

<div style="text-align: center;"> <h1>An Overview of Config4★</h1> </div> <p>www.config4star.org</p>	<div style="text-align: center;"> <h2>Copyright</h2> </div> <p>Copyright 2010 Ciaran McHale (www.CiaranMcHale.com).</p> <p>Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:</p> <p>The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.</p> <p>THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.</p> <hr/> <p>An Overview of Config4* 2</p>
<div style="text-align: center;"> <h2>What is Config4*?</h2> </div> <ul style="list-style-type: none"> ■ Config4* is pronounced "config for star": <ul style="list-style-type: none"> - The pedantically correct "config for asterisk" does not sound as good - The "*" is a placeholder for the name of a programming language ■ Config4* is a configuration-file parser for several languages: <ul style="list-style-type: none"> - Config4Cpp (C++ version) - Config4J (Java version) - Other programming languages may be added in the future ■ Config4* is open-source and uses the MIT license: <ul style="list-style-type: none"> - Compatible with most/all other open-source and proprietary licenses ■ The project website is www.config4star.org <ul style="list-style-type: none"> - Provides source code and comprehensive documentation <hr/> <p>An Overview of Config4* 3</p>	<div style="text-align: center;"> <h2>Why is Config4* so good?</h2> </div> <ul style="list-style-type: none"> ■ Config4* has useful features rarely found in other configuration technologies: <ul style="list-style-type: none"> - Each feature is useful in its own right - In addition, there is synergy in the interaction of the features ■ Structure of this presentation: <ul style="list-style-type: none"> - First, explain the basic features of Config4* (similar to features in competing technologies) - Then, discuss the useful features, and the synergies between them <hr/> <p>An Overview of Config4* 4</p>
<div style="text-align: center;"> <h2>1. Basic features</h2> </div>	<div style="text-align: center;"> <h2>Basic features of syntax</h2> </div> <pre># this is a comment @include "another_file.cfg"; string_variable = "value"; list_variable = ["a", "list of", "values"]; string_concatenation = string_variable + "..."; list_concatenation = list_variable + ["another", "list"]; scope_name { variable_inside_scope = "value"; nested_scope { yet_another_variable = "..."; } } # scoping operator is "." # Example: "scope_name.variable_inside_scope"</pre> <div style="border: 1px solid black; padding: 5px; width: fit-content; margin-left: auto; margin-right: auto;"> Keywords are prefixed with "@" to prevent clashes with names of variables or scopes </div> <hr/> <p>An Overview of Config4* 6</p>
<div style="text-align: center;"> <h2>Basic API (Java syntax)</h2> </div> <pre>import org.config4j.*; ... Configuration cfg = Configuration.create(); String configFile = "file.cfg"; String scope = "foo"; String[] fontList; String logFile; int logLevel; try { cfg.parse(configFile); logFile = cfg.lookupString(scope, "log_file"); fontList = cfg.lookupList(scope, "font_list"); logLevel = cfg.lookupInt(scope, "log_level"); } catch(ConfigurationException ex) { System.err.println(ex.getMessage()); }</pre> <hr/> <p>An Overview of Config4* 7</p>	<div style="text-align: center;"> <h2>Notes</h2> </div> <ul style="list-style-type: none"> ■ A lookup operation merges its scope and localName parameters to form a fully-scoped name: <ul style="list-style-type: none"> - Example, lookupString("foo", "bar") → "foo.bar" - The scope parameter is usually set from a command-line argument - One configuration file can contain scopes for many applications ■ Data-type conversion: <ul style="list-style-type: none"> - Some lookup operations call lookupString() and then convert to the desired type - Examples: lookupInt(), lookupFloat(), lookupBoolean(), ... ■ An existing configuration-file parser: <ul style="list-style-type: none"> - Might not have all the features shown on the previous two slides - But such features are sort-of common ■ Now let's look at useful Config4* features rarely found elsewhere... <hr/> <p>An Overview of Config4* 8</p>

<p style="text-align: center;">2. Centralised configuration</p>	<p style="text-align: center;">Description of a common problem</p> <ul style="list-style-type: none"> ■ The Acme company makes and sells software: <ul style="list-style-type: none"> - Small customers will run the software on one computer - These customers want the convenience of a configuration file - Large customers will run the software on hundreds of computers - They do not want to maintain hundreds of copies of a configuration file - They insist on having a centralised configuration repository <ul style="list-style-type: none"> ■ It might cost Acme a lot of time and money to implement a centralised configuration mechanism: <ul style="list-style-type: none"> - Complexity of a client-server architecture? - Use a database? Administration skills required. License costs? - Extra complexity and expense if fault tolerance is required <ul style="list-style-type: none"> ■ Config4*, with help from a utility called Curl, provides a zero-cost solution <hr/> <p>An Overview of Config4* 10</p>
<p style="text-align: center;">The synergy of Config4* and Curl</p> <ul style="list-style-type: none"> ■ Curl (an abbreviation of “Client for URL”): <ul style="list-style-type: none"> - Is an open-source utility, available for most operating systems - Retrieves a file from a specified URL and echoes it to the console - Supports many protocols: HTTP, FTP, LDAP, ... ■ Example: <code>curl -sS http://www.example.com/file.cfg</code> <ul style="list-style-type: none"> - The “-sS” option instructs curl to not print any diagnostics ■ Config4* can parse: <ul style="list-style-type: none"> - A file: <code>cfg.parse("file.cfg")</code> - The output of executing a command: <pre>cfg.parse("exec#curl -sS http://...")</pre> <div style="margin-left: 40px;"> </div> <hr/> <p>An Overview of Config4* 11</p>	<p style="text-align: center;">The synergy of Config4* and curl (cont')</p> <ul style="list-style-type: none"> ■ Benefit for Acme. It satisfies small and large customers: <ul style="list-style-type: none"> - Small customers use a “file.cfg” command-line argument - A large customer runs software on many computers, and specifies “exec#curl -sS http://centralisedHost/file.cfg” as a command-line argument ■ Benefits for large customers: <ul style="list-style-type: none"> - They can use any protocol supported by curl - They are not restricted to using only curl (they can use a utility that retrieves configuration from, say, a database) - They can use a fault-tolerant database or web server, if needed - As more and more applications use Config4*, the cost of maintaining a centralised database or web server is amortised <hr/> <p>An Overview of Config4* 12</p>
<p style="text-align: center;">3. Fallback configuration</p>	<p style="text-align: center;">The goal of “install and use” for applications</p> <ul style="list-style-type: none"> ■ “Install and use” applications are convenient to use: <ul style="list-style-type: none"> - Similar to “plug and play” hardware - Unfortunately, many applications requires a configuration step before use - The need to “configure before use” can irritate users ■ Ideally, an application will have an <i>optional</i> configuration file: <ul style="list-style-type: none"> - Embedded configuration means it can run <i>without</i> a configuration file (thus bypassing the need to “configure before use”) - The embedded configuration can be overridden with an external configuration file ■ Config4* enables developers to achieve this goal in two steps: <ul style="list-style-type: none"> - Run the <code>config2cpp</code> or <code>config2j</code> utility to create embedded configuration data - Call the <code>setFallbackConfiguration()</code> operation <hr/> <p>An Overview of Config4* 14</p>
<p style="text-align: center;">The config2cpp and config2j utilities</p> <ul style="list-style-type: none"> ■ The <code>config2cpp</code> and <code>config2j</code> utilities: <ul style="list-style-type: none"> - Read a configuration file, and - Generate a C++ or Java class that provides a snapshot of the file's contents ■ This provides “fallback” configuration data that can be embedded in an application ■ Examples of use: <pre>config2cpp -cfg Fallback.cfg -class FallbackConfig -singleton config2j -cfg Fallback.cfg -class FallbackConfig -singleton</pre> ■ The generated class provides a <code>getString()</code> operation that returns the “fallback” configuration data <hr/> <p>An Overview of Config4* 15</p>	<p style="text-align: center;">The setFallbackConfiguration() operation</p> <ul style="list-style-type: none"> ■ An application sets fallback configuration as follows: <pre>Configuration cfg = Configuration.create(); String cfgFile = ...; try { if (cfgFile != null) { cfg.parse(cfgFile); } cfg.setFallbackConfiguration(Configuration.INPUT_STRING, FallbackConfig.getString()); logFile = cfg.lookupString(scope, "log_file"); } catch(ConfigurationException ex) { System.err.println(ex.getMessage()); }</pre> ■ The Config4* lookup operations work as follows: <ul style="list-style-type: none"> - Search for the specified variable in the configuration object - If found, then return its value - Otherwise, search for the specified variable in the fallback configuration <hr/> <p>An Overview of Config4* 16</p>

Synergy of fallback and centralised configuration

- Fallback and centralised configuration are independent features
 - But they interact to provide synergy
- A Config4*-based application:
 - Can use fallback configuration to provide “install and use” convenience for new users
 - Can use an external configuration file to override fallback configuration
 - Can use “exec#curl ...” if the user deploys the application on many computers
- Thus:
 - Config4* scales from single-user to enterprise deployments (you could even use fallback configuration in an embedded system)
 - It is difficult to think of another configuration technology that provides this level of flexibility

An Overview of Config4*

17

4. Adaptable configuration

Description of a common problem

- Often, the contents of a configuration file change when:
 - Moving the application from one computer to another
 - Running the application under another user name
- Editing a configuration file to make such changes is tedious
- It would be better if a configuration file could automatically adapt to its runtime environment:
 - Then, the *same* configuration file could be used on multiple computers and by multiple users
 - Config4* provides excellent support for this

An Overview of Config4*

19

The `getenv()` and `exec()` functions

- The `getenv()` function:
 - Returns the value of an environment variable
 - Is typically used to access the name of the user or the installation directory for software: `getenv("USERNAME"), getenv("FOO_HOME")`
- The `exec()` function:
 - Executes a command and returns its standard output
 - Is typically used to determine the host name: `exec("hostname")`
 - A security mechanism prevents execution of malicious commands
- These operations, combined with the “+” operator, enable a configuration file to adapt to its runtime environment. Example:

```
log_dir = getenv("FOO_HOME") + "/logs/"
          + exec("hostname");
```

An Overview of Config4*

20

If-then-else statements and `osType()`

```
production_hosts = ["pizza", "pasta", "cheese"];
test_hosts = ["foo", "bar", "widget", "acme"];

@if (exec("hostname") @in production_hosts) {
  server_x.port = "5001";
  server_y.port = "5002";
  server_z.port = "5003";
} @elseif (exec("hostname") @in test_hosts) {
  server_x.port = "6001";
  server_y.port = "6002";
  server_z.port = "6003";
} @else {
  @error "This is not a production or test machine!";
}

@if (osType()) == "windows") {
  ...
} @else { # UNIX
  ...
}
```

An Overview of Config4*

21

Adapting to command-line options

- Users may want to use command-line options to override variables in a configuration file
- Config4* supports a two-step approach for doing this:
 - Before parsing a configuration file, the application calls `insertString()` to insert *name-value* pairs obtained from the command line
 - The configuration file uses the conditional assignment (“?”=) operator to provide default values for variables
- The following slides illustrate these steps

An Overview of Config4*

22

Using the `insertString()` operation

```
import org.config4j.*;
...
Configuration cfg = Configuration.create();
for (int i = 0; i < args.length; i++) {
  if (args[i].equals("-set")) {
    cfg.insertString(scope, args[i+1], args[i+2]);
    i = i + 2;
  }
}
try {
  cfg.parse(configFile);
  ... // calls to cfg.lookup<Type>() operations
} catch (ConfigurationException ex) {
  System.err.println(ex.getMessage());
}
```

An Overview of Config4*

23

Using the conditional assignment (“?”=) operator

- The “?”= operator assigns a value to a variable *only if* the variable does not already exist
- Example of syntax:


```
log_level ?= "0";
username ?= getenv("USERNAME");
password ?= "";
```
- Typically, such variables can be pre-set by command-line options that are processed as shown on the previous slide
- In this way, a configuration file can adapt to (be overridden by) command-line options

An Overview of Config4*

24

Synergy

- Adaptable configuration is independent of centralised configuration
- However, those features can interact to provide synergy.
- Example:
 - A large company deploys an application on 500 computers
 - A single configuration file is stored on a web server and accessed via `"exec#curl ..."`
 - That centralised configuration file can use `@if-then-@else, getenv()` and `exec()` to adapt its contents for each computer

5. Useful data types

Durations

- Some configuration files need to specify durations:
 - For example, connection timeout, idle timeout, transaction timeout...
 - In most configuration files, these are expressed as integer values (denoting seconds or milliseconds)
- Example from a product that does *not* use Config4*:


```
refresh: 28800
retry: 7200
expire: 1209600
minimum: 86400
```

Expressed in seconds
- Equivalent in Config4* syntax:


```
refresh = "8 hours";
retry = "2 hours";
expire = "2 weeks";
minimum = "1 day";
```

Units can also be minutes, seconds or milliseconds

Durations (cont')

- Config4* can convert durations into seconds or milliseconds:


```
x = cfg.lookupDurationSeconds(scope, "refresh");
y = cfg.lookupDurationMilliseconds(scope, "retry");
```
- The value "infinite" is converted into the value -1

Memory sizes

- Config4* also supports memory sizes
- Examples of syntax:


```
buffer_size = "512 bytes";
cache_size = "32 KB";
max_log_size = "1.5 GB";
```
- Config4* can convert memory sizes into bytes, KB or MB:


```
lookupMemorySizeBytes(scope, localName)
lookupMemorySizeKB(scope, localName)
lookupMemorySizeMB(scope, localName)
```

Other data-type conversions

- Other operations enable you to quickly write code to:
 - Convert string values to integer constants (like `enum` in C/C++)
 - Example: "red" → 0, "green" → 1, "blue" → 2
 - Parse values "`<units> <number>`" or "`<number> <units>`"
 - Examples:
 - "£19.99" → ("£", 19.99)
 - "42 cm" → ("cm", 42)
- You can also process a list as if it were a table. Example:


```
price_list = [
# description      unit price
#-----
"apple",           "£0.49",
"orange",          "£2.99",
];
```

6. Schema validation

Benefits of schema validation

- A *schema* is a blueprint or definition of a system.
- Examples:
 - A database schema defines the layout of a database
 - DTD, XML Schema and RELAX NG are competing schema languages for defining the permitted contents of an XML file
- Config4* has a schema language, which provides:
 - An intuitive, easy-to-use syntax
 - An easy-to-use API
 - The ability for developers to define new schema data-types.

<p style="text-align: center;">Example configuration scope</p> <ul style="list-style-type: none"> An application uses a configuration scope like that shown below: <pre>foo_server { idle_timeout = "2 minutes"; log_level = "3"; log_file = "/tmp/foo.log"; price_list = [# time colour price #----- "shirt", "green", "EUR 19.99", "jeans", "blue", "USD 39.99",]; }</pre> <ul style="list-style-type: none"> The next slide shows how to perform schema validation for such a scope 	<p style="text-align: center;">Example of schema validation (Java syntax)</p> <pre>Configuration cfg = Configuration.create(); SchemaValidator sv = new SchemaValidator(); String schema = new String[] { "@typedef colour = enum[red, green, blue]", "@typedef money = units_with_float[EUR, GBP, USD]", "idle_timeout = durationMilliseconds", "log_level = int[0,5]", "log_file = string", "price_list = table[string,item, colour,colour, money,price]" }; try { cfg.parse(configFile); sv.parseSchema(schema); sv.validate(cfg, "foo_server", ""); } catch(ConfigurationException ex) { System.err.println(ex.getMessage()); }</pre> <ul style="list-style-type: none"> A descriptive exception is thrown if schema validation fails
An Overview of Config4* 33	An Overview of Config4* 34

<p style="text-align: center;">Comparison with XML Schema</p> <ul style="list-style-type: none"> XML Schema: <ul style="list-style-type: none"> Is very verbose Has a steep learning curve: <ul style="list-style-type: none"> Syntax specification is written in impenetrable legalese (about 380 pages long if you convert it from HTML into PDF format) Good books on XML Schema are 400–500 pages long Provides difficult-to-understand error messages In contrast, the Config4* schema language: <ul style="list-style-type: none"> Is very concise Is intuitive and easy to learn: <ul style="list-style-type: none"> Syntax specification, with examples, is defined in 13 pages Provides easy-to-understand error messages 	<p style="text-align: center;">7. Reuse with the @copyFrom statement</p>
An Overview of Config4* 35	

<p style="text-align: center;">Description of problem</p> <ul style="list-style-type: none"> Some applications are related to other applications. Examples: <ul style="list-style-type: none"> Applications that are developed as part of the same project Replicas for a server application Such applications may have similar configuration settings: <ul style="list-style-type: none"> Most variables have identical values A few variables have different values Can such applications reuse the variables with identical values? <ul style="list-style-type: none"> Doing this can significantly reduce the size of configuration files The @copyFrom statement facilitates this 	<p style="text-align: center;">Example of the @copyFrom statement</p> <pre>server.defaults { timeout = "2 minutes"; log { dir = getenv("FOO_HOME") + "/logs"; level = "2"; } } foo_server { @copyFrom "server.defaults"; log.level = "1"; # override copied value } bar_server { @copyFrom "server.defaults"; timeout = "30 seconds"; # override copied value }</pre>
An Overview of Config4* 37	An Overview of Config4* 38

<p style="text-align: center;">Conditional @include and @copyFrom</p> <ul style="list-style-type: none"> Config4* provides conditional variations of the @include and @copyFrom statements These help a configuration file adapt to its environment. Examples: <pre>@include getenv("HOME") + "/.foo.cfg" @ifExists; override.pizza { ... } override.pasta { ... } foo_server { ... # set default values @copyFrom "override." + exec("hostname") @ifExists; }</pre>	<p style="text-align: center;">8. The "uid-" prefix</p>
An Overview of Config4* 39	

<p style="text-align: center;">The "uid-" prefix</p> <ul style="list-style-type: none"> ■ Let's assume you want to store details about employees ■ You might try the following: <pre>employee { name = "John Smith"; ... } employee { name = "Jane Doe"; ... }</pre> ■ That will not work: <ul style="list-style-type: none"> - Because the second occurrence of <code>employee</code> re-opens the <i>existing</i> scope, so the details of Jane Doe override those of John Smith ■ Config4* provides a "uid-" prefix for such situations: <ul style="list-style-type: none"> - "uid" is an abbreviation for "unique identifier". - Config4* expands each name of the form form <code>uid-employee</code> into <code>uid-<unique-number>-employee</code> <hr/> <p>An Overview of Config4* 41</p>	<p style="text-align: center;">The "uid-" prefix (cont')</p> <ul style="list-style-type: none"> ■ Fixed example: <pre>uid-employee { name = "John Smith; ... } uid-employee { name = "Jane Doe"; ... }</pre> ■ The Config4* API provides operations for processing uid entries ■ The "uid-" prefix makes Config4* suitable for use as a data-file format <hr/> <p>An Overview of Config4* 42</p>
<p style="text-align: center;">9. Comprehensive documentation</p>	<p style="text-align: center;">Comprehensive documentation</p> <ul style="list-style-type: none"> ■ Many open-source software projects provide minimal or no documentation: <ul style="list-style-type: none"> - Lack of documentation creates a frustrating learning curve for users - You might have to spend days or weeks "playing with" a project to learn if it suits your needs ■ In contrast, Config4* provides comprehensive documentation: <ul style="list-style-type: none"> - Getting Started guide - C++ and Java API guide - Practical Usage guide (this provides expert advice) - Maintenance guide (for people interesting in extending/porting Config4*) ■ The high-quality documentation (400 pages in total): <ul style="list-style-type: none"> - Significantly reduces the learning curve - Enables you to quickly decide if Config4* suits your needs <hr/> <p>An Overview of Config4* 44</p>
<p style="text-align: center;">10. Summary</p>	<p style="text-align: center;">Summary</p> <ul style="list-style-type: none"> ■ Config4* provides features common to many other configuration technologies ■ Config4* <i>also</i> provides: <ul style="list-style-type: none"> - Fallback (embedded) configuration for "install and run" applications - Centralised configuration (vital for large deployments) - Adaptable configuration: <ul style="list-style-type: none"> - <code>@if-then-@else</code>, <code>getenv()</code>, <code>exec()</code>, <code>osType()</code> - The <code>?=</code> operator enables integration with command-line options - Reusable configuration (the <code>@copyFrom</code> statement) - Scopes enable one file to store configuration for multiple applications - Useful data-types: duration, memory sizes, tables, ... - Schema validation - The "uid-" prefix (can use Config4* as a data-file format) - Comprehensive, high-quality documentation <hr/> <p>An Overview of Config4* 46</p>