



/Logreport/

Lire

Wessel Dankers
Joost van Baal

November 13, 2002

Topics

- ▶ Introduction
- ▶ What is Lire
- ▶ Lire components
- ▶ The distiller
- ▶ The query engine
- ▶ The typesetter
- ▶ Configuration
- ▶ Questions

Introduction

Stichting LogReport Foundation

Missions:

- ▶ Develop, maintain and propagate tools;
- ▶ Develop and introduce product independent standard log file formats.

The LogReport Foundation is currently sponsored by the Stichting NLNet but is moving towards a commercial model based on services.

Work is currently done by three payed part-time developers, from the Netherlands and Canada, as well as volunteers from the US and Australia, a.o.

Introduction

Why Lire?

Log files are often ignored (or just deleted) but contain valuable information.

- ▶ Schedule hardware upgrades;
- ▶ Detecting anomalies;
- ▶ Traffic based accounting;
- ▶ Marketing data extraction.

Tools already exist, but only examine one product's log files.

What is Lire

A generalized log analysis framework:

- ▶ Handles log files of any type
 - webserver, database server, firewall, ...
- ▶ Handles log files of any product
 - iptables, Cisco PIX, Watchguard, ...
- ▶ Extracts any kind of information
 - Top-10, per period, per user, per session, ...
- ▶ Produces many kinds of output
 - ASCII, (X)HTML, PDF, DocBook XML, Excel, RTF, ...


Being a framework means it is easy to add support for new products.

What is Lire

- ▶ Programmed in highly modularized Perl (ca. 25,000 LOC)
- ▶ Expat 1.9.x and XML::Parser 2.29
- ▶ xsltproc 1.0.4
- ▶ DocBook XML DTD V4.1.2
- ▶ Jade and JadeT_EX (for PDF and RTF output)
- ▶ PassiveT_EX (a modern alternative for Jade)
- ▶ Norman Walsh XSL stylesheets for DocBook (for DocBook, PDF and (X)HTML)

Lire is Free Software.

Lire components



The distiller

Moderately simple Perl program to transform application specific log files into space delimited plain text files.

```
1016698091 10.1.3.4 - 200 15029 /doc/  
1016698092 10.1.3.4 - 200 148 /icons/blank.gif  
1016698092 10.1.3.4 - 200 216 /icons/back.gif  
1016698092 10.1.3.4 - 200 225 /icons/folder.gif  
1016698105 10.1.3.4 - 200 5042 /doc/local/
```

This format is called DLF (Distilled Log Format).

There is a distilled log format for each type of program (web server, database server, firewall, ...)

For example, Postfix, Sendmail and Exim share the same DLF format.

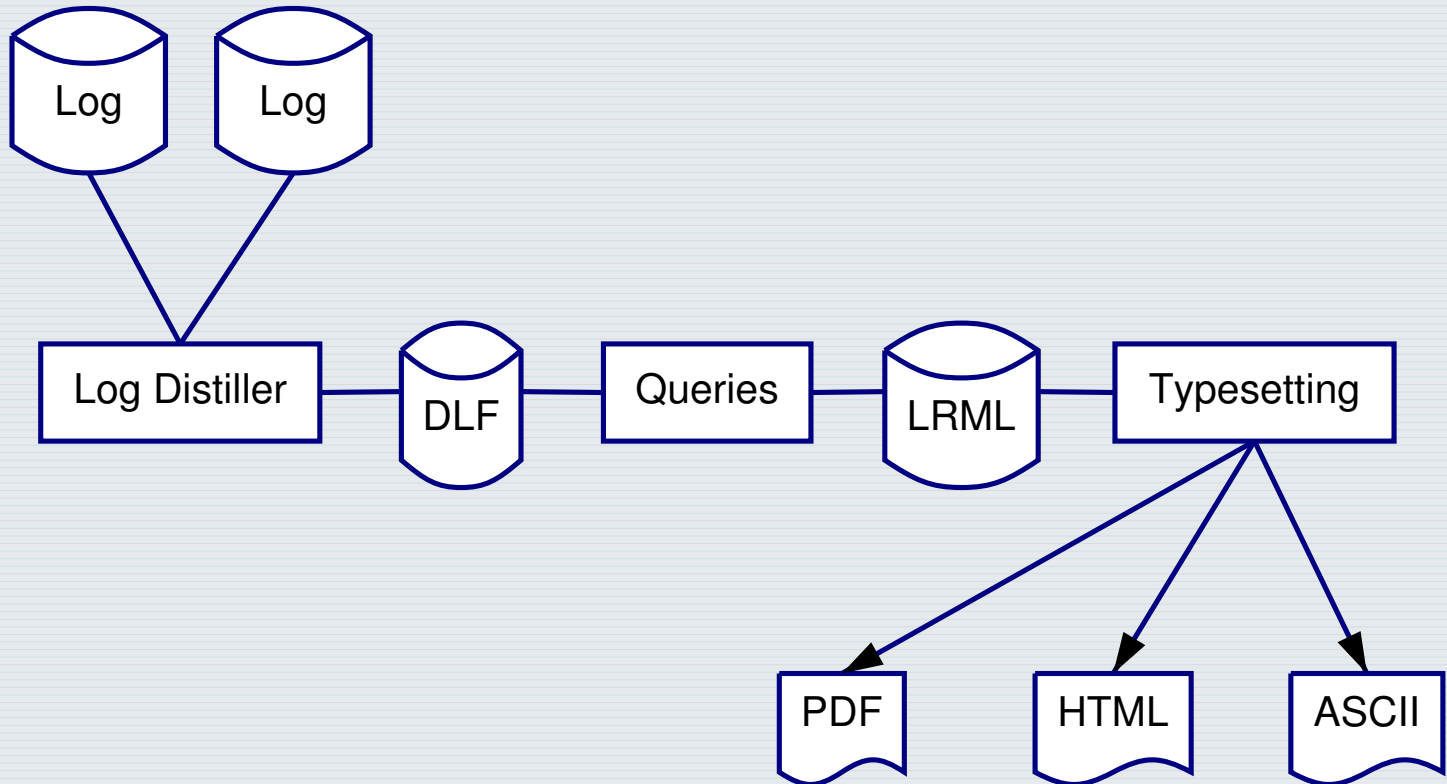
The distiller

XML files are used to specify what goes into a DLF file:

```
<lire:dlf-schema service="www"
  xmlns:lire="http://www.logreport.org/LDSML/">
  <lire:field name="time" type="timestamp">
    <lire:description>
      <para>The time of the request.</para>
    </lire:description>
  </lire:field>
  <lire:field name="client_host" type="hostname">
    <lire:description>
      <para>The hostname/ip address of the client</para>
    </lire:description>
  </lire:field>
</lire:dlf-schema>
```

Advantage: it's easy to change, add or delete fields.

Lire components



The query engine

Performs calculations on data which go into a report.

Consists of two phases:

- ▶ Analysis: extract derived information

For example, extract operating system information from common UserAgent values.

Merging data from external sources would take place here, too.

- ▶ Aggregation: calculate actual statistics to go into the report.

For example, the top 10 operating systems.

The results of many queries are concatenated into a report.

The query engine

In: a DLF file Out: fragments of LRML Report Markup Language.

```
<lire:report-spec id="top-last_page"  
  schema="www-user_session"  
  xmlns:lire="http://www.logreport.org/LRSML/">  
  
  <lire:param-spec>  
    <lire:param name="pages" type="int" default="10"/>  
  </lire:param-spec>  
  
  <lire:report-calc-spec>  
    <lire:group sort="-request_total" limit="$pages">  
      <lire:field name="last_page"/>  
      <lire:count name="request_total"/>  
    </lire:group>  
  </lire:report-calc-spec>  
</lire:report-spec>
```

The typesetter

Formats reports in different output formats:

- ▶ ASCII text:
Using a Perl module
- ▶ HTML pages:
Using xsltproc and Norman Walsh XSL stylesheets
- ▶ PDF documents: Using xsltproc, Jade and JadeT_EX

Intermediary XML reports can be stored and processed later (merging).

Configuration

Requirements:

- ▶ Must be “easily” accessible from several languages;
- ▶ Metadata for automatic documentation and for use by configuration tools;
- ▶ Types, defaults, complexity filtering, grouping, context;
- ▶ Multiple configuration frontends;
- ▶ Levels of configurability;
- ▶ Portability;
- ▶ No additional software requirements.

Configuration

Configuration consists of two parts:

- ▶ Specification
Denotes what options can occur and where
- ▶ Configuration
Holds actual values

Configuration

A sample specification:

```
<config-spec xmlns="http://www.logreport.org/LRCSML/">
  <boolean name="generate_images"/>
  <boolean name="keep_temp_dlf"/>

  <select name="target_user">
    <option>sysadmin</option>
    <option>manager</option>
  </select>

  <list name="path">
    <directory name="bin"/>
  </list>
</config-spec>
```

Configuration

A sample configuration:

```
<config xmlns="http://www.logreport.org/LRCML/">
  <global>
    <param name="generate_images">yes</param>
    <param name="keep_temp_dlf">yes</param>
    <param name="target_user">sysadmin</param>
    <param name="path">
      <param name="dir">/bin</param>
      <param name="dir">/usr/bin</param>
    </param>
  </global>
</config>
```

Contact us

Visit our website: <http://www.logreport.org/>

Commercial support: <http://www.logreport.com/>

Email us at logreport@logreport.org.

IRC: [#logreport](irc://irc.freenode.net/#logreport) on [irc.freenode.net](irc://irc.freenode.net).

Thank you for your attention!