters

users_to_show

This parameter controls the number of users to display in the report.

Defaults to 10.

Number Of Jobs For Each Number Of Sheets

ID: `jobs-per-sheets`

Chart: None

This report shows the number of jobs sorted by the number of sheets in them.

Parameters

range

This parameter controls the number of sheets into which the jobs are aggregated.

Defaults to 5.

Jobs per Period Print Report

ID: `jobs-per-period`

Chart: None

This report lists the number of jobs during different periods.

Parameters

period

This parameter controls the time period over which the jobs are aggregated.

Defaults to 1d.

Jobs per User per Period Print Report

ID: jobs-per-user-per-period

Chart: None

This report lists the daily total of jobs per user for each printer.

Parameters

period

This parameter controls the time period over which the usage is aggregated.

Defaults to 1d.

Jobs per Printer per Period Print Report

ID: jobs-per-printer-per-period

Chart: None

This report lists the daily total of jobs for each printer.

Parameters

period

This parameter controls the time period over which the jobs are aggregated.

Defaults to 1d.

Sheets per User Print Report

ID: sheets-per-user

Chart: bars

This report shows the number of sheets printed by each user.

This report doesn't have any parameters.

Sheets per Period Print Report

ID: sheets-per-period

Chart: None

This report shows the number of sheets split into periods.

Parameters

period

This parameter controls the time period over which the sheets are aggregated.

Defaults to 1d.

Sheets per User per Period Print Report

ID: sheets-per-user-per-period

Chart: None

This report shows the number of sheets printed by each user spliced out into periods.

Parameters

period

This parameter controls the time period over which the sheets are aggregated.

Defaults to 1d.

Billing Report

ID: billing

Chart: None

This report shows the billing totals.

This report doesn't have any parameters.

Billing per Printer Report

ID: billing-per-printer

Chart: None

This report shows the billing totals per printer.

This report doesn't have any parameters.

Filter Descriptions and Configuration

None.

Chapter 15. Proxy Reports

Supported Log Formats

Lire supports three different proxy log file formats allowing it to support a wide range of products.

Microsoft Internet Security and Acceleration Server

This product uses a format derived from the W3C Extended Log Format which is defined at <http://www.w3.org/TR/WD-logfile.html>. Information about the way Microsoft Internet Security and Acceleration Server uses that format can be found on the product's website

(http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/isa/proddocs/isadocs/M_S_C_LoggingFields.asp)

The format of

Lire can use the following fields of the format: *date*, *time*, *c-ip*, *c-host*, *cs-username*, *c-agent*, *time-taken*, *r-ip*, *r-host*, *sc-status*, *sc-protocol*, *sc-operation*, *s-object-source*, *sc-operation*, *rule#1*, *rule#2* and *cs-mime-type*. The other fields will be ignored.

Example 15-1. Microsoft Internet Security and Acceleration Server Log Sample

```
#Software: Microsoft(R) Internet Security and Acceleration Server 2000
#Version: 1.0
#Date: 2002-01-16 07:00:01
#Fields: c-ip cs-username c-agent date time s-computername \
        cs-referred r-host r-ip r-port time-taken cs-bytes\
        sc-bytes cs-protocol s-operation cs-uri s-object-source \
        sc-status
10.0.0.1 anonymous Mozilla/4.0 (compatible; MSIE 5.0; Win32)\
2002-01-16 07:00:01 GRO1SYX01 - - - -\
- 155 2569 - GET - - 200 \
10.0.0.1 anonymous Outlook Express/5.0 \
(MSIE 5.0; Windows 98; DigExt) 2002-01-16 07:00:04 \
GRO1SYX01 - 1.example.com
```

Squid

Lire can process native Squid access logs.

Example 15-2. Squid Log Sample

```
1011164724.171 1337 10.0.0.1 TCP_MISS/200 20110 GET \
http://images.google.com/images? - DIRECT/10.0.0.2 text/html
1011164724.965 740 10.0.0.1 TCP_MISS/200 26461 GET \
http://www.ia.hiof.no/informatikk/forelesning/historie/historie.html \
- DIRECT/10.0.0.3 text/html
1011164727.626 2580 10.0.0.1 TCP_MISS/200 111927 GET \
```

```

http://www.ia.hiof.no/informatikk/forelesning/historie/transistor.jpg \
- DIRECT/10.0.0.3 image/jpeg
1011164731.619    687 10.0.0.1 TCP_MISS/200 18191 GET \
http://images.google.com/images? - DIRECT/10.0.0.2 text/html
1011164734.972    3282 10.0.0.1 TCP_MISS/200 29595 GET \
http://www.hillnews.com/restaurants/rst_tosca.shtm - \
DIRECT/10.0.0.4 text/html
1011164735.482    467 10.0.0.1 TCP_MISS/200 7839 GET \
http://www.hillnews.com/global/banner_logo.gif - \
DIRECT/10.0.0.4 image/gif
1011164740.163    1004 10.0.0.1 TCP_MISS/200 19580 GET \
http://images.google.com/images? - DIRECT/10.0.0.2 text/html
1011164741.905    1687 10.0.0.1 TCP_MISS/200 17383 GET \
http://www.charlotteregional.com/speech.html - DIRECT/10.0.0.5 text/html
1011164742.214    275 10.0.0.1 TCP_MISS/200 8001 GET \
http://www.charlotteregional.com/images/st2.jpg - \
DIRECT/10.0.0.5 image/jpeg
1011164745.891    716 10.0.0.1 TCP_MISS/200 18796 GET \
http://images.google.com/images? - DIRECT/10.0.0.2 text/html

```

WebTrends Enhanced Format

The WELF format is a format developed by WebTrends and supported by many firewall vendors. Products can save log files in that format directly or can log through **syslog**. Either the WELF log files or **syslog**'s log files contain WELF information. This format can be used by packet filter firewalls, proxies or network intrusion detection devices. This Lire superservice will only process records that are related to proxy services (either application proxy like a web proxy or a transport proxy like for the telnet protocol).

Example 15-3. WELF Log Sample

```

WTsyslog[1998-08-01 00:04:11 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 00:08:52" fw=WebTrendsSample pri=6 proto=http \
src=10.0.0.2 dst=10.0.0.3 dstname=1.example.com \
arg=/selfupd/x86/en/WULPROTO.CAB op=GET result=304 sent=898
WTsyslog[1998-08-01 00:04:12 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 00:08:52" fw=WebTrendsSample pri=6 proto=http \
src=10.0.0.2 dst=10.0.0.3 dstname=1.example.com \
arg=/selfupd/x86/en/CUNPROT2.CAB op=GET result=304 sent=853
WTsyslog[1998-08-01 00:04:23 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 00:09:03" fw=WebTrendsSample pri=6 proto=http \
src=10.0.0.2 dst=10.0.0.3 dstname=1.example.com \
arg=/R510/v31content/90820/0x00000409.gng op=GET result=304 sent=2983
WTsyslog[1998-08-01 03:02:03 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 03:06:43" fw=WebTrendsSample pri=6 proto=http \
src=10.0.0.2 dst=10.0.0.4 dstname=2.example.com arg=/ op=POST \
result=200 sent=2195
WTsyslog[1998-08-01 16:25:33 ip=10.0.0.1 pri=6] id=firewall \
time="1998-08-01 06:30:09" fw=WebTrendsSample pri=6 proto=http \
src=10.0.0.5 dst=10.0.0.6 dstname=3.example.com \
arg=/portal/brand/images/logo_pimg.gif op=GET result=304 rcvd=1036

```

Report Descriptions and Configuration

Bytes by Cache Result

ID: bytes-by-cache_result

Chart: bars

This report shows the number of bytes transferred for each cache result. This gives an idea of how effective is the cache.

This report doesn't have any parameters.

Bytes by Object's Source

ID: bytes-by-result_src_code

Chart: bars

This report shows the number of bytes transferred from each source (Internet, Parent or sibling cache, etc.). This report is useful to analyze the performance of your array.

This report doesn't have any parameters.

Bytes Transferred By Period Proxy Report

ID: bytes-by-period

Chart: histogram

This report shows the number of bytes transferred by a user-configurable time period.

Parameters

period

This parameter controls the time period which is use to aggregate the records.

Defaults to 1d.

Bytes Transferred By Timeslot Proxy Report

ID: bytes-by-timeslot

Chart: histogram

This report shows the number of bytes transferred by user-configurable timeslot (hours of the day, days of the week, etc.). This report is useful to spot the most used period of the day for example.

Parameters

timeslot

This parameter controls the unit of time used to aggregate the records.

Defaults to 1h.

Requests by Cache Result

ID: `requests-by-cache_result`

Chart: bars

This report shows the number of requests for each cache result. This gives an idea of how effective is the cache.

This report doesn't have any parameters.

Requests By Period Proxy Report

ID: `requests-by-period`

Chart: histogram

This report shows the number of requests by a user-configurable time period.

Parameters

period

This parameter controls the time period which is use to aggregate the records.

Defaults to 1d.

Requests By Size Proxy Report

ID: `requests-by-size`

Chart: histogram

This report shows the number of requests by size. It shows the number of requests that were between 1k-5k, 5k-10k, etc.

Parameters

range_size

This parameter controls the size of the first range class.

Defaults to 1k.

Number of Requests By Timeslot Proxy Report

ID: `requests-by-timeslot`

Chart: histogram

This report shows the number of requests by user-configurable timeslot (hours of the day, days of the week, etc.).

Parameters

timeslot

This parameter controls the unit of time used to aggregate the records.

Defaults to 1h.

Requests By Request's Time Proxy Report

ID: requests-by-time

Chart: histogram

This report shows the number of requests by the time they took to process. It shows the number of requests that took between 0-1s, 1-2s, etc.

Parameters

range_size

This parameter controls the size of the first range class.

Defaults to 1s.

Top Clients by Destinations Proxy Report

ID: top-clients-by-destinations

Chart: None

This report lists the most popular destinations along with the clients that accessed them.

Parameters

clients_to_show

This parameter controls the number of clients for each destination to display in the report.

Defaults to 10.

dsts_to_show

This parameter controls the number of destinations to display in the report.

Defaults to 10.

Top Destinations by Number of Requests

ID: top-destinations

Chart: bars

This report lists most popular sites

Parameters

dsts_to_show

This parameter controls the number of destinations to display in the report.

Defaults to 10.

Top Destinations by Bytes Downloaded

ID: top-destinations-by-bytes

Chart: bars

This report lists most popular sites by traffic

Parameters

dsts_to_show

This parameter controls the number of destinations to display in the report.

Defaults to 10.

Top Destinations by Clients

ID: top-destinations-by-clients

Chart: None

This report lists most popular sites by clients

Parameters

dsts_to_show

This parameter controls the number of destinations to display in the report.

Defaults to 10.

clients_to_show

This parameter controls the number of clients to display in the report.

Defaults to 10.

Top Destinations by Users Proxy Report

ID: top-destinations-by-users

Chart: None

This report lists most popular destinations, grouped by users

Parameters

users_to_show

This parameter controls the number of users for each URL to display in the report.

Defaults to 10.

dsts_to_show

This parameter controls the number of destinations to display in the report for each user.

Defaults to 10.

Top Users by Destinations Proxy Report

ID: top-users-by-destinations

Chart: None

This report lists the most popular destinations along with the users that accessed them.

Parameters

users_to_show

This parameter controls the number of users for each destinations to display in the report.

Defaults to 10.

dsts_to_show

This parameter controls the number of destinations to display in the report.

Defaults to 10.

Top MIME types by Transferred Size

ID: top-types-by-bytes

Chart: bars

This report lists the MIME types that resulted in the most traffic.

Parameters

types_to_show

This parameter controls the number of MIME types to display in the report.

Defaults to 10.

Top Users by Bytes Proxy Report

ID: top-users-by-bytes

Chart: bars

This report lists the users who downloaded biggest volume using the proxy

Parameters

users_to_show

This parameter controls the number of users to display in the report.

Defaults to 10.

Top URLs by Users Proxy Report

ID: top-urls-by-users

Chart: None

This report lists most popular URLs, grouped by users

Parameters

users_to_show

This parameter controls the number of users for each URL to display in the report.

Defaults to 10.

urls_to_show

This parameter controls the number of URLs to display in the report.

Defaults to 10.

Filter Descriptions and Configuration

Select Cache Result Filter

ID: `select-cache_result`

This filter specification can be used to select only DLF records in the proxy superservice which have a particular `cache_result` value. For example, you could select only denied requests by using the value `TCP_DENIED`.

Parameters

result

This parameter is a regular expression which will be used to select the codes you want.

Defaults to `TCP_DENIED`.

Chapter 16. Syslog Reports

Supported Log Formats

Lire supports more than 7 different syslog file formats.

Report Descriptions and Configuration

Messages by Facility Syslog Report

ID: `messages-by-facility`

Chart: None

This report specification will generate a report showing the number of messages logged to each facility.

This report doesn't have any parameters.

Messages by Level Syslog Report

ID: `messages-by-level`

Chart: None

This report specification will generate a report showing the number of messages logged to each syslog level.

This report doesn't have any parameters.

Messages by Period Syslog Report

ID: `messages-by-period`

Chart: None

This report specification will generate a report showing the number of messages logged by configurable time period.

Parameters

period

The period used to aggregate the events.

Defaults to 1h.

Top Hosts Syslog Report

ID: `top-hosts`

Chart: None

This report specification will generate a report showing the hosts that logged the most messages.

Parameters

hosts_to_show

The number of hosts to include in the report.

Defaults to 10.

Top Messages by Period Syslog Report

ID: top-messages-by-period

Chart: None

This report specification will generate a report showing the messages logged most often during each period. The number of messages to include as well as the period over which events are aggregated are configurable.

Parameters

messages_to_show

The number of messages to include for each period in the report.

Defaults to 10.

period

The period used to aggregate the events.

Defaults to 1h.

Top Messages by Process Syslog Report

ID: top-messages-by-process

Chart: None

This report specification will generate a report showing the the messages that were logged the most number of times by each process. The number of messages to display and the number of processes to include in the report are customizable.

Parameters

processes_to_show

The number of processes to include in the report.

Defaults to 30.

messages_to_show

The number of messages to include for each process in the report.

Defaults to 10.

Top Messages Syslog Report

ID: top-messages

Chart: None

This report specification will generate a report showing the the messages that were logged the most number of times. The number of messages to display is customizable.

Parameters

messages_to_show

The number of messages to include in the report.

Defaults to 50.

Top Processes by Period Syslog Report

ID: top-processes-by-period

Chart: None

This report specification will generate a report showing the processes that logged the most messages during each period. The number of process to include as well as the period over which events are aggregated are configurable.

Parameters

processes_to_show

The number of processes to include for each period in the report.

Defaults to 10.

period

The period used to aggregate the events.

Defaults to 1h.

Top Processes Syslog Report

ID: top-processes

Chart: None

This report specification will generate a report showing the processes that logged the most messages.

Parameters

processes_to_show

The number of processes to include in the report.

Defaults to 10.

Filter Descriptions and Configuration

Exclude Message Filter

ID: `exclude-message`

This filter can be used to exclude syslog records based on the content of the *message* field.

Parameters

message_match

Records for which the *message* field matches that regular expression will be excluded. The match is case insensitive.

Defaults to `.*`.

Exclude Priority Filter

ID: `exclude-priority`

This filter can be used to exclude syslog records based on the priority (*facility* and *level*) of the message.

Parameters

facility_match

Regular expression used to select the facility that the event *must not* come from. The match is case insensitive.

Defaults to `.*`.

level_match

Regular expression used to select the level that the event *must not* come from. The match is case insensitive.

Defaults to `.*`.

Exclude Process Filter

ID: `exclude-process`

This filter can be used to exclude syslog records based on the value in the *process* field.

Parameters

process_match

Regular expression that excludes messages coming from matching processes. The match is case insensitive.

Defaults to `.*`.

Select Host Filter

ID: `select-host`

This filter can be used to select syslog records based on the host from which the event was received.

Parameters

host_match

Regular expression used to select the host that the event must come from. The match is case insensitive.

Defaults to `.*`.

Select Message Filter

ID: `select-message`

This filter can be used to select syslog records based on the content of the *message* field.

Parameters

message_match

Only records for which the *message* field matches that regular expression will be selected. The match is case insensitive.

Defaults to `.*`.

Select Priority Filter

ID: `select-priority`

This filter can be used to select syslog records based on the priority (*facility* and *level*) of the message.

Parameters

facility_match

Regular expression used to select the facility that the event must come from. The match is case insensitive.

Defaults to `.*`.

level_match

Regular expression used to select the level that the event must come from. The match is case insensitive.

Defaults to `.*`.

Select process Filter

ID: `select-process`

This filter can be used to select syslog records based on the value in the *process* field.

Parameters

process_match

Regular expression that select messages coming from matching processes. The match is case insensitive.

Defaults to `.*`.

Chapter 17. WWW Reports

Supported Log Format

The WWW superservice supports four log file formats which makes it possible to support a wide range of web servers like Apache, IIS or Boa.

Common Log Format

Common Log Format (CLF) is a standard log format that was originally implemented in the CERN httpd web server but that is supported nowadays by most web servers. Apache, IIS and Boa can be configured to log in that format.

The Common Log Format has the following format:

```
remotehost rfc931 authuser [date] "request" status bytes
```

where the fields have the following meaning:

remotehost

The host that made the request. This can be given as an IP address or a hostname.

rfc931

The result of an ident lookup on the host. This is usually not used.

authuser

The authenticated username.

date

The timestamp of the request.

request

The first line of the HTTP request. Usually in the format "*method file protocol*".

status

The result status of the request. i.e. 200, 301, 404, 500.

bytes

The size of the response sent back to the client.

Example of log lines in Common Log Format :

```
127.0.01 - - [11/03/2001 12:12:01 -0400] "GET / HTTP/1.0" 200 513
dsl1.myprovider.com - francis [11/03/2001 12:14:01 -0400] \
"GET /secret/ HTTP/1.0" 200 1256
```

Combined Log Format

The combined log format is an extension to the Common Log Format. It adds informations about the user agent and referer. It is also known as the extended common log format. It was first implemented in the NSCA httpd web server but is now supported in many web servers. Apache can be configured to use this log format.

Two fields are added at the end of the common log lines:

```
"referer" "useragent"
```

referer

The content of the Referer header of the request. This usually reflects the page the user visited before this request.

useragent

The content of the User-Agent header of the request. This usually reflects the browser that the user is using.

CLF With mod_gzip Extensions

Mod_gzip is another extension to the common log format. It is used by the mod_gzip Apache extension which can be used to compress the result of requests before sending them to the client.

mod_gzip is a module developed by RemoteCommunications, Inc. Sourcecode is freely available from http://www.RemoteCommunications.com/apache/mod_gzip/mod_gzip. More informations can be found in their FAQ (http://www.RemoteCommunications.com/apache/mod_gzip/mod_gzip_faq.htm).

mod_gzip can log information about the compression of pages. To enable this, one can configure Apache to log using the 'gzip' format which can be defined as follows:

```
LogFormat "%h %l %u %t \"%r\" %>s %b %{mod_gzip_result}n \
          %{mod_gzip_compression_ratio}n" gzip
```

This adds two fields at the end of each common log line:

gzip_result compression_ratio

gzip_result

The **gzip** result code. Usually OK.

compression_ratio

The ratio by which the content was compressed. A number from 0 to 100.

Referer Log Format

The Referer log format is an old format that was implemented in the NSCA httpd server. It was used to log information about the request's referer in a separate log file. The combined log format has made this log format obsolete.

Referer log files have the following format:

uridocument

uri

The referring URI. This is the content of the Referer header of the request which usually reflects the page where the user was before that request.

document

The local document that was referenced by that URI. This is the requested file without any query string.

Logs With Virtual Host Information

You may encounter log files that have a field containing the virtual host for which the requests was at the beginning of the line. The rest of the line is usually in the common or combined log format. This kind of logging is typically seen on web servers hosting several virtual servers.

Example of such a line:

```
www.example.com 1.7.2.21 - - [13/Oct/2000:10:30:16 +0200] \
  "GET / HTTP/1.0" 200 83
```

Although Lire doesn't directly support such logs, it is easy to split those logs into many log files in the common or combined log format which can subsequently be processed by Lire.

Example doing this in a shell:

```
$ mkdir apache-common.log
$ (while read virt rest; do echo $rest >> \
  apache-common.log/$virt; done) < /var/log/apache/common.log
$ for f in apache-common.log/*; do \
  lr_log2mail -s "$f" common joe@example.com < $f; done
```

W3C Extended Log Format

This is a log format defined by the W3C which can contain a variable amount of information. The format is defined at <http://www.w3.org/TR/WD-logfile.html>.

This log format uses a header to specify the order of the fields present in the log file.

Lire can use the following fields of the format: *date*, *time*, *c-ip*, *c-dns*, *cs-uri*, *cs-method*, *sc-bytes*, *sc-status*, *cs(User-Agent)*, *cs(Referer)*, *cs-uri-stem* and *cs-username*. The other fields will be ignored.

Report Descriptions and Configuration

Bytes By Period WWW Report

ID: bytes-by-period

Chart: histogram

This report calculates the sum of all transfers by a user-configurable time period.

Parameters

period

This parameter controls the time period which is use to aggregate the requests' size.

Defaults to 1d.

Bytes Per Directory WWW Report

ID: bytes-by-dir

Chart: bars

This report shows the amount of data requested by directory. This amount only includes the size of request in that directory, not of any child directories.

Parameters

bytesdir_to_show

This parameter controls the number of different directories to display in the report.

Defaults to 10.

Bytes By HTTP Result By Period WWW Report

ID: bytes-by-result-by-period

Chart: None

This report calculates the size of all requests by HTTP result by a user-configurable time period.

Parameters

period

This parameter controls the time period which is use to aggregate the records.

Defaults to 1d.

Bytes By HTTP Result WWW Report

ID: bytes-by-result

Chart: pie

This report calculates the size of all requests by HTTP result.

This report doesn't have any parameters.

Bytes Per Request WWW Report

ID: bytes-by-request

Chart: bars

This report shows the amount of data requested by request.

Parameters

url_to_show

This parameter controls the number of different files (or urls) to display in the report.

Defaults to 10.

Client Hosts By Period WWW Report

ID: clienthost-by-period

Chart: histogram

This report count the number of different hosts that made requests by a user-configurable time period.

Parameters

period

This parameter controls the time period which is use to aggregate the records.

Defaults to 1d.

Search Engines with Keywords Report

ID: keywords-by-search_engine

Chart: None

This report shows the search engines used to access your site along with the keywords used.

Parameters

keyword_to_show

This parameter controls the number of keywords to display in the report.

Defaults to 5.

engine_to_show

This parameter controls the number of search engines to display in the report.

Defaults to 10.

Requests By Browser WWW Report

ID: requests-by-browser

Chart: bars

This report shows the number of requests for each browser.

Parameters

browsers_to_show

This parameter controls the number of Browsers to display in the report.

Defaults to 10.

Number of Requests By Period WWW Report

ID: requests-by-period

Chart: histogram

This report shows the number of requests by a user-configurable time period.

Parameters

period

This parameter controls the time period which is use to aggregate the records.

Defaults to 1d.

Requests By Browser Language WWW Report

ID: requests-by-lang

Chart: bars

This report shows the number of requests for each browser language.

Parameters

lang_to_show

This parameter controls the number of Browsers to display in the report.

Defaults to 10.

Requests By HTTP Method WWW Report

ID: requests-by-method

Chart: pie

This report shows the number of requests for each HTTP method (POST, GET, PUT, etc.)

This report doesn't have any parameters.

Requests By OS WWW Report

ID: requests-by-os

Chart: bars

This report shows the number of requests for each browser OS.

Parameters

os_to_show

This parameter controls the number of OS to display in the report.

Defaults to 10.

Requests By Result By Period WWW Report

ID: requests-by-result-by-period

Chart: None

This report calculates the number of requests by HTTP result by a user-configurable time period.

Parameters

period

This parameter controls the time period which is use to aggregate the records.

Defaults to 1d.

Requests By HTTP Result WWW Report

ID: requests-by-result

Chart: pie

This report shows the number of requests for each HTTP result.

This report doesn't have any parameters.

Requests By Gzip Result WWW Report

ID: requests-by-gzip_result

Chart: pie

This report shows the number of requests for each Gzip result.

This report doesn't have any parameters.

Requests By Robot Report

ID: requests-by-robot

Chart: bars

This report shows the number of requests for each robot.

Parameters

robots_to_show

This parameter controls the number of Browsers to display in the report.

Defaults to 10.

Requests By Top Level Domain Report

ID: requests-by-country

Chart: pie

This report shows the number of requests for top level domain.

Parameters

country_to_show

This parameter controls the number of countries to show.

Defaults to 0.

Requests By Attack Report

ID: requests-by-attack

Chart: bars

This report shows the number of requests for each attack.

This report doesn't have any parameters.

Requests By Keywords Report

ID: requests-by-keywords

Chart: bars

This report shows the number of requests that resulted from keywords in search engine.

Parameters

keyword_to_show

This parameter controls the number of keywords to display in the report.

Defaults to 10.

Requests By User Agent WWW Report

ID: requests-by-useragent

Chart: bars

This report shows the number of requests for each user agent.

This report doesn't have any parameters.

Requests By Search Engines Report

ID: requests-by-search_engine

Chart: bars

This report shows the number of requests that resulted from search engines.

This report doesn't have any parameters.

Number of Requests By Size WWW Report

ID: requests-by-size

Chart: histogram

This report shows the number of requests by request's size.

Parameters

range_size

This parameter controls the size of the first range class.

Defaults to 1k.

Number of Requests By Timeslot WWW Report

ID: requests-by-timeslot

Chart: histogram

This report shows the number of requests by timeslot.

Parameters

timeslot

This parameter controls the unit of time used to aggregate the records.

Defaults to 1h.

Requests By HTTP Protocol Version WWW Report

ID: requests-by-version

Chart: pie

This report shows the number of requests for each HTTP Protocol Version.

This report doesn't have any parameters.

Average Compression By File Type WWW Report

ID: top-avg-compression-by-file-type

Chart: bars

This report lists show the file extension which are on average most compressed.

Parameters

file_type_to_show

This parameter controls the number of different extension to display in the report.

Defaults to 10.

Most Averaged Compressed Requested File WWW Report

ID: top-avg-compression-by-url

Chart: bars

This report lists show the files with the most compression.

Parameters

url_to_show

This parameter controls the number of URL to display in the report.

Defaults to 10.

Top Client By HTTP Result WWW Report

ID: top-client_host-by-result

Chart: None

This report lists the client hosts with the most requests for each HTTP results.

Parameters

client_to_show

This parameter controls the number of client hosts to display in the report.

Defaults to 5.

Top Client by Size WWW Report

ID: top-client_host-by-size

Chart: bars

This report lists the client hosts with the most nof bytes requested.

Parameters

client_to_show

This parameter controls the number of client hosts to display in the report.

Defaults to 10.

Top Client WWW Report

ID: top-client_host

Chart: bars

This report lists the client hosts with the most requests.

Parameters

client_to_show

This parameter controls the number of client hosts to display in the report.

Defaults to 10.

Last Pages By Session WWW Report

ID: top-last_page

Chart: None

This report shows the requested pages (this excludes pictures) which were the last page to be visited in a user session.

Parameters

page_to_show

This parameter controls the number of last pages to display.

Defaults to 10.

First Pages By Session WWW Report

ID: top-first_page

Chart: None

This report shows the requested pages (this excludes pictures) which were the first page to be visited in a user session.

Parameters

page_to_show

This parameter controls the number of last pages to display.

Defaults to 10.

Most Travelled Referer -> Page Connections WWW Report

ID: top-referer-page-connections

Chart: bars

This report lists the connections between a referers and a pages with the most requests.

Parameters

connection_to_show

This parameter controls the number of connections to display in the report.

Defaults to 10.

referer_exclusion

This parameter can be used to exclude referers from the report. For example, it can be used to exclude internal links by setting it to "`^http://www.yourdomain.com/$`". It defaults to "`^-$`" which excludes all records for which the referer wasn't specified.

Defaults to `^-$`.

target_exclusion

This parameter contains a regular expression which is used to filter out images as targets.

Defaults to `\.(png|gif|jpg)$`.

Top Referring Pages WWW Report

ID: top-referers

Chart: None

This report shows the top referring pages to your site.

Parameters

referer_to_show

This parameter controls the number of referring pages to show.

Defaults to 10.

referrer_exclusion

This parameter can be used to exclude referers from the report. For example, it can be used to exclude internal links by setting it to "`^http://www.yourdomain.com/$`". It defaults to "`^-$`" which excludes all records for which the referer wasn't specified.

Defaults to `^-$`.

Top Referring Pages By Requested Page WWW Report

ID: `top-referers-by-page`

Chart: None

This report shows the top referring pages for the top requested page.

Parameters

page_to_show

This parameter controls the number of requested pages to display.

Defaults to 10.

referrer_to_show

This parameter controls the number of referring pages to show for each requested page.

Defaults to 5.

not_page_re

This parameter contains a regular expression which is used to filter out images.

Defaults to `(png|gif|jpg|css|jpeg)$`.

referrer_exclusion

This parameter can be used to exclude referers from the report. For example, it can be used to exclude internal links by setting it to "`^http://www.yourdomain.com/$`". It defaults to "`^-$`" which excludes all records for which the referer wasn't specified.

Defaults to `^-$`.

Top Referring Sites WWW Report

ID: `top-referring_sites`

Chart: None

This report shows the top referring sites to your site.

Parameters

site_to_show

This parameter controls the number of referring sites to show.

Defaults to 10.

Most Requested Pages WWW Report

ID: top-requested-page

Chart: bars

This report lists the most requested pages.

Parameters

page_to_show

This parameter controls the number of pages to display in the report.

Defaults to 10.

Top Traversals WWW Report

ID: top-traversals

Chart: None

This report shows the most visited paths used by your users when they visit your website.

Parameters

level1_to_show

This parameter controls the number entry points to display.

Defaults to 10.

level2_to_show

This parameter controls the number of second page to show.

Defaults to 5.

level3_to_show

This parameter controls the number of third page to show.

Defaults to 5.

Top URLs By HTTP Result WWW Report

ID: top-urls-by-result

Chart: None

This report shows for each HTTP result the URLs with the most requests that ended in that result. This report is useful to find broken links and pages with access problem.

Parameters

url_to_show

This parameter controls the number of URLs to display in the report for each HTTP result.

Defaults to 10.

Most Requested URLs By Client Host WWW Report

ID: top-urls-by-client_host

Chart: None

This report shows the most requested URLs by client host.

Parameters

url_to_show

This parameter controls the number of URLs to display for each client host.

Defaults to 5.

client_to_show

This parameter controls the number of client host to display.

Defaults to 10.

User Sessions By Period WWW Report

ID: user_sessions-by-period

Chart: histogram

This report shows the number of user sessions by time period.

Parameters

period

This parameter controls the time period over which the user sessions are aggregated.

Defaults to 1d.

Recurring Visitors WWW Report

ID: user_session-visit_number

Chart: histogram

This report shows the number of sessions for recurring users.

This report doesn't have any parameters.

Visit times User Session WWW Report

ID: user_session-visit-times

Chart: histogram

This report shows the time a user took to visit the website.

Parameters

range_size

This parameter controls the time period over which the visit times are aggregated.

Defaults to 1m.

Page Counts User Session WWW Report

ID: user_session-page_counts

Chart: histogram

This report shows the number of pages for the visits.

Parameters

range_size

This parameter controls the ranges over which the number of pages per visit are aggregated.

Defaults to 1.

Filter Descriptions and Configuration

Select URL Filter

ID: `select-url`

This filter specification can be used to select only the requests for particular URL.

For example, this filter could be used to create subreports about downloadable files from your website.

Parameters

url_match

This parameter contains the regular expression that will be used to select the requests. Only requests where URL matches that regexp will be included in the subreports.

Defaults to `\.tar.gz$`.

Select Sessions by Page Filter

ID: `select-sessions-by-page`

This filter specification can be used to select sessions that include pages matching a regular expression. The filter is applied on all fields that contains page information.

Parameters

page_match

This parameter contains the regular expression that will be used to select the sessions. Sessions for which one of the page fields (`first_page`, `last_page`, `page_2`, etc.) match the regexp will be selected.

Select Client Host Filter

ID: `select-client_host`

This filter specification can be used to select only the requests by a particular client.

Parameters

client_match

This parameter contains the regular expression that will be used to select the client. Only requests made by a client host matching that regexp will be included in the subreports.

Defaults to `.*`.

Exclude URL Filter

ID: `exclude-url`

This filter specification can be used to exclude requests for particular URLs.

For example, this filter could be use to create subreports excluding images or other multimedia file.

Parameters

url_match

This parameter contains the regular expression that will be used to filter out the requests. Requests made for a URL matching that regexp will be excluded from the subreports.

Defaults to `\.(png|jpg|gif|jpeg)$`.

Exclude Sessions by Page Filter

ID: `exclude-sessions-by-page`

This filter specification can be used to exclude sessions that include pages matching a regular expression. The filter is applied on all fields that contains page information.

Parameters

page_match

This parameter contains the regular expression that will be used to filter out the sessions. Sessions for which one of the page fields (`first_page`, `last_page`, `page_2`, etc.) match the regexp will be excluded.

Exclude Client Host Filter

ID: `exclude-client_host`

This filter specification can be used to exclude requests coming from particular hosts from the repots.

Parameters

client_match

This parameter contains the regular expression that will be used to filter out the client. Requests made by a client host matching that regexp will be excluded from the subreports.

Defaults to `.*`.

Exclude Referrer Filter

ID: `exclude-referrer`

This filter specification can be used to exclude requests with a particular referrer.

For example, this filter could be used to exclude internal referral from your subreports.

Parameters

referer_match

This parameter contains the regular expression that will be used to filter out the referrer. Requests for which the referer field matches that regexp will be excluded from the subreports.

Defaults to `^-$`.

III. Lire Reference

Chapter 18. Installation Parameters

This chapter describes the various configuration variables that can be set when installing Lire. These can be set using options to **./configure** or by setting environment variables.

./configure parameters

--prefix

This option specifies where Lire will be installed.

Defaults to `/usr/local`.

--bindir

This option specifies where Lire's executables intended for users will be installed.

Defaults to `${prefix}/bin`.

--sysconfdir

This option specifies where Lire's configuration files will be installed. (Actually, they will be installed in a subdirectory named `lire`.)

Defaults to `${prefix}/etc`.

--libexecdir

This option specifies where Lire's internal executables and scripts will be installed. (Actually, they will be installed in a subdirectory of this one named `lire`.)

Defaults to `${prefix}/libexec`.

--sharedstatedir

This option specifies where Lire's data files will be installed. (Actually, they will be installed in a subdirectory of this one named `lire`.)

Defaults to `${prefix}/share`.

--mandir

This option specifies where Lire's man pages will be installed.

Defaults to `${prefix}/man`.

`--with-perl5libdir`

This option specifies where Lire's perl modules will be installed.

Defaults to `${prefix}/share/perl5`.

`--with-perl5archlibdir`

This option specifies where architecture dependent perl modules (like XML::Parser) will be installed.

Defaults to `${prefix}/lib/perl5`.

`--with-spooldir`

This option specifies the default value of `lr_spool_dir` which is the spool directory used by the responder. Unless you're running your own responder this variable is not relevant.

`--with-archivedir`

When you're archiving your reports and logs using the archive feature this sets the default value of `lr_archive_dir`.

Installation Environment Variables

Some environment variables can be set before running `./configure` to tune the installation process. This can be used to specify the locations of components which are installed but can't be found by `./configure` in "standard" locations. For example, you could pass the location of the DocBook DTD by running `./configure` as:

```
$ DBK_XML_DTD=/home/flacoste/xml/docbook-xml-4.1.2/docbookx.dtd \
./configure
```

The following list explains the purpose of each variable.

`PATHTOPERL`

Sets the path to the **perl** interpreter.

`PATHTOJADE`

Sets the path to the **jade** DSSSL interpreter.

`PATHTOPDFJADETEX`

Sets the path to the **pdfjagetex** command.

`PATHTOXSLTPROC`

Sets the path to the **xsltproc** XSLT processor.

`DBK_XML_DTD`

Sets the path to the DocBook XML Document Type Declaration. This should point to the XML V4.1.2 DTD.

DBK_XSL_STYLESHEETS

Sets the path to the directory which contains Norman Walsh's XSL stylesheets for DocBook. (This directory should contain subdirectories named `fo`, `html` or `xhtml`.)

DBK_DSSSL_STYLESHEETS

Sets the path to the directory which contains Norman Walsh's DSSSL stylesheets for DocBook. (This directory should contain a subdirectory named `print`.)

Chapter 19. Lire Logging and Error Messages

Logging

Lire can log its messages, and output them to either standard error (stderr) or to syslog using the **logger** program. Choosing between either one of them is done with the `LR_LOGGING` variable in `${sysconfdir}/lire/defaults` or `~/.lire/etc/defaults`. (See also **lr_run**.)

Log Messages

Each log message has a level, which is one of:

emerg

system is unusable

alert

action must be taken immediately

crit

critical conditions

err

error conditions

warning

warning conditions

notice

normal, but significant, condition

info

informational message

debug

debug-level message

See also `syslog(3)`.

A complete Lire message looks like

```
superservice service lr_tag program level message
```

where program is the name of the script producing the message. `lr_tag` is used to track different Lire jobs. E.g.

```
www apache lr_tag-20010826081801-31102 lr_log2mail notice storing \  
/tmp/lr_log2mail.apache.lr_tag-20010826081801-31102.report in \  
/var/lib/lire/data/report/ascii/www/apache/complete/example.com_20010826/2001081609224
```


Chapter 20. Lire Installation Layout

Service specific scripts should reside in *libexecdir/service*. Configuration in *sysconfdir/service*.