

Linux on System z Web 2.0



Setting up a Web 2.0 stack on SUSE Linux Enterprise Server 10 SP2 December 2008

Linux on System z Web 2.0



Setting up a Web 2.0 stack on SUSE Linux Enterprise Server 10 SP2 December 2008

Note

Before using this document, be sure to read the information in “Notices” on page 65.

Second Edition – December 2008

This edition applies to SUSE Linux Enterprise Server 10 SP2 only.

© **Copyright International Business Machines Corporation 2008.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

| | |
|--|-----|
| Summary of changes | vii |
| December 2008, Second Edition | vii |
| Chapter 1. Introduction | 1 |
| What is Web 2.0? | 1 |
| The Web 2.0 stack components | 1 |
| Applications exploiting the Web 2.0 stack | 2 |
| Web 2.0 on Linux on System z | 3 |
| System requirements for the Web 2.0 stack. | 4 |
| Assumptions for this white paper. | 4 |
| Where to find this document | 4 |
| Chapter 2. Setup of Programming Languages | 5 |
| Installation of PHP | 5 |
| Installing additional PHP modules using PEAR and PECL | 5 |
| Installing database connectors | 5 |
| References. | 5 |
| Installation of Perl | 6 |
| Installing additional Perl libraries using CPAN | 6 |
| Installing database connectors | 6 |
| References. | 7 |
| Installation of Python | 7 |
| Installing additional Python packages using easy_install | 7 |
| Installing database connectors | 7 |
| References. | 8 |
| Installation of Ruby | 8 |
| Installing additional Ruby modules using RubyGems | 9 |
| Installing database connectors | 9 |
| References | 10 |
| Installation of Java | 10 |
| Installing database connectors | 10 |
| Installation of JRuby | 11 |
| References | 12 |
| Installation of Groovy | 12 |
| References | 13 |
| Installation of Jython | 13 |
| References | 14 |
| Chapter 3. Setup of a Database server | 15 |
| Setup of MySQL | 15 |
| Installation of MySQL | 15 |
| Lifecycle of MySQL | 15 |
| Basic configuration of MySQL | 17 |
| Setup of PostgreSQL | 18 |
| Installation of PostgreSQL. | 19 |
| Lifecycle of PostgreSQL | 19 |
| Basic configuration of PostgreSQL. | 21 |
| Chapter 4. Setup of Apache HTTP Server | 23 |
| Installation of Apache HTTP Server | 23 |
| Verification of the Apache HTTP Server installation | 23 |
| Lifecycle of Apache HTTP Server | 23 |
| Overview of Apache HTTP server modules | 25 |

| | |
|---|-----------|
| Adding support for PHP | 25 |
| 'Hello World!' example using mod_php | 26 |
| Adding support for Perl | 26 |
| 'Hello World!' example using mod_perl | 27 |
| Adding support for Python | 27 |
| 'Hello World!' example using mod_python Publisher handler | 28 |
| 'Hello World!' example using mod_python PSP handler | 29 |
| Adding support for Ruby | 29 |
| 'Hello World!' example using mod_ruby | 31 |
| 'Hello World!' example using eRuby | 32 |
| Adding the Web Application Firewall ModSecurity | 32 |
| Requirements | 32 |
| Installation | 33 |
| Configuration | 33 |
| References | 34 |
| Using SSL in Apache HTTP server | 34 |
| Chapter 5. Setup of Apache Tomcat | 37 |
| Installation of Apache Tomcat | 37 |
| Verifying the Apache Tomcat installation. | 37 |
| Important folders in Apache Tomcat | 37 |
| Lifecycle of Apache Tomcat | 38 |
| The Apache Tomcat administration tools | 39 |
| Adding support for JSP and Servlet API. | 40 |
| Installation of JSP and Servlet API libraries | 40 |
| 'Hello World!' example as JSP | 40 |
| 'Hello World!' example as Servlet | 41 |
| Using SSL in Apache Tomcat | 43 |
| Chapter 6. Setup of Ruby on Rails | 45 |
| Installation of Ruby on Rails | 45 |
| Creation of a Ruby on Rails example. | 45 |
| Deployment of Ruby on Rails applications | 45 |
| Chapter 7. Setup of Grails | 49 |
| Installation of Grails | 49 |
| Creation of a Grails example. | 50 |
| Chapter 8. Setup of Caches | 51 |
| Setup of Squid | 51 |
| Installation of Squid | 51 |
| Lifecycle of Squid | 51 |
| Basic configuration of Squid | 53 |
| Setup of memcached | 54 |
| Installation of memcached. | 54 |
| Lifecycle of memcached | 54 |
| References | 55 |
| Chapter 9. Setup of AJAX support using Dojo | 57 |
| Installation of the Dojo Toolkit | 57 |
| Example for using Dojo. | 57 |
| References | 59 |
| Appendix. Packages for the Web 2.0 stack | 61 |
| Notices | 65 |

| | |
|----------------------|----|
| Trademarks | 66 |
|----------------------|----|

Summary of changes

This section describes the changes made in this documentation compared to previous editions. This edition may also include minor corrections and editorial changes that are not identified.

December 2008, Second Edition

- The Chapter 2, “Setup of Programming Languages,” on page 5 has been added to separate the installation of programming languages from the usage in Apache HTTP server and Apache Tomcat server.
- The sections “Installation of JRuby” on page 11, “Installation of Groovy” on page 12 and “Installation of Jython” on page 13 have been added.
- The sections “Using SSL in Apache HTTP server” on page 34, “Using SSL in Apache Tomcat” on page 43 and “Adding the Web Application Firewall ModSecurity” on page 32 have been added.
- The Chapter 7, “Setup of Grails,” on page 49 has been added.
- The instructions in section “Installation of Ruby” on page 8 to install Ruby and related components have significantly changed. These changes were related to a dependency for the Ruby connector for DB2®.
- Cleanups in chapter and section titles have been made to get a better documentation structure.
- Several changes to the version numbers of packages have been incorporated.

Chapter 1. Introduction

In recent years, the traditional way to use the Internet has changed significantly. Web administrators who publish information about the Internet invite their community to contribute to the company's Web pages. Applications such as Wikis, Blogs, Content Management Systems and a couple of others provides this new functionality to the Web administrators.

The variety of available open source Web 2.0 applications requires that server components are set up correctly. These server components and their setup for SUSE Linux Enterprise Server 10 SP2 on Linux® on System z® are described in this white paper.

What is Web 2.0?

The key question which needs to be addressed is: What is Web 2.0?

Tim O'Reilly defines the Web 2.0 as ¹

Web 2.0 is the business revolution in the computer industry caused by the move to the Internet as platform, and an attempt to understand the rules for success on that new platform.

This new perspective changed the usage of the Internet significantly. In the past, system administrators were required to prevent users to access their private application interfaces from Web applications to avoid security issues. Of course, security issues are still required to be prevented, but APIs have been defined to allow reuse by 3rd party Web applications.

Data that was earlier only available for customers who paid for it is now opened up for public use. Who imagined ten years ago that satellite images of the whole world would be freely available and accessible by a well defined, public API?

Additionally, Web applications are getting more dynamic and more flexible due to new technologies. AJAX is one example to increase the performance of a Web 2.0 application by receiving data asynchronously from different services. Functionality of desktop application such as drag and drop are available through JavaScript™ libraries.

All of this created a new spirit around the Internet and this is called Web 2.0.

The Web 2.0 stack components

In general, the Web 2.0 stack is based on the software solution stack called LAMP (acronym for "Linux, Apache, MySQL, PHP (Perl or Python)"). Some enhancements to the LAMP stack come into the spotlight when the discussion moves to Web 2.0.

The following figure gives an overview about the Web 2.0 stack components

1. see <http://radar.oreilly.com/archives/2006/12/web-20-compact-definition-tryi.html>

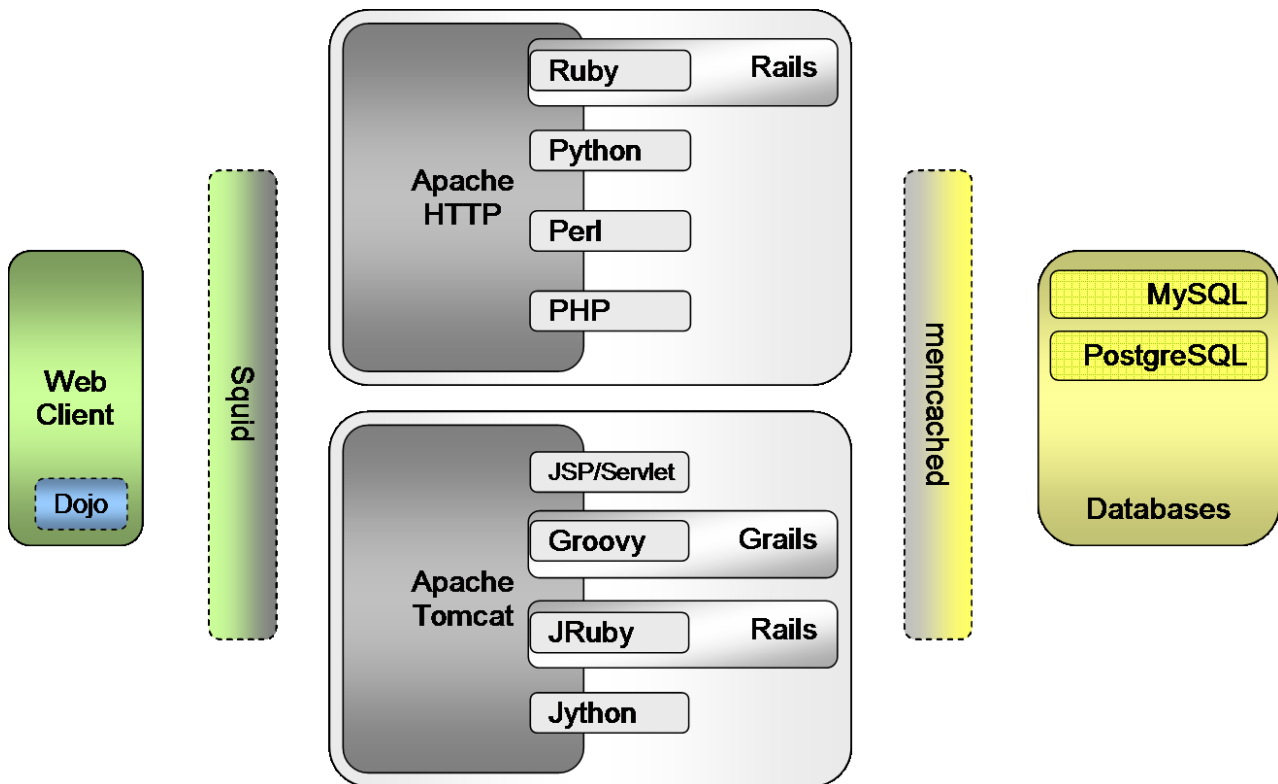


Figure 1. Web 2.0 stack components

Database

MySQL, PostgreSQL

Web server

Apache HTTP server, Apache Tomcat

Programming Languages

PHP, Perl, Python, Ruby, Java™, JRuby, Groovy, Jython

Frameworks

Dojo (AJAX support), Ruby on Rails, Grails

Caches

memcached, Squid

A system administrator does not have to install all Web 2.0 stack components at once. The selection is dependent on the Web 2.0 application and its functionality. For example if the Web 2.0 application offers information which is very static, it would make sense to establish a cache, but if the information is updated frequently then a cache is not the best choice.

The appendix gives an overview on all packages which are used in this document.

Applications exploiting the Web 2.0 stack

On the Internet, various open source Web 2.0 applications are available. In a follow-up white paper, the setup of some popular Web 2.0 applications running on Linux on System z will be described in detail.

The Web 2.0 stack which is described in this document enables a system administrator to choose the Web 2.0 application which fits best to the requirements. A list of very popular Web 2.0 applications is shown in the table below.

Table 1. Open Source Web 2.0 applications

| Application Type | Application name |
|----------------------------------|------------------|
| Wiki | MediaWiki |
| | MoinMoin |
| | XWiki |
| Blogs | WordPress |
| | Mephisto |
| | Movable Type |
| Content Management Systems (CMS) | Typo3 |
| | Drupal |
| | OpenCMS |
| | Joomla |
| e-Commerce | osCommerce |

Web 2.0 on Linux on System z

Why run the Web 2.0 stack on Linux on System z?

High availability

A System z enterprise mainframe offers a high available environment by default. This reduces system outages which are related to hardware issues.

Resource utilization

Rarely used Linux on System z virtual server machines running on z/VM[®] are swapped out and the resources are made available for other virtual machines. The swap in of a Linux on System z virtual server is hardly noticed by the user.

Vertical Scalability

By default, Linux on System z virtual servers running on z/VM are enabled for vertical scalability. This means a Linux on System z virtual server can be empowered with additional memory or processors on the fly.

Rapid deployment

If the vertical scalability is not sufficient, a new Linux on System z virtual server can be deployed within a couple of minutes.

Performance

Connections between virtual servers on System z can make use of the HiperSockets[™] technology to speed up communication.

Consolidation/TCO

The different Web 2.0 stack components might be running on dedicated server machines. Linux servers can be consolidated to run on one physical System z machine. This saves power, space in the computer center and reduces the administration effort.

System requirements for the Web 2.0 stack

The system requirements for a Web 2.0 stack setup are closely related to the requirements of the Web 2.0 application. An application which is based on a large database requires more system resources such as CPU and memory than an application which is running the application logic as client-side JavaScript.

The system requirements are investigated in more detail in a follow-up white paper which explains the Web 2.0 applications in more detail.

Assumptions for this white paper

System administrators who read this document should be familiar, how to setup YaST to include the SUSE Linux Enterprise Server 10 SP2 DVD image and the related SDK DVD image as installation sources. Additionally, the latest available security updates are applied to the system.

Where to find this document

The latest version of this document can be found on the developerWorks® Web site at http://www.ibm.com/developerworks/linux/linux390/distribution_hints.html

Chapter 2. Setup of Programming Languages

Web 2.0 applications make use of different programming languages. Beside of the core programming language, additional packages such as database connectors, are required to properly run Web 2.0 applications. This chapter explains the installation of programming languages and the setup of additional packages either included in SUSE Linux Enterprise Server 10 SP2 or taken from the Internet.

Installation of PHP

One of the general purposes for the development of PHP was to create a scripting language which suites Web development requirements perfectly. For this reason, PHP is supported in most Web server environments running on various operating systems.

To install PHP on SUSE Linux Enterprise Server 10 SP2, if this is not installed on the system already, process the following command:

```
# yast -i php5
```

Installing additional PHP modules using PEAR and PECL

Web 2.0 applications might require additional PHP modules to be installed, which are not part of SUSE Linux Enterprise Server 10 SP2. For this scenario the PEAR and PECL package managers can be used to add missing modules to the system in a comfortable way.

To install PECL and PEAR on SUSE Linux Enterprise Server 10 SP2, process the following command:

```
# yast -i php5-devel php5-pear
```

As an example, the installation of the memcached client library is processed with the following command:

```
# pecl install memcache
```

Installing database connectors

To install the MySQL connector for PHP on SUSE Linux Enterprise Server 10 SP2 use this command:

```
# yast -i php5-mysql
```

To install the PostgreSQL connector for PHP on SUSE Linux Enterprise Server 10 SP2 use this command:

```
# yast -i php5-pgsql
```

References

The following URLs provide more information about PHP:

- PHP project page, see <http://www.php.net/>

- PHP documentation, see <http://www.php.net/manual/en/>
- PEAR project page, see <http://pear.php.net/>

Installation of Perl

One of the most popular scripting languages is Perl. In general, Perl is based on other programming languages such as C, shell scripting, AWK and Lisp. This enables users to make use of a very powerful text manipulation engine.

To install Perl on SUSE Linux Enterprise Server 10 SP2, if this is not installed on the system already, process the following command:

```
# yast -i perl
```

Installing additional Perl libraries using CPAN

Web 2.0 applications might require additional Perl libraries to be installed, which are not part of SUSE Linux Enterprise Server 10 SP2. For this scenario the CPAN package manager can be used to add missing modules to the system in a comfortable way.

As an example, the installation of the memcached client library with CPAN is processed by using the following command:

```
# cpan Cache::Memcached
```

Note: When CPAN is started for the first time an initialization and configuration process is performed.

Installing database connectors

Perl provides database access based on a common module named DBI (Database Interface). To install this module on SUSE Linux Enterprise Server 10 SP2, use the following command:

```
# yast -i perl-DBI
```

To install the MySQL connector for Perl on SUSE Linux Enterprise Server 10 SP2 use the following command:

```
# yast -i perl-DBD-mysql
```

To install the PostgreSQL connector for Perl on SUSE Linux Enterprise Server 10 SP2 use the following command:

```
# yast -i pgperl
```

Note: The connector for PostgreSQL is available on the SUSE Linux Enterprise Server 10 SP2 SDK image. Therefore add the SDK image to the installation sources to run the given command properly.

References

The following URLs provide more information about Perl:

- Perl project page, see <http://www.perl.org>
- Perl documentation, see <http://perldoc.perl.org/perl.html>

Installation of Python

Python is another popular scripting language. Some key characteristics are that the language is very easy to learn and the code is required to be well structured which eases up the maintenance.

To install Python on SUSE Linux Enterprise Server 10 SP2, if this is not installed on the system already, process the following command:

```
# yast -i python
```

Installing additional Python packages using easy_install

Web 2.0 applications might require additional Python packages to be installed, which are not part of SUSE Linux Enterprise Server 10 SP2. For this scenario the `easy_install` package manager can be used to add missing modules to the system in a comfortable way.

To install the `easy_install` tool on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Download the `setuptools` for the Python version 2.4. While writing this documentation, the current version of `setuptools` is 0.6c9:

```
# wget -c http://pypi.python.org/packages/2.4/s/setuptools/setuptools-0.6c9-py2.4.egg
```

2. Create the new directory `/usr/local/lib64/python2.4/site-packages/`:

```
# mkdir -p /usr/local/lib64/python2.4/site-packages/
```

3. Run the downloaded file as a script to install the `setuptools`. This needs to be processed as user root:

```
# sh setuptools-0.6c9-py2.4.egg
```

As an example, the installation of the `memcached` client library with `easy_install` is processed by using the following command:

```
# easy_install python-memcached
```

Installing database connectors

To install the MySQL connector for Python on SUSE Linux Enterprise Server 10 SP2 use the following command:

```
# yast -i python-mysql
```

To install the PostgreSQL connector for Python on SUSE Linux Enterprise Server 10 SP2 use the following command:

```
# yast -i PyGreSQL
```

References

The following URLs provide more information about Python:

- Python project page, see <http://www.python.org>
- Python documentation, see <http://www.python.org/doc/>
- Python setuptools, see <http://pypi.python.org/pypi/setuptools/>

Installation of Ruby

Ruby is a object-oriented programming language. The language gained a large popularity boost in 2004 when the Web application framework “Ruby on Rails” was released. Additionally, implementations for virtual machines also became very popular. Such implementations are for example JRuby for the Java virtual machine or IronRuby for the .NET framework.

The number of third party libraries for Ruby is increasing day by day. These libraries, such as the IBM® DB2 driver, have special requirements to the Ruby programming language and therefore an upgrade of the Ruby installation might be required. While writing this documentation, the current version of Ruby is 1.8.7-p72.

To install Ruby on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Verify if there is already a Ruby installation available:

```
# rpm -qa|grep -i ruby
```

and

```
# ruby -v
```

Attention: Make sure to remove all Ruby related packages accordingly to their setup process before continuing with step 2.

2. Download the Ruby TAR file:

```
# wget -c ftp://ftp.ruby-lang.org/pub/ruby/1.8/ruby-1.8.7-p72.tar.gz
```

3. Extract the TAR file and change into the newly created folder:

```
# tar xzf ruby-1.8.7-p72.tar.gz
# cd ruby-1.8.7-p72
```

4. Setup the build environment and compile the Ruby sources:

```
# ./configure --enable-shared --enable-static
# make
```

5. Install Ruby. This needs to be processed as user root:

```
# make install
```

6. Verify the installation of Ruby:

```
# ruby -v
ruby 1.8.7 (2008-08-11 patchlevel 72) [s390x-linux]
```

Installing additional Ruby modules using RubyGems

RubyGems is used to manage libraries which extend the functionality of Ruby. While writing this documentation, the current version is 1.3.1.

To install RubyGems on SUSE Linux Enterprise Server 10 SP2, process the following steps:

1. Process the installation for Ruby as described in “Installation of Ruby” on page 8.
2. Download the RubyGems TAR file:

```
# wget -c http://rubyforge.org/frs/download.php/45905/rubygems-1.3.1.tgz
```

3. Extract the TAR file and change into the newly created folder:

```
# tar xzf rubygems-1.3.1.tgz
# cd rubygems-1.3.1
```

4. Run the 'setup.rb' script to install RubyGems. This needs to be processed as user root:

```
# ruby setup.rb
```

5. Verify the installation of RubyGems:

```
# gem -v
1.3.1
```

As an example, the installation of the memcached client library with RubyGems is processed by using the following command:

```
# gem install memcache-client
```

Installing database connectors

To install the MySQL connector for Ruby on SUSE Linux Enterprise Server 10 SP2, process the following steps

1. Process the installation for the MySQL development package:

```
# yast -i mysql-devel
```

2. Process the installation of the MySQL connector:

```
# gem install --remote mysql -- --with-mysql-config
```

To install the PostgreSQL connector for Ruby on SUSE Linux Enterprise Server 10 SP2, process the following steps:

1. Process the installation the PostgreSQL development package:

```
# yast -i postgresql-devel
```

2. Process the installation of the PostgreSQL connector:

```
# gem install --remote pg
```

References

The following URLs provide more information about Ruby and RubyGems:

- Ruby project page, see <http://www.ruby-lang.org/>
- RubyGems project page, see <http://www.rubygems.org/>

Installation of Java

Another popular programming language is Java. In this document Java 1.5.0 from IBM is used. The JRE and JDK are included in SUSE Linux Enterprise Server 10 SP2.

To install Java on SUSE Linux Enterprise Server 10 SP2, process the following steps:

1. Process the installation of Java:

```
# yast -i java-1_5_0-ibm java-1_5_0-ibm-devel jpackage-utils
```

2. Verify the installation of Java:

```
# java - version
java version "1.5.0"
Java(TM) 2 Runtime Environment, Standard Edition (build pxz64devifx-20080811a (SR8a))
IBM J9 VM (build 2.3, J2RE 1.5.0 IBM J9 2.3 Linux s390x-64 j9vmxz6423ifx-20080811 (JIT enabled))
J9VM - 20080809_21892_BHdSMr
JIT - 20080620_1845_r8
GC - 200806_19)
JCL - 20080811a
```

The output should look similar to the above.

Note: If the output reports a different version of Java, but the 1.5.0 version has been installed, it might be the case that the version 1.5.0 is not set to be the one used by the system. The tool `update-alternatives` is used to switch the Java version used by the system. Use `'update-alternatives --config java'` to select the right Java version.

3. Setup the `JAVA_HOME` environment variable. This environment variable is defined in the file `/etc/java/java.conf`. By default, the `JAVA_HOME` is commented out. To enable it, simply remove the `'#'` character at the beginning of the line. If there is no variable defined, add the following line to the file `/etc/java/java.conf`:

```
JAVA_HOME=/usr/lib/jvm/java/
```

Installing database connectors

To install the MySQL connector for Java, which is called 'MySQL Connector/J', on SUSE Linux Enterprise Server 10 SP2, process the following steps:

1. Download the TAR file from <http://dev.mysql.com/downloads/connector/j/5.1.html>. While writing this document, the current version is 5.1.7.
2. Extract the TAR file and change into the extracted folder:

```
# tar xzf mysql-connector-java-5.1.7.tar.gz
# cd mysql-connector-java-5.1.7
```

3. Copy the JAR file into the folder /usr/share/java:

```
# cp mysql-connector-java-5.1.7-bin.jar /usr/share/java/
```

4. Change the file permissions of the copied JAR file:

```
# chmod 644 /usr/share/java/mysql-connector-java-5.1.7-bin.jar
```

5. Create a link between the JAR file including the version number to a JAR file without the using a version number:

```
# ln -s /usr/share/java/mysql-connector-java-5.1.7-bin.jar /usr/share/java/mysql-connector-java.jar
```

To install the PostgreSQL connector for Java on SUSE Linux Enterprise Server 10 SP2 use this command:

```
# yast -i postgresql-jdbc
```

Installation of JRuby

JRuby is an implementation of the Ruby interpreter written in Java. While writing this documentation, the current version of JRuby is 1.1.5.

To install JRuby on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Process the installation of Java as described in “Installation of Java” on page 10.
2. Download the JRuby TAR file:

```
# wget -c http://dist.codehaus.org/jruby/jruby-bin-1.1.5.tar.gz
```

3. Extract the TAR file:

```
# tar xzf jruby-bin-1.1.5.tar.gz
```

4. Move the extracted folder to a proper place in the file system:

```
# mv jruby-1.1.5 /opt/jruby-1.1.5
```

5. Setup the PATH environment variable to get extended by the folder /opt/jruby-1.1.5/bin.

Create the new file /etc/profile.d/jruby.sh with the following content:

```
if ! echo ${PATH} | /bin/grep -q /opt/jruby-1.1.5/bin ; then
    PATH=${PATH}:/opt/jruby-1.1.5/bin
fi
```

Create the new file /etc/profile.d/jruby.csh with the following content:

```
if ( "${path}" !~ */opt/jruby-1.1.5/bin* ) then
    set path = ( /opt/jruby-1.1.5/bin $path )
endif
```

6. Load the environment variables:

```
# . /etc/profile.d/jruby.sh
```

7. Verify the installation of JRuby:

```
# jruby -v
jruby 1.1.5 (ruby 1.8.6 patchlevel 114) (2008-11-03 rev 7996) [s390x-java]
```

References

The following URLs provide more information about JRuby:

- JRuby project page, see <http://jruby.codehaus.org/>
- Getting started guide, see http://wiki.jruby.org/wiki/Getting_Started

Installation of Groovy

Groovy is a scripting language that provides seamless integration with Java. While writing this documentation, the current version of Groovy is 1.5.7.

To install Groovy on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Process the installation of Java as described in “Installation of Java” on page 10.
2. Download the Groovy archive:

```
# wget -c http://dist.groovy.codehaus.org/distributions/groovy-binary-1.5.7.zip
```

3. Extract the archive:

```
# unzip -qq groovy-binary-1.5.7.zip
```

4. Move the extracted folder to a proper place in the file system:

```
# mv groovy-1.5.7 /opt/groovy-1.5.7
```

5. Setup the GROOVY_HOME environment variable to point to the folder /opt/groovy-1.5.7 and the PATH environment variable to include the folder /opt/groovy-1.5.7/bin.

Create the new file /etc/profile.d/groovy.sh with the following content:

```
if ! echo ${PATH} | /bin/grep -q /opt/groovy-1.5.7/bin ; then
    PATH=${PATH}:/opt/groovy-1.5.7/bin
fi
export GROOVY_HOME=/opt/groovy-1.5.7
```

Create the new file /etc/profile.d/groovy.csh with the following content:

```
if ( "${path}" !~ */opt/groovy-1.5.7/bin* ) then
    set path = ( /opt/groovy-1.5.7/bin $path )
endif
setenv GROOVY_HOME "/opt/groovy-1.5.7"
```

6. Reload the environment variables:

```
# . /etc/profile.d/groovy.sh
```

7. Verify the installation of Groovy:

```
# groovy -v
Groovy Version: 1.5.7 JVM: 1.5.0
```

References

The following URLs provide more information about Groovy:

- Groovy project page, see <http://groovy.codehaus.org/>
- Getting started guide, see <http://groovy.codehaus.org/Getting+Started+Guide>

Installation of Jython

Jython is an implementation of the Python interpreter written in Java. While writing this documentation, the current version of Jython is 2.2.1.

To install Jython on SUSE Linux Enterprise Server 10 SP2, process the following steps:

1. Process the installation of Java as described in “Installation of Java” on page 10.
2. Download the Jython JAR file:

```
# wget -c http://downloads.sourceforge.net/jython/jython_installer-2.2.1.jar
```

3. Run the Jython installer using the following command:

```
# java -jar jython_installer-2.2.1.jar -c
```

The installer requests some information about the environment. For the installation target folder, use `/opt/jython-2.2.1`.

4. Setup the PATH environment variable to include the folder `/opt/jython-2.2.1`.

Create the new file `/etc/profile.d/jython.sh` with the following content:

```
if ! echo ${PATH} | /bin/grep -q /opt/jython-2.2.1 ; then
    PATH=${PATH}:/opt/jython-2.2.1
fi
```

Create the new file `/etc/profile.d/jython.csh` with the following content:

```
if ( "${path}" !~ */opt/jython-2.2.1* ) then
    set path = ( /opt/jython-2.2.1 $path )
endif
```

5. Reload the environment variables:

```
# . /etc/profile.d/jython.sh
```

6. Verify the installation of Jython:

```
# jython --version
Jython 2.2.1 on java1.5.0
```

References

The following URLs provide more information about Jython:

- Jython project page, see <http://www.jython.org/>

Chapter 3. Setup of a Database server

In general, Web applications deal with a lot of information which is displayed to the user. To get a structured and stable environment in place which guarantees consistency and persistency, databases are used. By default, SUSE Linux Enterprise Server 10 SP2 includes two popular databases: MySQL and PostgreSQL.

MySQL is one of the key components of the software solution stack called LAMP. Many Web 2.0 applications support both databases, so it is up to the system administrator which one to use.

Setup of MySQL

MySQL is a SQL database management system (DBMS), which has become one of the most popular open source database systems.

Some key characteristics - ease of use, reliability, security and performance at close to zero cost - increased the usage of MySQL, especially in Web applications.

Detailed information about MySQL is available at <http://www.mysql.org>

Installation of MySQL

The installation of MySQL is straight forward as the packages are included in SUSE Linux Enterprise Server 10 SP2. Install the MySQL server and the MySQL client packages by using the following command:

```
# yast -i mysql mysql-client
```

The mysql-client package includes a command line client for MySQL. This client is required for administration purposes of the database server.

To access the MySQL database from a Web application, a client for the related programming language is required. For the programming languages covered in this document, different client implementations are available. The setup of these clients is described in the related sections where the setup of the programming languages is described.

Verifying the MySQL installation

To verify MySQL has been installed correctly, run the following command:

```
# mysql -V  
mysql Ver 14.12 Distrib 5.0.26, for ibm-linux (s390x) using readline 5.1
```

The output should look similar to the above.

Lifecycle of MySQL

Lifecycle operations of Linux services like starting/stopping are basic functionality. In general, the lifecycle functionality is supported by a command line tool named service. The tool allows, dependent on the supported functionality of the service specific script, to start, stop and restart a server. Apart from that, the current status

of a server can be displayed. In the following, a walkthrough is shown for the MySQL server. This walkthrough assumes that this is the initial startup right after the packages have been installed.

1. Verifying the current status of the server

To get the current status of the MySQL server, run the following command:

```
# service mysql status
Checking for service MySQL:                                     unused
```

This indicates that the MySQL server is not started.

2. Starting the server for the first time

To start the MySQL server, run the following command:

```
# service mysql start
Creating/Updating MySQL privilege database...
Installing all prepared tables
Fill help tables
(.....)
Updating MySQL privilege database...
Fixing privilege tables...
Starting service MySQL                                         done
```

The output displays some initialization steps which indicate that the MySQL server has been started for the first time. After the startup is completed, verify the status of the MySQL server by running:

```
# service mysql status
Checking for service MySQL:                                     running
```

As shown, the MySQL server is running as expected.

3. Restarting the server

Another useful functionality is the restart where it is possible to stop and start the server with one single command. To restart the MySQL server, run the following command:

```
# service mysql restart
Shutting down service MySQL
Starting service MySQL                                         done
done
```

Once the restart completes, verify the status of the MySQL server by running:

```
# service mysql status
Checking for service MySQL:                                     running
```

As expected, the MySQL server is running again.

4. Stopping the server

To stop the MySQL server, run the following command:

```
# service mysql stop
Shutting down service MySQL                                     done
```

Again, the status can be verified by running:

```
# service mysql status
Checking for service MySQL: unused
```

As expected, the MySQL server is not running anymore.

To start the MySQL server at boot time, the command `chkconfig` is used. Decide first in which runlevel the MySQL server should start. Runlevels are used to coordinate the startup of services during boot time. In the following example, the runlevels 3 and 5 are used as both support networking:

```
# chkconfig --level 35 mysql on
```

To verify the setup use:

```
# chkconfig --list mysql
mysql          0:off  1:off  2:off  3:on   4:off  5:on   6:off
```

Now, the MySQL server starts during boot time of runlevels 3 and 5. To deactivate this behavior again, run:

```
# chkconfig mysql off
```

Again, to verify the setup use:

```
# chkconfig --list mysql
mysql          0:off  1:off  2:off  3:off  4:off  5:off  6:off
```

Basic configuration of MySQL

MySQL delivers several example configuration files in its packages. At least four different flavours of the server configuration are available. These additional configuration files are located in the folder `/usr/share/mysql` and are named:

- `my-small.cnf`
- `my-medium.cnf`
- `my-large.cnf`
- `my-huge.cnf`

The MySQL documentation offers additional information about which configuration file to choose for a given scenario. For example, the `my-small.cnf` file can be used as the configuration for the MySQL server. The following steps should be processed to activate the predefined configuration:

```
# cp /usr/share/mysql/my-small.cnf /etc/my.cnf
# chown root:root /etc/my.cnf
# chmod 644 /etc/my.cnf
```

One of the first issues when dealing with any common resource is to ensure its security. At the beginning the password for the database superuser is set properly by running the following commands:

```
# mysqladmin -u root password <'new-password'>
# mysqladmin -u root -h <your_host_name> password <'new-password'>
```

or by using the MySQL console:

```
# mysql -u root
mysql> SET PASSWORD FOR 'root'@'localhost' = PASSWORD('<new_password>');

mysql> SET PASSWORD FOR 'root'@'<your_host_name>' = PASSWORD('<new_password>');
mysql> quit
```

Note: As this is the superuser for all databases, use a strong password

The second recommended security measure is to remove anonymous accounts from the server. Access is granted only to those users specifically enabled:

```
# mysql -u root -p
mysql> DELETE FROM mysql.user WHERE User = '';
mysql> FLUSH PRIVILEGES;
mysql> SELECT Host, User FROM mysql.user;
mysql> quit
```

To get the list of users which have database access privileges, run the following command:

```
# mysql -u root -p
mysql> select user,host from mysql.user;
+-----+-----+
| user | host |
+-----+-----+
| root | localhost |
| root | machine.example.com |
+-----+-----+
mysql> quit
```

Note: Applications connecting to the MySQL server's databases should use an account with minimum privileges required for its actions. This helps to prevent malicious code from accessing other databases and data.

The configuration shown in this chapter does not cover all security aspects of MySQL. For more detailed information, refer to the MySQL documentation.

Setup of PostgreSQL

PostgreSQL is an advanced object-relational database management system.

Numerous features such as stored procedures, functions and triggers are included. All of these are blocks of code to be executed by the server and can be written in SQL or compatible languages such as C, C++, Java, PHP, Perl, Python, Ruby, and so on. PostgreSQL provides features such as MVCC (Multi-Version Concurrency Control), rules, triggers, indexes and so on.

Detailed information about PostgreSQL and the supported functionality is available at <http://www.postgres.org>

Installation of PostgreSQL

PostgreSQL has been included in various Linux distributions for some time. The installation process is straight forward. The database server can be installed by running the following command:

```
# yast -i postgresql postgresql-server
```

Next to the PostgreSQL server, the client which is required to maintain the databases is installed also.

To access the PostgreSQL database from a Web application, a client for the related programming language is required. For the programming languages mentioned in this document, different client implementations are available. The setup of these clients are described in the section where the setup of the programming languages is described.

Verifying the PostgreSQL installation

Checking whether the PostgreSQL package has been installed is done in the same way as with other servers. To retrieve the version of PostgreSQL the following command is performed:

```
# postgres -V
postgres (PostgreSQL) 8.1.11
```

The output should look similar to the above.

Lifecycle of PostgreSQL

Once these basic PostgreSQL packages are installed, lifecycle operations can be performed like starting/stopping the server. In the following, a walkthrough is shown for the PostgreSQL server. This walkthrough assumes that this is the initial startup right after the packages have been installed.

1. Verifying the current status of the server

To get the current status of the PostgreSQL server, run the following command:

```
# service postgresql status
Checking for PostgreSQL:                                     unused
```

This indicates that the PostgreSQL server is not started.

2. Starting the server for the first time

To start the PostgreSQL server, run the following command:

```
# service postgresql start
Initializing the PostgreSQL database at location /var/lib/pgsql/data  done
Starting PostgreSQL                                                  done
```

The output displays an initialization step which indicates that the PostgreSQL server has been started for the first time. After the startup is completed, verify the status of the PostgreSQL server by running:

```
# service postgresql status
Checking for PostgreSQL:                                     running
```

As shown, the PostgreSQL server is running as expected.

3. Restarting the server

Another useful functionality is the restart where it is possible to stop and start the server with one single command. To restart the PostgreSQL server, run the following command:

```
# service postgresql restart
Shutting down PostgreSQL:postmaster stopped
Starting PostgreSQL
```

done
done

Once the restart completes, verify the status of the PostgreSQL server by running:

```
# service postgresql status
Checking for PostgreSQL:
```

running

As expected, the PostgreSQL server is running again.

4. Stopping the server

To stop the PostgreSQL server, run the following command:

```
# service postgresql stop
Shutting down PostgreSQL:postmaster stopped
```

done

Again, the status can be verified by running:

```
# service postgresql status
Checking for PostgreSQL:
```

unused

As expected, the PostgreSQL server is not running anymore.

To start the PostgreSQL server at boot time, the command `chkconfig` is used. Decide first in which runlevel the PostgreSQL server should start. Runlevels are used to coordinate the startup of services during boot time. In the following example, the runlevels 3 and 5 are used as both support networking:

```
# chkconfig --level 35 postgresql on
```

To verify the setup use:

```
# chkconfig --list postgresql
postgresql          0:off 1:off 2:off 3:on  4:off 5:on  6:off
```

Now, the PostgreSQL server starts during boot time of runlevels 3 and 5. To deactivate this behavior again, run:

```
# chkconfig postgresql off
```

Again, to verify the setup use:

```
# chkconfig --list postgresql
postgresql                                0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

Basic configuration of PostgreSQL

PostgreSQL does not allow remote connections by default. To enable remote connections a parameter needs to be set in the file `/etc/sysconfig/postgresql`. This file governs optional PostgreSQL startup parameters and needs to be created if it does not already exist on the system.

In the file find the reference to `POSTGRES_OPTIONS`. This should be changed to set the `-i` parameter at startup. The example below shows an extract of the file with `POSTGRES_OPTIONS` set to `-i`:

```
## Path: Applications/PostgreSQL
## Description: The PostgreSQL Database System
## Type: string()
## Default: ""
## ServiceRestart: postgresql
#
# The options that are given to the PostgreSQL master daemon on startup.
# See the manual pages for postmaster and postgres for valid options.
#
# Don't put "-D datadir" here since it is set by the startup script
# based on the variable POSTGRES_DATADIR above.
#
POSTGRES_OPTIONS="-i"
```

Run the following command to restart the PostgreSQL server:

```
# service postgresql restart
```

The client authentication is controlled by a configuration file, which traditionally is named `pg_hba.conf`. This file is stored in the database folder `/var/lib/pgsql/data`.

One possible scenario is to allow connections from local and from a specific subnet. The following example shows how this works for the subnet `192.168.12.0/24`. Local users can connect only to their own databases (databases with the same name as their database user name), administrators can connect to all databases and client connections from a subnet of `192.168.12.x` can connect to the databases `db1` and `db2`. For all users, authentication is required:

| # | TYPE | DATABASE | USER | CIDR-ADDRESS | METHOD |
|---|-------|----------|---------|-----------------|--------|
| | local | sameuser | all | | md5 |
| | local | all | @admins | | md5 |
| | host | db1,db2 | all | 192.168.12.0/24 | md5 |

The configuration shown in this chapter does not cover all security aspects of PostgreSQL. For more detailed information, refer to the PostgreSQL documentation.

Chapter 4. Setup of Apache HTTP Server

The Apache HTTP Server, also known as Apache, is a secure, reliable and efficient Web server. It is highly configurable and extensible with lots of third party modules and Apache is available for a wide variety of operating systems including Linux, UNIX®, Mac OS X, Microsoft® Windows® and several others.

Apache became the Web server component of the popular software solution stack called LAMP. Extending this concept, the Apache HTTP Server is a key component of the defined Web 2.0 stack also.

The Apache HTTP server is well documented. A good reference is the Apache Web page at <http://httpd.apache.org>

Installation of Apache HTTP Server

Apache is one of the key components of the WorldWideWeb, and is included in SUSE Linux Enterprise Server 10 SP2. The installation is performed using the following command:

```
# yast -i apache2 apache2-doc apache2-example-pages apache2-prefork
```

This installs the Apache HTTP server with some documentation and an example page.

Verification of the Apache HTTP Server installation

To verify the installation of Apache, run the following command:

```
# apache2ctl -v
Server version: Apache/2.2.3
Server built: Sep 25 2008 10:55:57
```

The output should look similar to the above.

To connect with a Web browser to the Apache HTTP server, the server needs to be started. Use the following command to start the server:

```
# service apache2 start
```

Open a Web browser and enter the URL <http://<server-name>>. A Web page gets displayed indicating that the Apache HTTP server is running properly.

Note: If the system is protected by a firewall additional configuration might be necessary to allow client access to the Apache HTTP server

Lifecycle of Apache HTTP Server

Once these basic Apache HTTP server packages are installed, lifecycle actions can be performed like starting/stopping the server. In the following, a walkthrough is shown for the Apache HTTP server. This walkthrough assumes that this is the initial startup right after the packages have been installed and the server is not running.

1. Verifying the current status of the server

To get the current status of the Apache HTTP server, run the following command:

```
# service apache2 status
Checking for httpd2:                                     unused
```

This indicates that the Apache HTTP server is not started.

2. Starting the server for the first time

To start the Apache HTTP server, run the following command:

```
# service apache2 start
Starting httpd2 (prefork)                                done
```

After the startup is completed, verify the status of the Apache HTTP server by running:

```
# service apache2 status
Checking for httpd2:                                     running
```

As shown, the Apache HTTP server is running as expected.

3. Restarting the server

Another useful functionality is the restart where it is possible to stop and start the server with one single command. To restart the Apache HTTP server, run the following command:

```
# service apache2 restart
Shutting down httpd2 (waiting for all children to terminate)  done
Starting httpd2 (prefork)                                     done
```

Once the restart completes, verify the status of the Apache HTTP server by running:

```
# service apache2 status
Checking for httpd2:                                     running
```

As expected, the Apache HTTP server is running again.

4. Stopping the server

To stop the Apache HTTP server, run the following command:

```
# service apache2 stop
Shutting down httpd2 (waiting for all children to terminate)  done
```

Again, the status can be verified by running:

```
# service apache2 status
Checking for httpd2:                                     unused
```

As expected, the Apache HTTP server is not running anymore.

To start the Apache HTTP server at boot time, the command `chkconfig` is used. Decide first in which runlevel the Apache HTTP server should start. Runlevels are

used to coordinate the startup of services during boot time. In the following example, the runlevels 3 and 5 are used as both support networking:

```
# chkconfig --level 35 apache2 on
```

To verify the setup use:

```
# chkconfig --list apache2
apache2          0:off 1:off 2:off 3:on  4:off 5:on  6:off
```

Now, the Apache HTTP server starts during boot time of runlevels 3 and 5. To deactivate this behavior again, run:

```
# chkconfig apache2 off
```

Again, to verify the setup use:

```
# chkconfig --list apache2
apache2          0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

Overview of Apache HTTP server modules

The Apache server modules are software elements which extend Apache's functionality. There are modules available to add server-side programming language support, authentication schemes and a lot of additional functionality.

By default, the Apache HTTP server has already some modules loaded, statically and dynamically. To display the list of configured modules, run the following command:

```
# apache2ctl -t -D DUMP_MODULES
```

During the installation of Apache HTTP server extensions additional modules are installed, but those might not be enabled. To get a list of all available modules, run:

```
# a2enmod -l
```

A module is enabled by running:

```
# a2enmod <name_of_module>
```

A module is disabled by running:

```
# a2dismod <name_of_module>
```

Adding support for PHP

To add the PHP support for Apache HTTP on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Process the installation of PHP as described in "Installation of PHP" on page 5.
2. Install the Apache HTTP module for PHP support:

```
# yast -i apache2-mod_php5
```

3. Finally, restart the Apache HTTP server:

```
# service apache2 restart
```

During the installation process of PHP, one additional configuration file is added to the Apache server configuration. This file is located at `/etc/apache2/conf.d/php5.conf` and includes a definition for the module to configure the interpreter to do its work when a PHP page is requested.

'Hello World!' example using mod_php

To verify whether PHP is working correctly within Apache, create the new file `/srv/www/htdocs/hello-world.php` with the following content:

```
<html>
  <head>
    <title>Hello, World! (PHP)</title>
  </head>
  <body>
    <?php
      print 'Hello, World! (PHP)<br>';
      print 'The time on the server is ';
      print strftime('%H:%M:%S');
    ?>
  </body>
</html>
```

This newly created PHP script can be accessed by opening up a Web browser at the following URL:

- `http://<server-name>/hello-world.php`

A Web page is shown which displays 'Hello World!' along with the current time.

Adding support for Perl

To add the Perl support for Apache HTTP on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Process the installation of Perl as described in "Installation of Perl" on page 6.
2. Install the Apache HTTP module for Perl support:

```
# yast -i apache2-mod_perl
```

3. Enable the Apache HTTP module for Perl:

```
# a2enmod mod_perl
```

4. Finally, restart the Apache HTTP server:

```
# service apache2 restart
```

The installation of `apache2-mod_perl` integrates two new configuration files into the Apache HTTP setup

`/etc/apache2/mod_perl-startup.pl`

includes a list of additional Perl modules which are pre-loaded by Apache.

This mechanism speeds up the processing of Perl scripts since basic modules which are required for processing are already ready to use.

/etc/apache2/conf.d/mod_perl.conf

includes the setup for the Perl interpreter. In the default configuration, the folder /srv/www/cgi-bin is enabled to serve Perl scripts. This location is accessible at following two URLs:

- `http://<server-name>/perl/<script-name>`
- `http://<server-name>/cgi-perl/<script-name>`

Dependent on the URL used, different handlers are used to run the Perl script.

To activate the changes of the configuration files, restart the Apache HTTP server:

```
# service apache2 restart
```

'Hello World!' example using mod_perl

Based on the previously explained configuration, setup a 'Hello World!' example. Create the new file /srv/www/cgi-bin/hello-world.pl with the following content:

```
use CGI qw(:standard);
use CGI::Carp qw(fatalsToBrowser);

use strict;
use warnings;

my ($sec,$min,$hr) = localtime();

my $cgi = new CGI;

print $cgi->header(),
      $cgi->start_html({title=>"Hello World! - Perl"}),
      $cgi->p("Hello World! - Perl<br>
            The time on the server is $hr:$min:$sec"),
      $cgi->end_html();
```

This newly created Perl script can be accessed by opening up a Web browser at the following two URLs:

- `http://<server-name>/perl/hello-world.pl`
- `http://<server-name>/cgi-perl/hello-world.pl`

A Web page is shown which displays 'Hello World!' along with the current time.

Adding support for Python

To add the Python support for Apache HTTP on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Process the installation of Python as described in "Installation of Python" on page 7.
2. Install the Apache HTTP module for Python support:

```
# yast -i apache2-mod_python
```

3. Enable the Apache HTTP module for Python:

```
# a2enmod mod_python
```

4. Finally, restart the Apache HTTP server:

```
# service apache2 restart
```

The Python support for Apache HTTP server needs to be set up manually. Python offers three different standard handlers to run Python code:

Python Publisher handler

To enable the Publisher handler, a folder needs to be configured in the Apache HTTP server to which the publisher handler is assigned. Create the new file `/etc/apache2/conf.d/mod_python-publisher.conf` with the following content:

```
<Location /python>
    SetHandler mod_python
    PythonHandler mod_python.publisher
</Location>
```

This configuration enables the folder `/srv/www/htdocs/python` to serve Python scripts by making use of the Publisher handler.

Python PSP handler

To enable the Python Server Pages create the new file `/etc/apache2/conf.d/mod_python-psp.conf` with the following content:

```
<Directory /srv/www/htdocs>
    AddHandler mod_python .psp
    PythonHandler mod_python.psp
</Directory>
```

This configuration enables the folder `/srv/www/htdocs/` to deliver Python Server Pages using the PSP handler.

Python CGI handler

Using the Python CGI handler is not recommend. It is only intended to migrate legacy code away from CGI.

The folder `/srv/www/htdocs/python` does not exist by default, therefore create the folder by using the following command:

```
# mkdir -p /srv/www/htdocs/python
```

To activate the changes of the configuration files, restart the Apache HTTP server:

```
# service apache2 restart
```

'Hello World!' example using mod_python Publisher handler

Based on the previously explained configuration for the Publisher handler, setup a 'Hello World' example. Create the new file `/srv/www/htdocs/python/hello-world.py` with the following content:

```

from mod_python import apache
import time

def index(req):
    req.content_type = 'text/html'
    req.write('<html>')
    req.write('<head><title>Hello, World! (Python Publisher handler)</title></head>')
    req.write('<body>Hello, World! (Python Publisher handler)<br>')
    req.write('The time on the server is ')
    req.write(time.strftime("%H:%M:%S"))
    req.write('</body></html>')

```

This newly created Python script can be accessed by opening a Web browser at the following URLs:

- `http://<server-name>/python/hello-world`
- `http://<server-name>/python/hello-world/index`

A Web page is shown which displays 'Hello World!' along with the current time.

'Hello World!' example using mod_python PSP handler

Based on the previously explained configuration for the Python PSP handler, setup a 'Hello World' example. Create the new file `/srv/www/htdocs/hello-world.psp` is created with the following content:

```

<html>
  <head>
    <title>Hello, World! (Python PSP)</title>
  </head>
  <body>
    <% import time %>
    Hello, World! (Python PSP)<br>
    The time on the server is <%=time.strftime("%H:%M:%S")%>
  </body>
</html>

```

This newly created Python Server Page can be accessed by opening a Web browser at the following URL:

- `http://<server-name>/hello-world.psp`

A Web page is shown which displays 'Hello World!' along with the current time.

Adding support for Ruby

`mod_ruby` with `eruby` enables the execution of Ruby scripts by the Apache HTTP server. It embeds the Ruby interpreter to run scripts natively in the Apache HTTP server. The `mod_ruby` package is not part of SUSE Linux Enterprise Server 10 SP2 and needs to be installed manually. While writing this documentation, the current version of `mod_ruby` is 1.3.0.

To add the Ruby support for Apache HTTP on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Process the installation of Ruby as described in "Installation of Ruby" on page 8.
2. Install development packages which are required to compile `mod_ruby` and `eruby`:

```
# yast -i apache2-devel libapr1 libapr1-devel libapr-util1 libapr-util1-devel
```

3. Download the mod_ruby TAR file:

```
# wget -c http://www.modruby.net/archive/mod_ruby-1.3.0.tar.gz
```

4. Extract the TAR file and change into the extracted folder:

```
# tar xzf mod_ruby-1.3.0.tar.gz  
# cd mod_ruby-1.3.0
```

5. Set up the build environment:

```
# ./configure.rb --with-apr-includes=/usr/include/apr-1
```

6. Process the build of the mod_ruby source files:

```
# make
```

7. Process the installation. This needs to be processed as user root:

```
# make install
```

8. Download the eRuby TAR:

```
# wget -c http://www.modruby.net/archive/eruby-1.0.5.tar.gz
```

9. Extract the TAR file and change into the extracted folder:

```
# tar xzf eruby-1.0.5.tar.gz  
# cd eruby-1.0.5
```

10. Set up the build environment:

```
# ./configure.rb --with-apr-includes=/usr/include/apr-1
```

11. Process the build of the eRuby source files:

```
# make
```

12. Process the installation. This needs to be processed as user root:

```
# make install
```

13. Finally, restart the Apache HTTP server:

```
# service apache2 restart
```

The Ruby handlers must be set up manually. Create a new file `/etc/apache2/conf.d/modruby.conf` with the following content:

```

LoadModule ruby_module /usr/lib/apache2/mod_ruby.so
<IfModule mod_ruby.c>
  RubyRequire apache/ruby-run
  RubyRequire apache/eruby-run

  <Location /ruby>
    SetHandler ruby-object
    RubyHandler Apache::RubyRun.instance
    Options +ExecCGI
  </Location>

  <Files *.rbx>
    SetHandler ruby-object
    RubyHandler Apache::RubyRun.instance
  </Files>

  <Location /eruby>
    SetHandler ruby-object
    RubyHandler Apache::ERubyRun.instance
  </Location>

  <Files *.rhtml>
    SetHandler ruby-object
    RubyHandler Apache::eRubyRun.instance
  </Files>
</IfModule>

<Directory /var/www/html/ruby>
  Options ExecCGI
</Directory>

<Directory /var/www/html/eruby>
  Options ExecCGI
</Directory>

```

To activate the changes of the configuration files, restart the Apache HTTP server:

```
# service apache2 restart
```

'Hello World!' example using mod_ruby

Based on the previously explained environment, setup a 'Hello World!' example for CGI Ruby.

Create the new folder /srv/www/htdocs/ruby:

```
# mkdir /srv/www/htdocs/ruby
```

In this folder, create the new file /srv/www/htdocs/ruby/hello-world.rbx with the following content:

```

require "cgi"
cgi = CGI.new("html4")
cgi.out {
  cgi.html {
    cgi.head { cgi.title {"Hello, World! (Ruby)"} } +
    cgi.body {
      "Hello, World! (Ruby)<br>" +
      Time.now.strftime("The time on the server is %H:%M:%S")
    }
  }
}

```

This file requires the executable flag to be set:

```
# chmod +x /srv/www/htdocs/ruby/hello-world.rbx
```

The newly created Ruby CGI script can be accessed by opening a Web browser at the following URL:

- `http://<server-name>/ruby/hello-world.rbx`

A Web page is shown which displays 'Hello World!' along with the current time.

'Hello World!' example using eRuby

Based on the previously explained environment, setup a 'Hello World!' example for eRuby. Create the new folder `/srv/www/htdocs/eruby`:

```
# mkdir /srv/www/htdocs/eruby
```

In this folder, create the new file `/srv/www/htdocs/eruby/hello-world.rhtml` with the following content:

```
<html>
<head>
  <title>Hello, World! (eRuby)</title>
</head>
<body>
  <%= "Hello, World! (eRuby)" %><br>
  <%= Time.now.strftime("The time on the server is %H:%M:%S") %>
</body>
</html>
```

This newly created eRuby page can be accessed by opening a Web browser at the following URL:

- `http://<server-name>/hello-world.rhtml`

A Web page is shown which displays 'Hello World!' along with the current time.

Adding the Web Application Firewall ModSecurity

ModSecurity, <http://www.modsecurity.org>, is a Web Application Firewall (WAF) which enables a System administrator to implement an additional security layer without modifying the Web application itself. While writing this documentation, the current version is 2.5.7.

Requirements

To compile and run ModSecurity properly, several dependent packages must be installed. These are:

- `curl-devel`
- `pcre-devel`
- `libapr1-devel`
- `libapr-util1-devel`
- `apache2-devel`

All of these dependent packages are available in the SUSE Linux Enterprise Server 10 SP2 distribution and are installed, using the following command:

```
# yast -i curl-devel pcre-devel libapr1-devel libapr-util1-devel apache2-devel
```

Installation

To install ModSecurity on SUSE Linux Enterprise Server 10 SP2 follow these steps:

1. Download the ModSecurity TAR file:

```
# wget -c http://www.modsecurity.org/download/modsecurity-apache_2.5.7.tar.gz
```

2. Extract the TAR file:

```
# tar xzf modsecurity-apache_2.5.7.tar.gz
```

3. Setting up the build process and compiling the sources:

```
# cd modsecurity-apache_2.5.7/apache2
# ./configure
# make
```

4. Installation of ModSecurity:

```
# make install
```

5. Copy the configuration for the Apache HTTP server into the folder /etc/apache2/conf.d:

```
# cd ..
# cp modsecurity.conf-minimal /etc/apache2/conf.d/modsecurity.conf
```

Configuration

The configuration of ModSecurity includes the adjustment of configuration files for the Apache HTTP server and the integration of rules. Both are described in the following steps:

1. Add the LoadModule statement for the ModSecurity module configuration file. Add the following line to the /etc/apache2/conf.d/modsecurity.conf file:

```
LoadModule security2_module /usr/lib64/apache2/mod_security2.so
```

2. Adjust the values for the attributes SecDebugLog and SecAuditLog in the file /etc/apache2/conf.d/modsecurity.conf to be set as the following:

```
SecDebugLog /var/log/apache2/modsec_debug.log
SecAuditLog /var/log/apache2/modsec_audit.log
```

3. Enable the module mod_unique_id, if this is not the case. First, verify if the module is already loaded by using:

```
# apache2ctl -t -D DUMP_MODULES
```

If this the module is not loaded, process the following command:

```
a2enmod unique_id
```

4. Enable the Core rule set to be processed by ModSecurity. The 'Core rule set' is used to prevent the Web 2.0 applications for several known vulnerabilities. First, create the new folder where the rules are stored:

```
# mkdir -p /etc/apache2/sysconf.d/modsecurity
```

Copy the Core rules set from the extracted TAR file to the folder /etc/apache2/sysconf.d/modsecurity:

```
# cp rules/*.conf /etc/apache2/sysconf.d/modsecurity
```

Add the following line to the file /etc/apache2/httpd.conf:

```
Include sysconf.d/modsecurity/*.confmodsecurity/*.conf
```

Finally, restart the Apache HTTP server to enable the ModSecurity Web Application Firewall.

References

The following URLs provide more information about ModSecurity:

- ModSecurity project page, see <http://www.modsecurity.org>
- Further documentation, see <http://www.modsecurity.org/documentation/>

Using SSL in Apache HTTP server

The Secure Sockets Layer protocol (SSL) is used to encrypt the communication over TCP/IP. By default, the Apache HTTP server is not configured to support SSL. This section explains how to create a server certificate and configure the Apache HTTP server to support SSL.

To enable the SSL support for the Apache HTTP server, process the following steps:

1. Generate the server certificate by using the `gensslcert` script delivered with the Apache HTTP server package:

```
# gensslcert
```

This will generate a certificate together with a server key. Run '`gensslcert -h`' to get an overview about command line options to submit more specific information.

2. Create a virtual host for the SSL configuration. Therefore use the template to create a default configuration:

```
# cp /etc/apache2/vhosts.d/vhost-ssl.template /etc/apache2/vhosts.d/vhost-ssl.conf
```

By default, the previously generate certificate and server key are pre-configured in the configuration.

3. Verify if the module for SSL is loaded:

```
# a2enmod -l | grep -i ssl
```

If the SSL module is not already enabled, process:

```
# a2enmod ssl
```

4. Verify if the Apache HTTP server flags are set properly to make use of SSL. Open the file /etc/sysconfig/apache2 and look for the statement

'APACHE_SERVER_FLAGS'. This should include the directive 'SSL'. The following shows an example where the server flags are set only to SSL:

```
APACHE_SERVER_FLAGS="SSL"
```

5. Enable hardware support for encryption. Open the file `/etc/apache2/ssl-global.conf` and add the `SSLCryptoDevice` directive for `ibmca`:

```
<IfModule mod_ssl.c>
    SSLCryptoDevice ibmca
    ...
</IfModule>
```

Note: To enable the support for using the cryptographic adapter, make sure the Linux on System z guest is setup to access the required adapter. At IBM Redbooks a document is available which includes more detailed information about SSL support for Apache HTTP server - Using Cryptographic Adapters for Web Servers with Linux on IBM system z9® and zSeries®, see <http://www.redbooks.ibm.com/abstracts/redp4131.html>.

Chapter 5. Setup of Apache Tomcat

Apache Tomcat is a Servlet container which is compliant with the official specifications for the Java Servlet and JavaServer Pages technologies, providing an environment for Java code to run in cooperation with a Web server. The Java Servlet and JavaServer Pages specifications are developed by Sun under the Java Community Process. The two Java Specification Requests, JSR53 and JSR154 specify the Servlet API and the Java ServerPages API.

Further information about Apache Tomcat is available at the project Web page <http://tomcat.apache.org>.

Installation of Apache Tomcat

Java is a requirement to run the Apache Tomcat server. Process the installation of Java as described in “Installation of Java” on page 10.

When the correct Java version is available, install Apache Tomcat by running:

```
# yast -i tomcat5 tomcat5-admin-webapps tomcat5-webapps
```

This installs the Apache Tomcat server and some administrative Web applications which allow some basic configuration of the Apache Tomcat server. Additionally, several example Web applications for Apache Tomcat are installed.

Verifying the Apache Tomcat installation

To verify if the Apache Tomcat server is working properly, the server itself needs to be started. Therefore, run the following command:

```
# service tomcat5 start
```

Open a Web browser on the server where the Apache Tomcat server is running and open the URL <http://<server-name>:8080>. The default Apache Tomcat Web page is shown which offers some examples, documentation and administration pages.

Note: If the system is protected by a firewall additional configuration might be necessary to allow client access to the Apache Tomcat server

Important folders in Apache Tomcat

There is an environment variable with special importance, \$CATALINA_HOME. This environment variable is defined in /etc/tomcat5/tomcat5.conf and points to the root of the Apache Tomcat server installation – the default is /usr/share/tomcat5. These are some of the key Apache Tomcat folders, all relative to \$CATALINA_HOME:

- \$CATALINA_HOME/bin - Startup, shutdown, and other scripts.
- \$CATALINA_HOME/conf - Configuration files and related DTDs. The most important file is the server.xml which is the main configuration file for Apache Tomcat.
- \$CATALINA_HOME/logs - Default location of log files.
- \$CATALINA_HOME/webapps - This is where additional Web applications are installed.

Important folders for storing library files are:

- \$CATALINA_HOME/common/lib - Library files used by Apache Tomcat and Web applications.
- \$CATALINA_HOME/shared/lib - Library files used across Web applications only.

Lifecycle of Apache Tomcat

Once these basic Apache Tomcat server packages are installed, lifecycle actions can be performed like starting/stopping the server. In the following, a walkthrough is shown for the Apache Tomcat server. This walkthrough assumes that this is the initial startup right after the packages have been installed and the server is not running.

1. Verifying the current status of the server

To get the current status of the Apache Tomcat server, run the following command:

```
# service tomcat5 status
Checking for Tomcat (/srv/www/tomcat5/base/)          unused
```

This indicates that the Apache Tomcat server is not started.

2. Starting the server

To start the Apache Tomcat server, run the following command:

```
# service tomcat5 start
Starting Tomcat (/srv/www/tomcat5/base/)              done
```

After the startup is completed, verify the status of the Apache Tomcat server by running:

```
# service tomcat5 status
Checking for Tomcat (/srv/www/tomcat5/base/)          running
```

As shown, the Apache Tomcat server is running as expected.

3. Restarting the server

Another useful functionality is the restart where it is possible to stop and start the server with one single command. To restart the Apache Tomcat server, run the following command:

```
# service tomcat5 restart
Shutting down tomcat (/srv/www/tomcat5/base/)         done
Starting Tomcat (/srv/www/tomcat5/base/)              done
```

Once the restart completes, verify the status of the Apache Tomcat server by running:

```
# service tomcat5 status
Checking for Tomcat (/srv/www/tomcat5/base/)          running
```

As expected, the Apache Tomcat server is running again.

4. Stopping the server

To stop the Apache Tomcat server, run the following command:

```
# service tomcat5 stop
Shutting down tomcat (/srv/www/tomcat5/base/) done
```

Again, the status can be verified by running:

```
# service tomcat5 status
Checking for Tomcat (/srv/www/tomcat5/base/) unused
```

As expected, the Apache Tomcat server is not running anymore.

To start the Apache Tomcat server at boot time, the command `chkconfig` is used. Decide first in which runlevel the Apache Tomcat server should start. Runlevels are used to coordinate the startup of services during boot time. In the following example, the runlevels 3 and 5 are used as both support networking:

```
# chkconfig --level 35 tomcat5 on
```

To verify the setup use:

```
# chkconfig --list tomcat5
tomcat5          0:off 1:off 2:off 3:on  4:off 5:on  6:off
```

Now, the Apache Tomcat server starts during boot time of runlevels 3 and 5. To deactivate this behavior again, run:

```
# chkconfig tomcat5 off
```

Again, to verify the setup use:

```
# chkconfig --list tomcat5
tomcat5          0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

The Apache Tomcat administration tools

Two main administration tools are included in the package `tomcat5-admin-webapps` called `admin` and `manager`.

The `admin` tool is used to configure the Apache Tomcat server itself. It provides functionality for user, group and role management. Additionally the debug level can be modified for an application or the whole server.

The `manager` tool provides functionality for the Web application lifecycle. Web applications can be deployed, started, stopped and Web application specific options can be modified.

By default, user access to these tools is not configured and has to be granted manually. The installation of Apache Tomcat includes a file called `/etc/tomcat5/base/tomcat-users.xml` which is used to assign roles to groups or users. For example, the user `tomcat` is assigned to the roles `tomcat`, `admin` and `manager`. Another user `admin` is assigned to the role `admin` and the user `manager` to the role `manager`. Such a setup is shown in the following:

```
<?xml version='1.0' encoding='utf-8'?>
<tomcat-users>
  <role rolename="tomcat"/>
  <role rolename="admin"/>
  <role rolename="manager"/>
  <user username="tomcat" password="tomcat1" roles="tomcat,admin,manager"/>
  <user username="admin" password="tomcat2" roles="admin"/>
  <user username="manager" password="tomcat3" roles="manager"/>
</tomcat-users>
```

Afterward the two administration applications can be accessed by using the URL `http://<server-name>:8080/admin` and `http://<server-name>:8080/manager/html`.

Adding support for JSP and Servlet API

Installation of JSP and Servlet API libraries

The Apache Tomcat installation requires the JSP and Servlet libraries. Therefore an installation of additional packages is not required.

In general the JSP implementation classes and the JSP Standard Tag Library are included in the following packages:

- jakarta-taglibs-standard
- servletapi5

The content of these packages is included into the Apache Tomcat environment during their installation.

'Hello World!' example as JSP

The verification of the JSP installation is a bit more complex than the other verification processes. First, create a new folder where the JSP file is stored. This example uses the folder named `/usr/share/tomcat5/webapps/sample-jsp`:

```
# mkdir /usr/share/tomcat5/webapps/sample-jsp
```

Once this folder is created, create the new file `/usr/share/tomcat5/webapps/sample-jsp/hello.jsp` with the following content:

```
<%@page import="java.util.Date" %>
<%@page import="java.text.DateFormat" %>
<DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN">
<html>
  <head>
    <title>Hello, World! - JSP</title>
  </head>
  <body>
    <% String hwStr = "Hello World!\n"; %>
    <% String timeStr = "The time on the server is "; %>
    <% DateFormat currentTime = DateFormat.getInstance(DateFormat.FULL); %>
    <%= hwStr %><br>
    <%= timeStr + currentTime.format(new Date()) %>
  </body>
</html>
```

As well as the JSP file, create a new folder which includes information about the Web application:

```
# mkdir /usr/share/tomcat5/webapps/sample-jsp/WEB-INF
```

This folder includes one additional file - web.xml - which can define various settings for the Web application. To keep this example simple it only includes the application name. The file /usr/share/tomcat5/webapps/sample-jsp/WEB-INF/web.xml needs to be created with the following content:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE web-app
    PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.2//EN"
    "http://java.sun.com/dtd/web-app_2_2.dtd">

<web-app>

    <display-name>JSP - Hello, World!</display-name>

</web-app>
```

To complete the setup the new Web application needs to be deployed into the Apache Tomcat server. This is done by using the manager application. In the page at <http://<server-name>:8080/manager/html>, the application stored in the path sample-jsp needs to be started. Once the application is running, it can be accessed at <http://<server-name>:8080/sample-jsp/hello.jsp>.

Note: Depending on the setup of Apache Tomcat the Web application example might be auto-deployed. In this case, the Web application is activated once the Apache Tomcat server is restarted.

'Hello World!' example as Servlet

The verification of the Servlet follows the same procedure as the JSP example. First, create a new folder where the Servlet and related files are stored. This example uses the folder named /usr/share/tomcat5/webapps/sample-servlet:

```
# mkdir /usr/share/tomcat5/webapps/sample-servlet
# mkdir /usr/share/tomcat5/webapps/sample-servlet/WEB-INF
# mkdir /usr/share/tomcat5/webapps/sample-servlet/WEB-INF/classes
# mkdir /usr/share/tomcat5/webapps/sample-servlet/WEB-INF/classes/sample
```

The Servlet file is named /usr/share/tomcat5/webapps/sample-servlet/WEB-INF/classes/sample/HelloWorldServlet.java and includes the following content:

```

package sample;

import java.io.*;
import java.util.Date;
import java.text.*;

import javax.servlet.http.*;
import javax.servlet.*;

public class HelloWorldServlet extends HttpServlet {
    public void doGet (HttpServletRequest request, HttpServletResponse response)
        throws ServletException, IOException {

        PrintWriter out = response.getWriter();

        out.println("<html>");
        out.println("<head><title>Hello, World! - Servlet</title></head>");
        out.println("<body>");
        out.println("Hello, World!<br>");

        DateFormat currentTime = DateFormat.getTimeInstance(DateFormat.FULL);
        out.println("The time on the server is " + currentTime.format(new Date()));
        out.println("</body></html>");

        out.close();
    }
}

```

This Java source file needs to be compiled by using the following commands:

```

# cd /usr/share/tomcat5/webapps/sample-servlet/WEB-INF/classes
# javac -cp /usr/share/java/servletapi5.jar sample/HelloWorldServlet.java

```

In the folder /usr/share/tomcat5/webapps/sample-servlet/WEB-INF, create the new file web.xml with the following content:

```

<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE web-app
    PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.2//EN"
    "http://java.sun.com/dtd/web-app_2_2.dtd">

<web-app>
    <display-name>Servlet - Hello, World!</display-name>

    <servlet>
        <servlet-name>HelloWorld</servlet-name>
        <servlet-class>sample.HelloWorldServlet</servlet-class>
    </servlet>

    <servlet-mapping>
        <servlet-name>HelloWorld</servlet-name>
        <url-pattern>hello-servlet</url-pattern>
    </servlet-mapping>
</web-app>

```

To complete the setup, the new Web application needs to be deployed into the Apache Tomcat server. This is done by using the manager application. In the page at <http://<server-name>:8080/manager/html>, the application stored in the path sample-servlet needs to be started. Once the application is running, it can be accessed at <http://<server-name>:8080/sample-servlet/hello-servlet>.

Note: Dependent on the setup of Apache Tomcat, the Web application example might be auto-deployed. In this case, the Web application is activated once the Apache Tomcat server is restarted.

Using SSL in Apache Tomcat

The Secure Sockets Layer protocol (SSL) is used to encrypt the communication over TCP/IP. By default, the Apache Tomcat server is not configured to support SSL. This section explains how to create a server certificate and configure the Apache Tomcat server to support SSL.

To enable the SSL support for the Apache Tomcat server, process the following steps:

1. Enable the user tomcat to write content to the folder /usr/share/tomcat5:

```
# chown tomcat:tomcat /usr/share/tomcat5/
```

2. Becoming user tomcat:

```
# su - tomcat
```

3. Generate a keystore and together with a key:

```
# keytool -genkey -alias tomcat -keyalg RSA
```

Follow the instructions on the command line to enter proper information for the key generation.

4. As user root, open the file /usr/share/tomcat5/conf/server.xml and enable the connector for SSL. Uncomment the configuration similar to the following:

```
<Connector port="8443" maxHttpHeaderSize="8192"
  maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
  enableLookups="false" disableUploadTimeout="true"
  acceptCount="100" scheme="https" secure="true"
  clientAuth="false" sslProtocol="TLS" />
```

Adjust the connector for SSL with IBM Java specific settings. Set the 'sslProtocol' and the 'algorithm' attributes to the following values:

```
sslProtocol="SSL"
algorithm="IbmX509"
```

The previously created keystore is secured by a password. This password needs to be defined in the connector for SSL configuration. Set the 'keystorePass' attribute as follows:

```
keystorePass="<your-password-for-keystore>"
```

Chapter 6. Setup of Ruby on Rails

Ruby on Rails is a ruby-based framework for developing Web applications. It emphasizes "convention over configuration", is following the MVC (Model-View-Controller) pattern, making Web development faster and more efficient and Web applications easier to maintain.

To get more information about Ruby on Rails refer to the project Web page at <http://rubyonrails.org>.

Installation of Ruby on Rails

To install Ruby on Rails on SUSE Linux Enterprise Server 10 SP2 follow these steps

1. Setup Ruby as explained in chapter "Installation of Ruby" on page 8.
2. Process the installation of Ruby on Rails:

```
# gem install rails --version 2.2.2 --remote
```

Beside the installation of the current version of Ruby on Rails, all dependent packages are installed also. During writing this documentation, the current version of Ruby on Rails is 2.2.2.

3. Verify the Ruby on Rails installation:

```
# rails -v
Rails 2.2.2
```

Creation of a Ruby on Rails example

One way to run a Ruby on Rails application in a production environment is to start the default Ruby application server on port 80. This approach - while being sufficient for a small installation - has the drawback of the default HTTP port now being bound exclusively to one single Ruby on Rails application. The following example session shows how to create and run a Ruby on Rails application:

```
# service apache2 stop
# cd /srv/www/htdocs/
# rails www.railstest.example.com
....
# cd www.railstest.example.com
# RAILS_ENV=production ruby script/server -p 80
```

To verify the Ruby on Rails example, start a Web browser on the machine where the Ruby on Rails application server is running and direct to the URL <http://<server-name>>. A 'Welcome aboard' screen comes up with some information how to proceed further activities on the Web application.

Deployment of Ruby on Rails applications

By default, the Ruby on Rails application server is bound exclusively to a single port. This indicates that, if a Ruby on Rails application is bound to port 80, the Apache HTTP server is not able to make use of this port anymore. A number of deployment options have been established to address this issue, including: running

the Web application as a FastCGI process, using JRuby to deploy to a Java application server or forwarding requests from Apache or another Web server to a Ruby application server such as Mongrel.

The latter method is used in the configuration example given below: requests are forwarded by the Apache HTTP server to a single Mongrel backend server using `mod_proxy`.

For further information about Mongrel visit the project Web site at <http://mongrel.rubyforge.org>

To install Mongrel on SUSE Linux Enterprise Server 10 SP2 follow this step

1. Process the installation of Mongrel:

```
# gem install mongrel --remote
```

Beside the installation of the current version of Mongrel, all dependent packages are installed also. During writing this documentation, the current version is 1.1.5.

Then, the configuration for the Apache HTTP needs to be modified. Add a new virtual host by creating a file named `/etc/apache2/vhosts.d/railstest.conf` with the following content²:

```
ProxyRequests off
<VirtualHost *:80>
    ServerName www.railstest.example.com

    ProxyPass / http://127.0.0.1:8000/
    ProxyPassReverse / http://127.0.0.1:8000
    ProxyPreserveHost on
</VirtualHost>
```

Additionally, the Apache HTTP server needs to be setup to map requests for the new server name to the new virtual host. Therefore, the `NameVirtualHost` attribute needs to be configured in the file `/etc/apache2/listen.conf`. Since this is dependent on the individual Network setup, refer to the documentation of Apache HTTP server.

Afterward make sure `mod_proxy` and `proxy_http` are enabled:

```
# a2enmod mod_proxy
# a2enmod proxy_http
```

Start the Ruby on Rails application in production mode on port 8000 by executing the following commands:

```
# cd /srv/www/htdocs/www.railstest.example.com
# mongrel_rails start -d -p 8000 -e production
```

Note: When multiple Ruby on Rails Web applications are intended to run simultaneously on the same system, each Web application must be assigned to a unique port number. Therefore the configuration of the virtual host given in the `railstest.conf` file needs to be duplicated with another `ServerName`.

2. This configuration is taken from <http://mongrel.rubyforge.org/wiki/Apache> where additional information is available about the integration of Mongrel into Apache

Additionally the ports of the ProxyPass and ProxyPassReverse attributes have to match the port specified during startup of the Web application. Finally, restart Apache:

```
# service apache2 restart
```

Make sure the application is running by opening `http://<server-name>` in a browser. The default "Welcome aboard" page is now displayed.

Chapter 7. Setup of Grails

Installation of Grails

To install Grails on SUSE Linux Enterprise Server 10 SP2, process the following steps:

1. Process the installation of Java as described in “Installation of Java” on page 10.

2. Download the Grails TAR file:

```
# wget -c http://dist.codehaus.org/grails/grails-bin-1.0.4.tar.gz
```

3. Extract the TAR file:

```
# tar xzf grails-bin-1.0.4.tar.gz
```

4. Move the extracted folder to /opt/grails-1.0.4:

```
# mv grails-1.0.4 /opt/grails-1.0.4
```

5. Setup the GRAILS_HOME environment variable to point to /opt/grails-1.0.4 and the PATH environment variable to include the folder /opt/grails-1.0.4/bin.

Create the new file /etc/profile.d/grails.sh with the following content:

```
if ! echo ${PATH} | /bin/grep -q /opt/grails-1.0.4/bin ; then
    PATH=${PATH}:/opt/grails-1.0.4/bin
fi
if [ -z "$JAVA_HOME" ]; then
    . /etc/java/java.conf
    export JAVA_HOME
fi
export GRAILS_HOME=/opt/grails-1.0.4
```

Create the new file /etc/profile.d/grails.csh with the following content:

```
if ( "${path}" !~ */opt/grails-1.0.4/bin* ) then
    set path = ( /opt/grails-1.0.4/bin $path )
endif
if [ -z "$JAVA_HOME" ]; then
    . /etc/java/java.conf
    setend JAVA_HOME "$JAVA_HOME"
endif
setenv GRAILS_HOME "/opt/grails-1.0.4"
```

6. Reload the environment variables:

```
# . /etc/profile.d/grails.sh
```

7. Verify the installation of Grails:

```
# grails
Welcome to Grails 1.0.34 - http://grails.org/
Licensed under Apache Standard License 2.0
Grails home is set to: /opt/grails-1.0.4

No script name specified. Use 'grails help' for more info or 'grails interactive' to enter interactive mode
```

Creation of a Grails example

1. Stop the Apache Tomcat service to free port 8080:

```
# service tomcat5 stop
```

2. Create the Grails application:

```
# grails create-app grailstest
```

3. Change into the new created folder for the application 'grailstest':

```
# cd grailstest
```

4. Run the Grails application:

```
# grails run-app
```

To verify the Grails example, start a Web browser and direct to the URL `http://<server-name>:8080/grailstest`.

Chapter 8. Setup of Caches

Setup of Squid

Squid is a caching proxy for the Web supporting various protocols such as HTTP, HTTPS, FTP and more. It is mainly used for two different purposes, Web caching reducing bandwidth and response times on the client side, and speeding up the delivery of Web elements by caching frequently-repeated requests on the server side. Nevertheless, it fits perfectly with roles such as proxying Secure Sockets Layer (SSL) requests and caching of Domain Name Server (DNS) lookups, and perform transparent caching. Squid also supports a wide variety of caching protocols, such as Internet Cache Protocol, (ICP) the Hyper Text Caching Protocol, (HTCP) the Cache Array Routing Protocol (CARP), and the Web Cache Coordination Protocol (WCCP).

Squid has extensive granular access control mechanisms and allows monitoring of critical parameters via the Simple Network Management Protocol (SNMP).

To get more information about Squid refer to the documentation at <http://www.squid-cache.org>

Installation of Squid

The Squid server is also part of SUSE Linux Enterprise Server 10 SP2. Use the following command to install the package:

```
# yast -i squid
```

Note: The Squid server requires the network setup to provide a fully qualified domain name (FQDN). Review `/etc/hosts` to make sure that a FQDN is provided.

Verification of the Squid installation

To verify the installation of Squid, use the following command:

```
# squid -v
Squid Cache: Version 2.5.STABLE12
(...)
```

The output in this document shows the version number. From the output of the command, several configuration parameters and values are displayed also.

Lifecycle of Squid

Once these basic Squid server packages are installed, lifecycle actions can be performed like starting/stopping the server. In the following, a walkthrough is shown for the Squid server. This walkthrough assumes that this is the initial startup right after the packages have been installed.

1. Verifying the current status of the server

To get the current status of the Squid server, run the following command:

```
# service squid status
Checking for WWW-proxy squid                                unused
```

The status unused indicates that the Squid server is not started.

2. Starting the server for the first time

To start the Squid server, run the following command:

```
# service squid start
Starting WWW-proxy squid (/var/cache/squid) done
```

After the startup is completed, verify the status of the Squid server by running:

```
# service squid status
Checking for WWW-proxy squid running
```

As shown, the Squid server is running as expected.

3. Restarting the server

Another useful functionality is the restart where it is possible to stop and start the server with one single command. To restart the Squid server, run the following command:

```
# service squid restart
Shutting down WWW-proxy squid done
Starting WWW-proxy squid (/var/cache/squid) done
```

Once the restart completes, verify the status of the Squid server by running:

```
# service squid status
Checking for WWW-proxy squid unused
```

As expected, the Squid server is running again.

4. Stopping the server

To stop the Squid server, run the following command:

```
# service squid stop
Shutting down WWW-proxy squid done
```

Again, the status can be verified by running:

```
# service squid status
Checking for WWW-proxy squid unused
```

As expected, the Squid server is not running anymore.

To start the Squid server at boot time, the command `chkconfig` is used. Decide first in which runlevel the Squid server should start. Runlevels are used to coordinate the startup of services during boot time. In the following example, the runlevels 3 and 5 are used as both support networking:

```
# chkconfig --level 35 squid on
```

To verify the setup use:

```
# chkconfig --list squid
squid          0:off  1:off  2:off  3:on   4:off  5:on   6:off
```

Now, the Squid server starts during boot time of runlevels 3 and 5. To deactivate this behavior again, run:

```
# chkconfig squid off
```

Again, to verify the setup use:

```
# chkconfig --list squid
squid          0:off  1:off  2:off  3:off  4:off  5:off  6:off
```

Basic configuration of Squid

In SUSE Linux Enterprise Server 10 SP2, a version of Squid 2.5 is included. The documentation for this version is delivered within the Squid RPM package. It includes several examples of possible configurations and is stored in `/usr/share/doc/packages/squid`.

To cover all the configuration scenarios is beyond the scope of this white paper. As an example for a Squid configuration, the setup for a reverse proxy is given. Refer to the documentation for additional information.

Configuring the Squid server as a reverse proxy

Reverse proxy cache, also known as Web server acceleration, is a method of reducing the load on a busy Web server by using a Web cache between the server and the Internet. Another benefit is improved security. Additionally it is one of many ways to improve the scalability without increasing the maintenance of the server too much. A good use case of a reverse proxy is to reduce the workload on a Web server that provides both static and dynamic content. The static content can be cached on the reverse proxy while the Web server is freed up to better handle the dynamic content.

In the scenario where the Web server is running on a different machine, the configuration of the Squid server in `/etc/squid/squid.conf` looks like the following:

```
http_port 80 # Port of Squid proxy
httpd_accel_host <your_webservers_ip> # IP address of Web server
httpd_accel_port 80 # Port of Web server
httpd_accel_single_host on # Forward uncached requests to single host
httpd_accel_with_proxy on
httpd_accel_uses_host_header off
```

If the Web server runs on the same machine as the Squid server is running, the Web server must be re-configured to run on a different port than 80, for example 81. The reason is that clients connect to the Squid server which acts between the clients and the Web server. Therefore, the configuration in `/etc/squid/squid.conf` needs to be modified to redirect requests to port 81 of the local machine:

```
http_port 80 # Port of Squid proxy
httpd_accel_host localhost # IP address of Web server
httpd_accel_port 81 # Port of Web server
httpd_accel_single_host on # Forward uncached requests to single host
httpd_accel_with_proxy on
httpd_accel_uses_host_header off
```

Setup of memcached

Memcached is a distributed memory system for caching purposes. In general it is used to speed up communication between a Web application and a database. The result is to reduce response time for highly frequented content of Web pages and lesser load of the database server.

There are several client APIs available to access the memcached server. All of the programming languages described in this book, Java, PHP, Perl, Python and Ruby, are supported with a memcached client API.

Various public Web sites are available which make use of memcached such as SourceForge, Wikipedia, YouTube, Facebook and many others.

More information is available on the memcached project Web page at <http://www.danga.com/memcached>.

Installation of memcached

Memcached, <http://www.danga.com/memcached/download.bml>, is not part of SUSE Linux Enterprise Server 10 SP2. At the moment, packages for memcached are only available from the Internet. At time of writing, the latest version is 1.2.6.

To install memcached on SUSE Linux Enterprise Server 10 SP2 process the following steps:

1. Download the memcached TAR file:

```
# wget -c http://www.danga.com/memcached/dist/memcached-1.2.6.tar.gz
```

2. Extract the TAR file and change into the extracted folder:

```
# tar xzf memcached-1.2.6.tar.gz
# cd memcached-1.2.6
```

3. Compile the memcached sources:

```
# ./configure --prefix="" && make
```

4. Install memcached:

```
# make install
```

5. Verify the compiled memcached:

```
# make test
```

Lifecycle of memcached

The memcached service can be started up with several command line options.

Memcached daemons can be started on as many spare machines as required. For example to start memcached as a daemon process using 2GB of memory and listening on IP 10.0.0.40 on port 11211. The user name needs to be submitted only for the case when running as root. Start memcached by running the following command:

```
# memcached -d -m 2048 -l 10.0.0.40 -p 11211 -u root
```

To get an overview for all command line options, refer to the memcached documentation.

Note: Memcached lacks authentication and security features, meaning it should only be used on servers with an appropriate firewall set up. By default, memcached uses the port 11211.

The memcached daemon process can be stopped by running the following command:

```
# killproc memcached
```

References

The following URLs provide more information about memcached:

- memcached project page, see <http://www.danga.com/memcached/>
- Documentation, see <http://www.socialtext.net/memcached/index.cgi>

Chapter 9. Setup of AJAX support using Dojo

Web 2.0 applications are getting more flexible and more dynamic due to the enablement of functionality such as "Asynchronous JavaScript and XML" (AJAX) or "Drag and Drop". Several frameworks have been created to offer a set of these functionality in a bundled package.

The Dojo toolkit, <http://dojotoolkit.org>, is a very popular framework which offers a lot of support to the developer of Web 2.0 applications.

Installation of the Dojo Toolkit

Right now there are no RPM packages available, but a tar file is available from the Dojo project Web site at <http://dojotoolkit.org>. While writing this documentation, the current version is 1.2.3.

To install the Dojo Toolkit on SUSE Linux Enterprise Server 10 SP2 follow these steps:

1. Download the Dojo Toolkit TAR file:

```
# wget -c http://download.dojotoolkit.org/release-1.2.0/dojo-release-1.2.3.tar.gz
```

2. Extract the TAR file:

```
# tar xzf dojo-release-1.2.3.tar.gz
```

3. Move the extracted folder into a folder which the Apache Web server can access:

```
# mv dojo-release-1.2.3 /srv/www/htdocs
```

Example for using Dojo

This section shows how to create an example which displays the current time, updated every second.

The Apache HTTP server needs to be setup to support PHP and Dojo as previously described. PHP scripts require the executable flag to be set.

To give an example, two files need to be created. The first, `time.html`, should be placed in the `/srv/www/htdocs` folder. It should include the following source code:

```

<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
<title>Dojo example: AJAX clock</title>

<script type="text/javascript" src="dojo-release-1.2.3/dojo/dojo.js"
    djConfig="parseOnLoad: true"></script>

<style type="text/css">
    @import "dojo-release-1.2.3/dijit/themes/tundra/tundra.css";
    @import "dojo-release-1.2.3/dojo/resources/dojo.css";
</style>

<style type="text/css">
#error, #main {
    margin: auto;
    margin-top: 120px;
    text-align: center
}
#error {
    font-size: 120%;
}
#main input,button {
    font-size: 400%; width: 250px
}

</style>

<script type="text/javascript">
dojo.require("dojo.parser");
dojo.require("dojox.timing");
dojo.require("dijit.form.TextBox");
dojo.require("dijit.form.Button");

var timer;

var startClock = function() {
    timer.start();
}
var stopClock = function() {
    timer.stop();
}

var getCurrentTime = function() {
    console.debug("Timer ticked: Requesting time from server")

    // Performs the AJAX request to the URL specified;
    // After the request is sent and the response is received,
    // the load event is triggered, which in turn sets the new
    // value of the clock widget.
    dojo.xhrGet( {
        url: "http://localhost/time.php",
        handleAs: "text",
        timeout: 4000,

        load: function(response) {
            dojo.byId("error").innerHTML = "";
            dijit.byId("clock").setValue(response);
        },
        error: function(response) {
            dojo.byId("error").innerHTML = response;
        }
    } );
}

```

```

dojo.addOnLoad(function() {
    // Create a new timer that fires a tick every second
    timer = new dojox.timing.Timer(1000);
    // Every tick will lead to updating the text box with the current
    //time value from the server
    timer.onTick = getCurrentTime;

    dojo.connect(dojo.byId("clock-start"), 'onclick', startClock);
    dojo.connect(dojo.byId("clock-stop"), 'onclick', stopClock);

});
</script>
</head>

<body class="tundra">
    <div id="error"></div>

    <div id="main">
        <input id="clock" dojoType="dijit.form.TextBox" style="" value="" /><br/>
        <button id="clock-start" dojoType="dijit.form.Button">Start the clock</button><br/>
        <button id="clock-stop" dojoType="dijit.form.Button">Stop the clock</button>
    </div>
</body>

```

The second file, `time.php`, is stored in the `/srv/www/htdocs` folder also. It includes the following source code:

```

<?php
    // Returns the server time as HH:MM:SS
    echo (date("H").":".date("i").":".date("s"));
?>

```

To run the example, a Web browser needs to be opened on the machine where the Web server is running and pointed to the URL `http://<server-name>/time.html`

References

The following URLs provide more information about Dojo Toolkit:

- Dojo project page, see <http://dojotoolkit.org/>
- Getting started guide, see <http://sitepen.com/labs/guides/?guide=DojoQuickStart>

Appendix. Packages for the Web 2.0 stack

The following table includes an overview of all packages used in this document

Table 2. Package overview of Web 2.0 stack in SUSE Linux Enterprise Server 10 SP2

| Component | Package Name | Version | Source |
|--------------------|-----------------------|--------------|-------------------|
| MySQL | mysql | 5.0.26-12.22 | SLES10 SP2 |
| | mysql-client | 5.0.26-12.22 | SLES10 SP2 |
| | mysql-devel | 5.0.26-12.22 | SLES10 SP2 |
| | | | |
| PostgreSQL | postgresql | 8.1.11-0.2 | SLES10 SP2 |
| | postgresql-server | 8.1.11-0.2 | SLES10 SP2 |
| | postgresql-devel | 8.1.11-0.2 | SLES10 SP2 |
| | | | |
| Apache HTTP server | apache2 | 2.2.3-16.19 | SLES10 SP2 |
| | apache2-devel | 2.2.3-16.19 | SLES10 SP2 |
| | apache2-doc | 2.2.3-16.19 | SLES10 SP2 |
| | apache2-example-pages | 2.2.3-16.19 | SLES10 SP2 |
| | apache2-prefork | 2.2.3-16.19 | SLES10 SP2 |
| | apache2-mod_php5 | 5.2.5-9.9 | SLES10 SP2 |
| | apache2-mod_perl | 2.0.2-14.2 | SLES10 SP2 |
| | apache2-mod_python | 3.1.3-60.9 | SLES10 SP2 |
| | mod_ruby | 1.3.0 | Internet Download |
| | mod_security2 | 2.5.7 | Internet download |
| | libpar1 | 1.2.2-13.2 | SLES10 SP2 |
| | libapr1-devel | 1.2.2-13.2 | SLES10 SP2 |
| | libapr-util1 | 1.2.2-13.2 | SLES10 SP2 |
| | libapr-util1-devel | 1.2.2-13.2 | SLES10 SP2 |
| | | | |
| Apache Tomcat | tomcat5 | 5.0.30-27.32 | SLES10 SP2 |
| | tomcat5-webapps | 5.0.30-27.32 | SLES10 SP2 |
| | tomcat5-admin-webapps | 5.0.30-27.32 | SLES10 SP2 |
| | | | |
| PHP | php5 | 5.2.5-9.9 | SLES10 SP2 |
| | php5-devel | 5.2.5-9.9 | SLES10 SP2 |
| | php5-pear | 5.2.5-9.9 | SLES10 SP2 |
| | php5-mysql | 5.2.5-9.9 | SLES10 SP2 |
| | php5-pgsql | 5.2.5-9.9 | SLES10 SP2 |
| | | | |
| Perl | perl | 5.8.8-14.10 | SLES10 SP2 |
| | perl-DBI | 1.50-13.2 | SLES10 SP2 |

Table 2. Package overview of Web 2.0 stack in SUSE Linux Enterprise Server 10 SP2 (continued)

| Component | Package Name | Version | Source |
|-----------------|-----------------------------|----------------|-------------------|
| | perl-DBD-mysql | 3.0002-15.2 | SLES10 SP2 |
| | pgperl | 2.1.11-16.2 | SLES10 SDK |
| | | | |
| Python | python | 2.4.2-18.22 | SLES10 SP2 |
| | easy_install/ setuptools | 0.6c9 | Internet download |
| | python-mysql | 1.2.0-17.2 | SLES10 SP2 |
| | PyGreSQL | 3.7-14.2 | SLES10 SP2 |
| | | | |
| Java | java-1_5_0-ibm | 1.5.0_sr8a-1.1 | SLES10 SP2 |
| | java-1_5_0-ibm-devel | 1.5.0_sr8a-1.1 | SLES10 SP2 |
| | jpackage-utils | 1.6.3-18.4 | SLES10 SP2 |
| | mysql-connector-java | 5.1.7 | Internet download |
| | postgresql-jdbc | 8.1-12.2 | SLES10 SP2 |
| | | | |
| Ruby | ruby | 1.8.7-p72 | Internet download |
| | rubygems | 1.3.1 | Internet download |
| | rubygem-rake | 0.8.3 | Internet download |
| | rubygem-rails | 2.2.2 | Internet download |
| | rubygem-activesupport | 2.2.2 | Internet download |
| | rubygem-actionpack | 2.2.2 | Internet download |
| | rubygem-actionmailer | 2.2.2 | Internet download |
| | rubygem-activerecord | 2.2.2 | Internet download |
| | rubygem-activeresource | 2.2.2 | Internet download |
| | rubygem-daemons | 1.0.10 | Internet download |
| | rubygem-fastthread | 1.0.1 | Internet download |
| | rubygem-gem_plugin | 0.2.3 | Internet download |
| | rubygem-mongrel | 1.1.5 | Internet download |
| | rubygem-mysql | 2.7 | Internet download |
| | rubygem-pg | 0.7.9 | Internet download |
| | eruby | 1.0.5 | Internet download |
| | | | |
| JRuby | JRuby | 1.1.5 | Internet download |
| | | | |
| Groovy / Grails | Groovy | 1.5.7 | Internet download |
| | Grails | 1.0.4 | Internet download |
| | | | |
| Jython | Jython | 2.2.1 | Internet download |

Table 2. Package overview of Web 2.0 stack in SUSE Linux Enterprise Server 10 SP2 (continued)

| Component | Package Name | Version | Source |
|-----------|--------------|-------------------|-------------------|
| | | | |
| Squid | squid | 2.5STABLE12-18.13 | SLES10 SP2 |
| | | | |
| memcached | memcached | 1.2.6 | Internet download |
| | | | |
| Dojo | dojo-release | 1.2.3 | Internet download |

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information about the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY
10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

DB2, HiperSockets, IBM, OS/2®, System z, developerWorks, z/VM, z9, zSeries

The following terms are trademarks of other companies:

Java, JavaScript, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

