

Clarion Magazine

Clarion News

- ◊ » [gNotes 3.1](#)
- ◊ » [DMC Viewer Licensing Changes](#)
- ◊ » [DMC Runtime Version](#)
- ◊ » [Build Automator 1.30.150](#)
- ◊ » [SetupBuilder 6.8 Build 2266 Developer Edition Hotfox](#)
- ◊ » [Aussie DevCon Pics](#)
- ◊ » [Aussie DevCon Comments](#)
- ◊ » [DMC 1.4.0.1](#)
- ◊ » [Function Points Strategy for Business Information System Estimating](#)

[More news]

- ◊ » [Clarion.NET FAQ](#)
- ◊ » [Clarion# Language Comparison](#)
- ◊ » [Creating A UserControl In Clarion#](#)
- ◊ » [Coloring Template Source In C7/Clarion#](#)

[More Clarion & .NET]

[More Clarion 101]

Latest Free Content

- ◊ » [OSCON: Orbiting Pushes Data To Web Pages](#)
- ◊ » [OSCON: Open-Sourcing Identity](#)
- ◊ » [OSCON: Open Source Gains Ground In The Enterprise](#)
- ◊ » [OSCON: Pro PostgreSQL](#)
- ◊ » [ClarionMag At OSCON 2008](#)
- ◊ » [Source Code Library 2008.06.30 Available](#)

[More free articles]

Clarion Sites

Clarion Blogs

Save up to **50% Off** ebooks.
Subscription has its rewards.



Latest Subscriber Content

OSCON: Even Faster Web Sites

Steve Souder is the former Chief Performance Yahoo! and is now doing similar performance-related work at Google. Souder demonstrated tools and techniques for speeding web page loading by working on client side issues.

Posted Thursday, July 31, 2008

OSCON: Using Cascading Style Sheets

Gavin Doughtie covered a number of cascading style sheet (CSS) features that can take the place of Javascript code.

Posted Thursday, July 31, 2008

OSCON: Orbiting Pushes Data To Web Pages

Orbiting is an implementation of Comet, and in programming terms Comet is a web application model where the client, via Javascript, keeps an open connection back to the server. This allows the server to push data back to the client without having to first wait for a request.

Posted Friday, July 25, 2008

OSCON: Secrets of Javascript Libraries

Resig's presentation covered a variety of Javascript programming techniques for cross-browser support but he also touched on some language fundamentals.

Posted Thursday, July 24, 2008

OSCON: A/B Testing

A/B testing is the process of creating two versions of a web page (or other code) and evaluating users responses in order to optimize the user's interaction. For many sites this translates as maximizing conversion of visitors to customers.

Posted Thursday, July 24, 2008

OSCON: Open-Sourcing Identity

Atlassian's Dave O'Flynn announced a new (and as yet unnamed) open source identify management project at OSCON. Atlassian's effort is intended to address the shortcomings of existing identification systems such as LDAP, SAML and OpenID.

Posted Wednesday, July 23, 2008

OSCON: Open Source Gains Ground In The Enterprise

An O'Reilly Radar report released today estimates that perhaps 10% of all new IT jobs involve at least some open source software. The report cites explosive growth in open source projects, draws on analysis of job posting data and considers a number of case studies.

Posted Wednesday, July 23, 2008

OSCON: Pro PostgreSQL

In his Pro PostgreSQL tutorial session Robert Treat covered a variety of topics including understanding the nature of the PostgreSQL community, performance optimizations, upgrade paths, availability and hardware.

Posted Tuesday, July 22, 2008

ClarionMag At OSCON 2008

This week (July 21-25) Clarion Magazine is at OSCON 2008, O'Reilly's open source conference held annually in Portland Oregon. Stay tuned for daily reports...

Posted Sunday, July 20, 2008

Creating A UserControl In Clarion#

UserControls in Clarion# are much like control templates; they are a way to bundle up multiple controls and related source code in a single package which can be easily reused. Randy Rogers shows how it's done.

Posted Saturday, July 19, 2008

Coloring Template Source In C7/Clarion#

The new IDE's editor is a vast improvement over the C6 source editor, making it an appealing choice for template editing even in advance of the AppGen. Dave Harms shows how easy it is to set up coloring rules for template source.

Posted Friday, July 18, 2008

Source Code Library 2008.06.30 Available

The Clarion Magazine Source Code Library has been updated to include the latest source. Source code subscribers can download the June 2008 update from the [My ClarionMag](#) page. If you're on Vista please run Lindersoft's Clarion detection patch first.

Posted Wednesday, July 02, 2008

[Last 10 articles] [Last 25 articles] [All content]

Source Code

The ClarionMag Source Code Library

Clarion Magazine is more than just a great place to learn about Clarion development techniques, it's also home to a massive collection of Clarion source code. Clarion subscribers already know this, but now we've made it easier for subscribers and non-subscribers alike to find the code they need.

The Clarion Magazine Source Library is a single point download of all article source code, complete with an article cross-reference.

[More info](#) • [Subscribe now](#)

Printed Books & E-Books

E-Books

E-books are another great way to get the information you want from Clarion Magazine. Your time is valuable; with our e-books, you spend less time hunting down the information you need. We're constantly collecting the best Clarion Magazine articles by top developers into themed PDFs, so you'll always have a ready reference for your favorite Clarion development topics.

Printed Books

As handy as the Clarion Magazine web site is, sometimes you just want to read articles in print. We've collected some of the best ClarionMag articles into the following print books:



- o » Clarion Tips & Techniques Volume 4 - ISBN 978-0-9784034-09
- o » Clarion Tips & Techniques Volume 3 - ISBN: 0-9689553-9-8
- o » Clarion 6 Tips & Techniques Volume 1 - ISBN: 0-9689553-8-X
- o » Clarion 5.x Tips and Techniques, Volume 1 - ISBN: 0-9689553-5-5
- o » Clarion 5.x Tips and Techniques, Volume 2 - ISBN: 0-9689553-6-3
- o » Clarion Databases & SQL - ISBN: 0-9689553-3-9

We also publish Russ Eggen's widely-acclaimed [Programming Objects in Clarion](#), an introduction to OOP and ABC.

From The Publisher

About Clarion Magazine

Clarion Magazine is your premier source for news about, and in-depth articles on Clarion software development. We publish articles by many of the leading developers in the Clarion community, covering subjects from everyday programming tasks to specialized techniques you won't learn anywhere else. Whether you're just getting started with Clarion, or are a seasoned veteran, Clarion Magazine has the information *you* need.

Subscriptions

While we do publish some free content, most Clarion Magazine articles are for subscribers only. Your [subscription](#) not only gets you premium content in the form of new articles, it also includes all the back issues. Our [search engine](#) lets you do simple or complex searches on both articles and news items. Subscribers can also post questions and comments directly to articles.

Satisfaction Guaranteed

For just pennies per day you can have this wealth of Clarion development information at your fingertips. Your Clarion magazine subscription will more than [pay for itself](#) - you have my personal guarantee.

Dave Harms

ISSN

Clarion Magazine's ISSN

Clarion Magazine's [International Standard Serial Number \(ISSN\)](#) is 1718-9942.

About ISSN

The ISSN is the standardized international code which allows the identification of any serial publication, including electronic serials, independently of its country of publication, of its language or alphabet, of its frequency, medium, etc.

Copyright © 1999-2008 by CoveComm Inc. All Rights Reserved. Reproduction in any form without the express written consent of CoveComm Inc., except as described in the [subscription agreement](#), is prohibited.

Clarion Magazine

Clarion News

[Search the news archive](#)

RPMxt Update and Drill Down Support

A free update for current RPMxt users will be released shortly. Shortly thereafter the initial release of the drill down, clickable link, add-on will be released. This new release is expected by mid-September.

Posted Thursday, August 07, 2008

Lodestar 30% Off, Birthday Celebration

For the remainder of August there is a 30% discount on all RPM and AFE products including upgrades. Help Lee celebrate his birthday (some time this month) and save a few bucks in the process.

Posted Thursday, August 07, 2008

gNotes 3.1

A re-branded version of gNotes 3.1 has been released. This version comes with precompiled DLL / LIB for C63 9054-9056, partially re-written template code, new installer and more. No black boxes - source code is included. Price for gNotes 3.1 Source Code is USD 79.00. Until July 15 2008 an introductory price of USD 39.00 is available to users of any of the following: gCal, gNotes, gReg, gCalc, gFileFind, gQ, gSec, and LGP. Until July 15th, owners of any above mentioned utilites can purchase gCal 3.11 Source Code at USD 49.00. New customers may purchase one product at normal price and then order next one at special price.

Posted Wednesday, July 02, 2008

DMC Viewer Licensing Changes

DMC Viewer licensing has changed from one copy to unlimited copies (all registered users will receive an updated activation license). Some prices have also been updated to reflect thischange.

Posted Wednesday, July 02, 2008

DMC Runtime Version

DMC can now be called with command line arguments. DMC closes automatically at the end of processing. All tasks are logged along with errorcodes. This feature will be included in DMC from the next version on.

Posted Wednesday, July 02, 2008

Build Automator 1.30.150

Icetips Creative, Inc. has released Build Automator version 1.30.150. If you are a licensed user, you can use the "Help | Check for updates" from the main menu in the Build Automator to download and update the software automatically. You can purchase the Build Automator with a 60 day Maintenance Plan for US\$ 99.00 or you can purchase it with a 1 year Maintenance Plan for US\$ 149.00. Icetips Creative, Inc. has recently set up a partnership with Indigo Rose, makers of Setup Factory and MSI Factory and will start developing actions for those products in July. Icetips has also entered a technological partnership with CodeGear,

makers of Delphi, RAD Studio and C++ Builder, now being acquired from Borland by Embarcadero Technologies.

Posted Wednesday, July 02, 2008

SetupBuilder 6.8 Build 2266 Developer Edition Hotfix

A downloadable hotfix for SetupBuilder 6.8 Build 2266 (Developer Edition) is available.

Posted Wednesday, July 02, 2008

Aussie DevCon Pics

Stu Andrews has posted some pics from the 2008 Aussie Devcon.

Posted Wednesday, July 02, 2008

Aussie DevCon Comments

Read comments from DevCon participants at Geoff Spillane's message board.

Posted Wednesday, July 02, 2008

DMC 1.4.0.1

DMC 1.4.0.1 is now available. Changes include: SQL query environment; Clone to DAT/DBF; User option for a Sum Up display screen of the Main Wizard (with canceling possibility); Direct access in Viewer to the last 25 most recently used tables (all formats); All Data Management screens redesigned to accommodate more space (better visibility); Data Management Task wizard - define step by step your task (change source and / or destination settings without losing pace); Create ODBC wizard; User option to read only first 50 records before reaching the Mapping tab; Define the Color of all List Box records; List of all SQL Key's in Structure Tab - Structures stored in clipboard; List by ... processes changed to working on queues (speed); Processing of Columns to define if DATETIMES are present also changed to working on a queue (speed); Bug fixes.

Posted Wednesday, July 02, 2008

Function Points Strategy for Business Information System Estimating

A new Whitemarsh Information Systems short paper on function point analysis is now available for download. See item #13.

Posted Wednesday, July 02, 2008

New DMC Video

A preview video of the new DMC features is now available. This is a 15MB download.

Posted Friday, June 20, 2008

Lindersoft Vacation Schedule

The Lindersoft office will be closed for vacation from Saturday, June 28 to Sunday, July 6. During this time the online store will remain open. Orders for new licenses will be processed immediately. Upgrade orders and subscription extension orders will be processed when the office reopens. Emergency support will be provided on a limited basis.

Posted Friday, June 20, 2008

SetupBuilder 6.8 June 2008 Update

The SetupBuilder 6.8 June 2008 Update is now available. This release is free of charge to all SetupBuilder customers who have an active SetupBuilder maintenance subscription plan. If you do not have an active subscription plan, please contact your account manager at sales at lindersoft.com.

Posted Friday, June 20, 2008

iQ-Notes 3.64

New in iQ-Notes 3.64: Play a generic Open/Close/Delete/Alarm sounds when opening, closing and deleting notes; Set non-active notes to a transparent background; Switch to confirm when deleting notes; More flexibility and options for Find features including drop-downs of history; Title can automatically be created when you start typing the details of a note; Integrated HTML help and better layout of User Options tab; Preview contents of notes while in Manage Notes detail screen; New FTP socket tool with better synchronization with Vista.

Posted Friday, June 20, 2008

SimGlobalButtons Template Enhanced

The SimGlobalButtons Template now let you exclude procedures from the automatically setting of button icons and options throughout an application. The upgrade to version 1.01 is free to registered users. The template can be purchased for \$29.50 US from ClarionShop.

Posted Friday, June 20, 2008

Clarion Developer Hosting Sale

Now through the end of June Oak Park is running a sale on its shared hosting plans, starting as low as \$3.99 per month for Clarion Developers and their clients.

Posted Friday, June 20, 2008

Ready-to-go Logos \$49

1st Logo Design offers a selection of ready-to-go logos at \$49. These logos are only sold once to retain exclusivity, once sold they are removed from the available list and never sold again. Logos delivered in 300 dpi B&W AI and CDR (vector logos only), PSD, TIFF, PDF, PNG, 72dpi PNG, BMP, JPG, GIF in 800px, 700px, 600px, 500px, 400px, 300px, 200px, 150px and 100px widths.

Posted Friday, June 20, 2008

BuildAutomator Goes Gold

Icetips Creative has released the gold version (1.30.100) of Build Automator. For the final release the Standard Edition has been dropped and the price of the Developer Edition has been lowered to US\$149 with a one year Maintenance Plan or US \$99 including a 60 day Maintenance Plan. The Maintenance plan is required in order to install updates to the software. Additional features will be available as plugins.

Posted Friday, June 20, 2008

Clarion Version Switcher, Ver 2

Dirt simple application to manage multiple Clarion versions on the same machine. Works with any version of Clarion including multiple versions of Clarion 7 and Clarion.NET. Full source provided, APP file created with C6.1.9033. v2 adds comment, tool tips and rapid selection across multiple versions.

Posted Monday, June 09, 2008

CapeSoft Price Increases June 17th

A number of CapeSoft products will go up in price on June 17th, including: AnyFont (from \$49 to \$69); MessageBox (from \$59 to \$69); NetTalk 4 Upgrade (from \$199 to \$249); NetTalk 4 (from \$399 to \$499); RightReports (from \$249 to

\$349) (still in beta); Tearoff (from \$49 to \$69).

Posted Thursday, June 05, 2008

CapeSoft Releases TabTree

TabTree allows you to add a Sheet-Management control to a window. It creates a list of all the tabs in a list box. You can then navigate between tabs using either the keyboard or mouse or the list itself. TabTree is available for an introductory special of \$69 until June 17, 2008, when the price goes up to \$99. Users upgrading from SimTabTree can purchase an upgrade for \$49 until June 17, 2008, when the price will go up to \$69.

Posted Thursday, June 05, 2008

NeatMessage 2.15

NeatMessage 2.15 includes an urgent fix to a problem with the Enter key not always firing the selected button. NeatMessage is a Message Box replacement from Huenuleufu Development.

Posted Thursday, June 05, 2008

J-Zip 1.01

J-Zip 1.01 adds an optional progress window.

Posted Thursday, June 05, 2008

J-Skype Public Chat

The new public chat room feature in J-Skype 2.5 lets you participate in Skype chats from within a Clarion application.

Posted Thursday, June 05, 2008

J-Zip Released

Strategy Online has released J-Zip for Clarion. This is a template wrapper around the 7Zip library. For now there is a single code template to add files to a ZIP archive, and another code template to extract files. Demo available; price is \$15.

Posted Thursday, June 05, 2008

SealSoft Discount Ends June 10

Anyone purchasing a SealSoft product at ClarionShop by June 10 can receive a 40% discount. Include 'SealSoft DevCon08 Discount' in the comment field. Buyers at Motleysoft.com will also receive the discount.

Posted Thursday, June 05, 2008

Clarion Desktop 4.11

Clarion Desktop 4.11 is now available.

Posted Thursday, June 05, 2008

J-Skype Week

The gold release of J-Skype 2.50 is due out this Friday.

Posted Thursday, June 05, 2008

SetupBuilder 6.8 May 2008 Hotfix

Lindersoft has released a SetupBuilder 6.8 hotfix which addresses a command line compiler issue. Web Update can be accessed via "Help | Check for Updates..." in your SetupBuilder IDE.

Posted Thursday, June 05, 2008

Clarion Magazine

Creating A UserControl In Clarion#

by Randy Rogers

Published 2008-07-19

This is the second in a series of articles in which I will share much of what I discovered along the way to creating a 'standard' Clarion menu, browse, form, report type of application in Clarion#. In the [first article](#) I created a Keystone.UIControls assembly to contain my custom controls and explained how to create a RoundPanel control by deriving from the System.Windows.Forms.Panel class. In this article I will demonstrate how to create a custom datatype class and use that class to create a CountryPicker control derived from the System.Windows.Forms.UserControl class.

User controls

The Microsoft documentation for the UserControl Class states:

The UserControl gives you the ability to create controls that can be used in multiple places within an application or organization. You can include all the code needed for validation of common data you ask the user to input; some examples of this are e-mail addresses, telephone numbers, and postal codes. Another efficient use of the user control is to simply preload a ComboBox or ListBox with static items you commonly use in almost every application; some examples of this are countries/regions, cities, states, and office locations.

My design for the CountryPicker control is simple. The International Standards Organization defines [two character](#) and [three character](#) country codes for each country name. I will use the two character code, but my control will only display the country name to the user. I will use properties to allow the user to get or set the CountryName; the CountryCode will be read-only. I will also use a virtual method to initialize the array with country objects; this will allow other developers to derive the class and populate their own list of countries.

The control will trigger an event when the user selects a new country so I have a way to execute code that assigns the selected country code to a variable.

I plan to initialize a ComboBox with an ArrayList of Country objects to create my CountryPicker control. To represent the Country objects I first need to create a class with CountryCode and CountryName properties.

This is a continuation of the previous article, so I start by opening the Keystone.UIControls project.

Creating a Custom Datatype Class

In the Projects pad I right click and select Add | New Item& from the context menu. In the New File dialogue window I select the Clarion.Net category and the Member File quick start template. I enter country.cln as the filename and press the Create button. My starting point looks like this:

```
MEMBER("")

NAMESPACE(Keystone.UIControls)
MAP
END
```


Next I declare the Country class and CountryCode and CountryName properties. As a matter of style I use PRIVATE member variables prefixed with m_ to store the values, and I provide GETONLY property accessors for those variables. I use the Construct method to initialize the properties and I add a ToString method to return the CountryName property.

```
Country    CLASS,TYPE,PUBLIC
m_ShortName    String,PRIVATE
CountryCode    PROPERTY,String,PUBLIC,GETONLY
m_LongName    String,PRIVATE
CountryName    PROPERTY,String,PUBLIC,GETONLY
Construct    PROCEDURE(String strLongName, String strShortName),PUBLIC
ToString    PROCEDURE(),String,DERIVED,PUBLIC
            END

Country.Construct PROCEDURE(String strLongName, String strShortName)
CODE
    SELF.m_ShortName = strShortName
    SELF.m_LongName = strLongName

Country.Get_CountryCode PROCEDURE()
CODE
    RETURN SELF.m_ShortName

Country.Get_CountryName PROCEDURE()
CODE
    RETURN SELF.m_LongName
Country.ToString PROCEDURE()
CODE
    RETURN SELF.m_LongName
```

Creating the CountryPicker control

Again, in the Projects pad I right click and select Add | New Item& from the context menu. In the New File dialogue window I select the Clarion.Net | Windows Applications category and the User Control quick start template. I enter CountryPicker.cln as the filename and press the Create button. My starting point looks like this:

```
MEMBER("")

NAMESPACE(Keystone.UIControls)
USING(System)
USING(System.Drawing)
USING(System.Windows.Forms)
!!! <summary>
!!! Description of CountryPicker.
!!! &lt;/summary>

CountryPicker CLASS(System.Windows.Forms.UserControl),TYPE,NETCLASS,PARTIAL
CONSTRUCT    PROCEDURE(),PUBLIC
            END
```

```

CountryPicker.CONSTRUCT PROCEDURE()
    CODE
    !
    ! The InitializeComponent() call is required for Windows Forms designer support.
    !
    SELF.InitializeComponent()

```

In design mode the user control looks like a resizable window design panel with grid marks (see Figure 1).

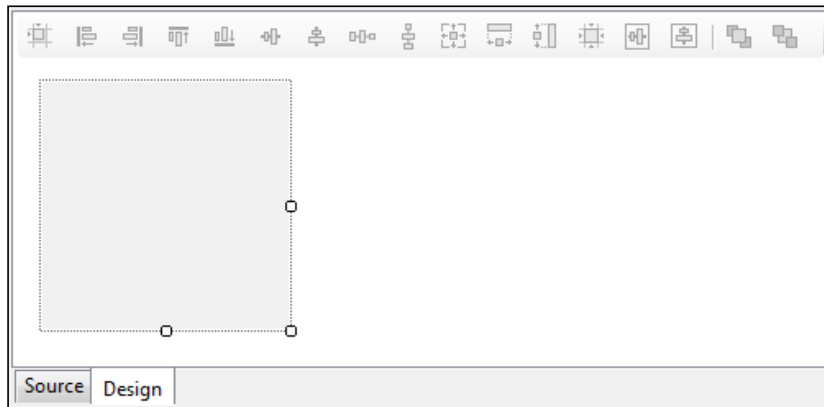


Figure 1. The empty user control in design mode

I change the default UserControl1 name to CountryPicker and populate a Label control named CountryLabel and a ComboBox control named CountryComboBox, setting the DropDownStyle property to DropDownList. Finally I resize the CountryPicker user control so that it is the same height as the ComboBox and just wide enough to contain the two controls.

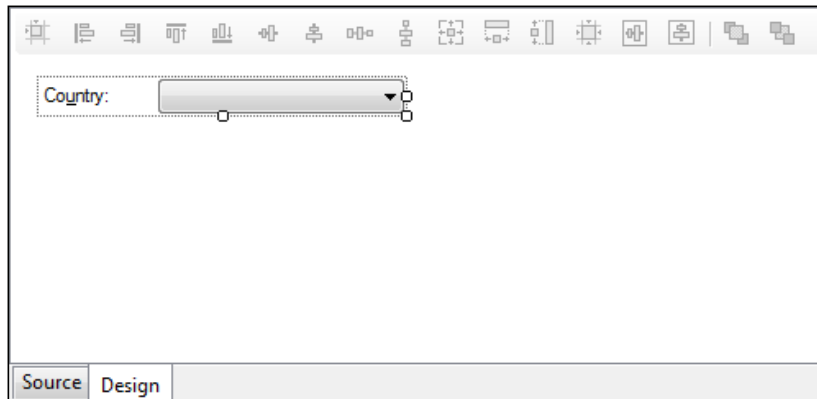


Figure 2. The user control with controls populated

To handle resizing of the user control when the user resizes the form, I switch to the Events view on the properties pad and double click on the SizeChanged event to create the CountryPicker_SizeChanged event handler.

```

CountryPicker.CountryPicker_SizeChanged PROCEDURE(System.Object sender, System.EventArgs e)
    CODE
    SELF.CountryComboBox.Width = SELF.Width - 72
    RETURN

```

Adding the SelectedIndexChanged event

The CountryComboBox has a SelectedIndexChanged event that I will use to trigger a similar event for the CountryPicker control. To do this I first declare a DELEGATE within the module MAP structure. A DELEGATE acts like a Clarion 6 procedure prototype.

```
MAP
    SelectedIndexChangedEventHandler(System.Object sender, |
        System.EventArgs e) ,DELEGATE,PUBLIC
END
```

Then I declare the EVENT inside the CountryPicker class.

```
SelectedIndexChanged EVENT,SelectedIndexChangedEventHandler,PUBLIC
```

Finally in design view I select the CountryComboBox, switch the Properties pad to the Events view and double click on the SelectedIndexChanged event to create the event stub:

```
CountryPicker.CountryComboBox_SelectedIndexChanged PROCEDURE(|
    System.Object sender, System.EventArgs e)
CODE
IF NOT SELF.SelectedIndexChanged = NULL
    SELF.SelectedIndexChanged(sender, e)
END
```

Because I don't have a default handler for the event I first check for an event handler and only call it if it exists.

Since I want to use an ArrayList to initialize the ComboBox, I need to add another USING statement at the start of the module.

```
USING(System.Collections)
```

The final class declaration looks like this:

```
CountryPicker    CLASS(System.Windows.Forms.UserControl)|
                ,TYPE,PUBLIC,NETCLASS,PARTIAL
m_Countries     ArrayList,PRIVATE
CountryCode     PROPERTY,String,PUBLIC,GETONLY
CountryName     PROPERTY,String,PUBLIC
SelectedIndexChanged EVENT,SelectedIndexChangedEventHandler|
                ,PUBLIC
Construct       PROCEDURE(),PUBLIC
AddCountries    PROCEDURE(),PUBLIC,VIRTUAL
CountryPicker_SizeChanged PROCEDURE(System.Object sender, |
                System.EventArgs e),PUBLIC
CountryComboBox_SelectedIndexChanged PROCEDURE(System.Object sender,|
                System.EventArgs e)
END
```

I want the user of the control to be able to set the CountryName property so I removed the GETONLY attribute from the property declaration and added the setter method to the class.

In the Construct method, after the call to SELF.InitializeComponent, I NEW an ArrayList large enough to hold the 239 country objects I will add and then call the method to add the countries to the ArrayList.

```
SELF.m_Countries = NEW ArrayList(239)
SELF.AddCountries()
```

Then I configure the CountryComboBox by setting the DataSource property to the m_Countries ArrayList, the DisplayMember to 'CountryName' and, the SelectedText to 'Canada'. I frame the code with calls to the BeginUpdate() and EndUpdate() methods to prevent the control from repainting while the properties are being set.

```
SELF.CountryComboBox.BeginUpdate()
SELF.CountryComboBox.DataSource = SELF.m_Countries
SELF.CountryComboBox.DisplayMember = 'CountryName'
SELF.CountryComboBox.SelectedText = 'Canada'
SELF.CountryComboBox.EndUpdate()
```

Next I create the required getter property methods. Here I use a Country object to reference the SelectedItem and return the requested property. Again, this is a matter of style; a less verbose implementation might use the AS or TRYAS operators.

```
CountryPicker.Get_CountryCode PROCEDURE()
SelectedCountry Country
CODE
SelectedCountry = SELF.CountryComboBox.SelectedItem
RETURN SelectedCountry.CountryCode
```

```
CountryPicker.Get_CountryName PROCEDURE()
SelectedCountry Country
CODE
SelectedCountry = SELF.CountryComboBox.SelectedItem
RETURN SelectedCountry.CountryName
```

To set the CountryName property I use the FOREACH operator to search the ArrayList for the matching name and, when found, I set the SelectedIndex property for the ComboBox control.

```
CountryPicker.Set_CountryName PROCEDURE(String strCountryName)
SelectedCountry Country
ndx System.Int32(-1)
arr ArrayList
CODE
arr = SELF.CountryComboBox.DataSource AS ArrayList
FOREACH SelectedCountry IN arr
  ndx += 1
  IF SelectedCountry.CountryName = strCountryName
    BREAK
END
END
SELF.CountryComboBox.SelectedIndex = ndx
```

Now I just need to write the code to fill the ArrayList with countries. The example shown is only partial; the downloadable code contains all 289 country name and code combinations. The code is straight forward; I create a NEW Country object and then Add it to the ArrayList.

```
CountryPicker.AddCountries PROCEDURE()
CODE
SELF.m_Countries.Add(NEW Country('Afghanistan','AF'))
SELF.m_Countries.Add(NEW Country('Albania','AL'))
...
SELF.m_Countries.Add(NEW Country('Zambia','ZM'))
SELF.m_Countries.Add(NEW Country('Zimbabwe','ZW'))
RETURN
```

At this point I do a clean build. There are no errors, so it's time to see if the control actually works.

Testing the CountryPicker control

I close the Keystone.UIControls project and open up the RoundPanelTest project that accompanied the last article. In the previous article I described how to configure the ToolBox. I need to repeat this step again to add my new CountryPicker control to the sidebar. Also, the Keystone.UIControls project was compiled with a newer version of Clarion# so I remove the old reference from the sidebar and add the new Keystone.UIControls assembly and display its controls. I now have both controls in my toolbox (Figure 3).

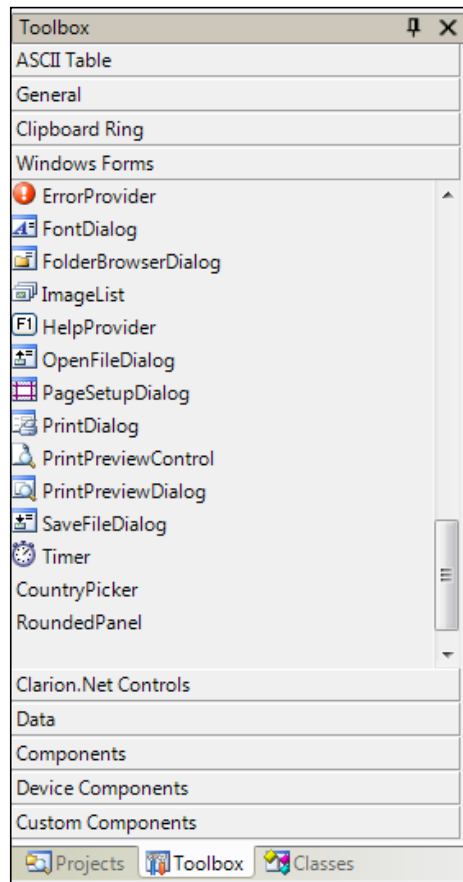


Figure 3. The custom controls added to the toolbox

I select a CountryPicker control then drag and drop it onto my form (countryPicker1). I also select a Label control and place it on the form (renamed to MessageArea). I will use the label to display the country name and code selected by the user as a test for the event handler.

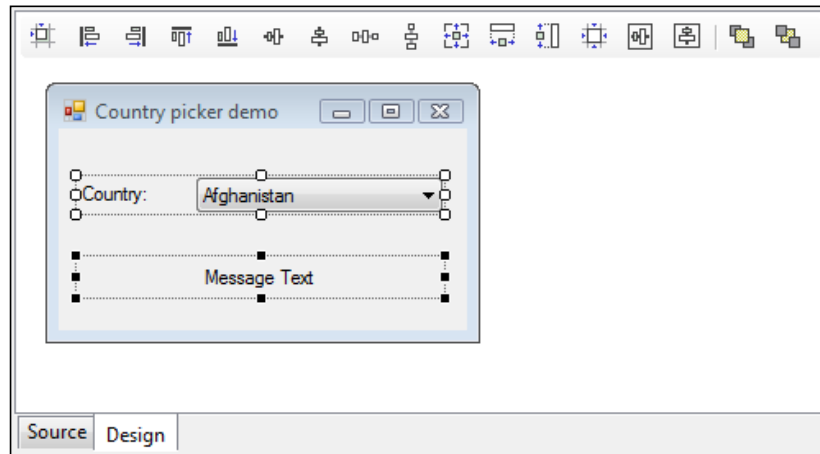


Figure 4. The CountryPicker control and a label control

I select the CountryPicker control. On the Properties pad I switch to the Event view, double click the SelectedIndexChanged event, and add the code to display the CountryName and CountryCode as the MessageArea text .

```
MainForm.CountryPicker1_SelectedIndexChanged PROCEDURE(
    System.Object sender, System.EventArgs e)
CODE
SELF.MessageArea.Text = SELF.countryPicker1.CountryName |
    & ' [' & SELF.countryPicker1.CountryCode & ' ]'
```

Time to do a clean build and see what happens. There are no errors and the program works as expected with the CountryName and CountryCode displayed in the MessageArea when the user makes a new selection.

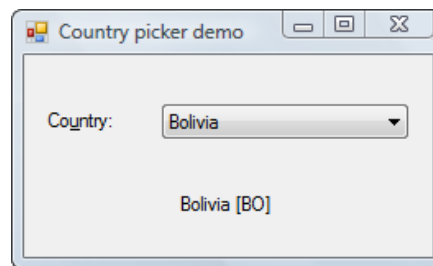


Figure 5. The demo program

There are many enhancements that could be made to the control, such as adding properties for positioning the sub-controls or setting the prompt text. In my live implementation I created a control that lets the user pick a country, then, if US or Canada, the user picks State/Province as appropriate and enters the zip or postal code with proper formatting. Now, anytime I need an address, I just drop my control onto the form and most of the work is done.

In my next article I will cover the application frame and main menu. Until then I hope this article encourages you to explore user controls further.

[Download the source](#)

[Randy Rogers](#) is a data processing professional with over 35 years of experience in a wide variety of industries including accounting, municipal government, insurance, printing, and pharmacoeconomics. He has a degree in Mathematics from Florida State University and is the president of [Keystone Computer Resources](#). Randy is the author of [ClassViewer](#), a utility for browsing the Clarion class hierarchies. He is also the creator of NetTools, Queue Edit-in-Place, and Screen Capture Tools for Clarion application developers.

Reader Comments

[Add a comment](#)

Clarion Magazine

Coloring Template Source In C7/Clarion#

by Dave Harms

Published 2008-07-18

Although the release of Clarion7/Clarion.NET I'm currently using doesn't yet have the AppGen, I've found it useful for writing C6 templates. There are several reasons for this. The new IDE provides a tabbed interface to multiple open files which is a lot easier to work with than a Window menu pick list. And with the project pad I can set up a list of template files I regularly edit.

But the best part of writing templates in the new IDE is color coding. In this article I'll show you how easy it is to add custom coloring rules to your templates (and to any other source files you edit).

Creating a project

Although you don't have to set up a project in order to edit templates, I recommend that you do so. This way you'll be able to easily manage the list of templates you regularly edit.

In this example I'm creating a dedicated project just for template editing but you can also add template files to an existing project. Figure 1 shows the new project dialog. I've highlighted the Clarion.NET category which provides one option: a plain, empty project with nothing in it except references to the System, System.Data and System.Xml namespaces. If you're using Clarion 7 you'll have the choice of Win32 DLL, Win32 EXE or Win32 Library projects.

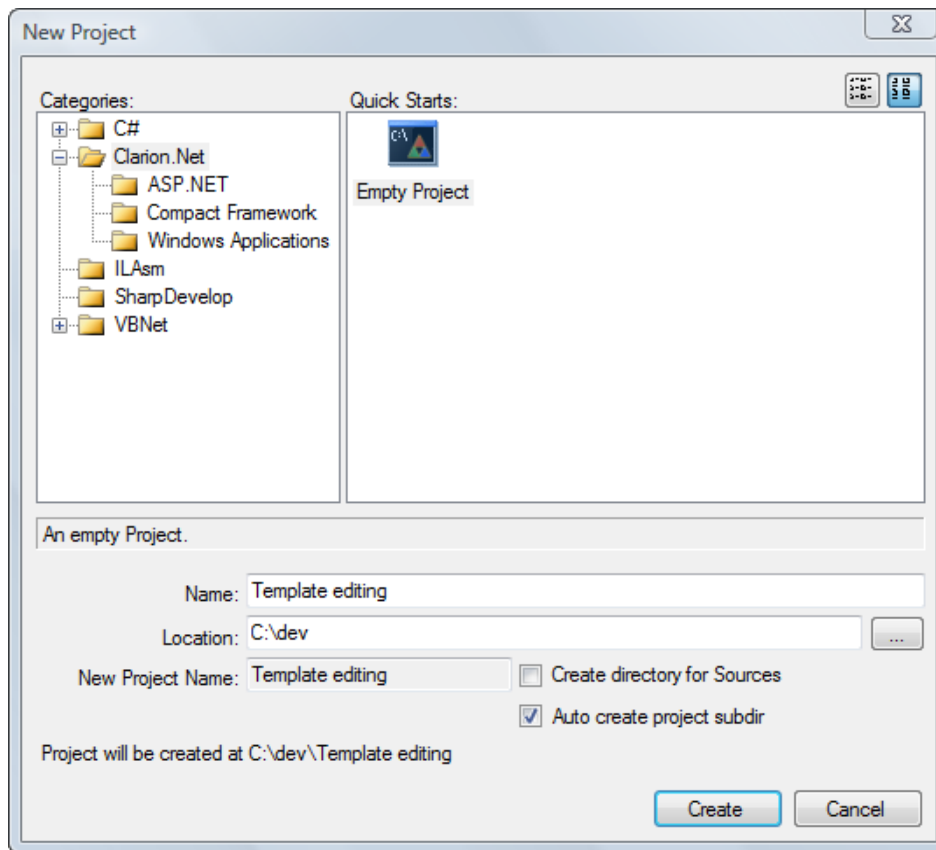


Figure 1. Creating a new project.

If you're using this project just for template editing then you can go ahead and create a new template file or add an existing template file. In either case the template files will not be compiled. But if your project is more than just a way to organize templates, then you may want to create one or more folders first. I like to create one called Templates.

In the project pad right-click on the project and choose Add | New Folder (Figure 2). Give the folder a name.

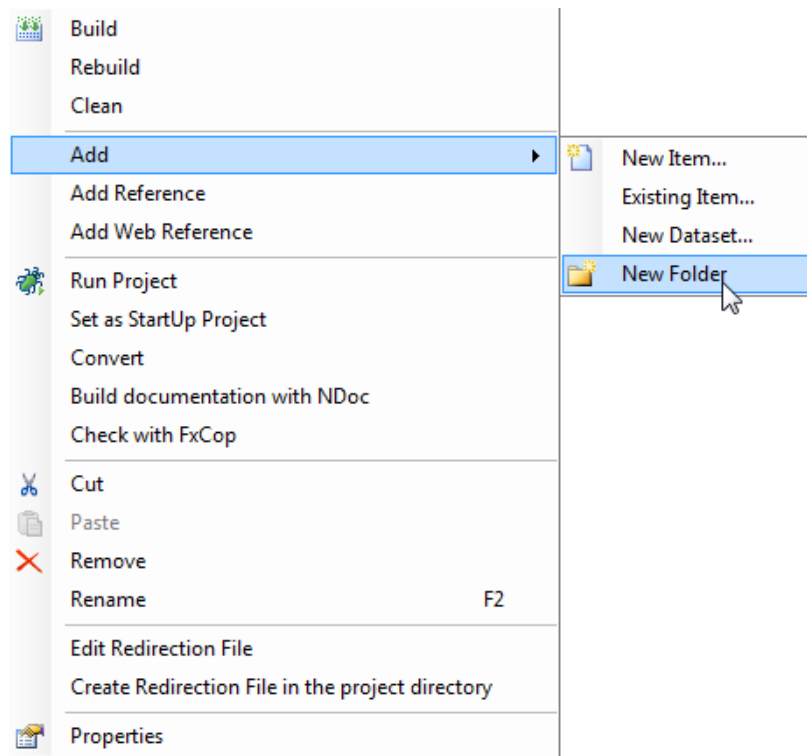
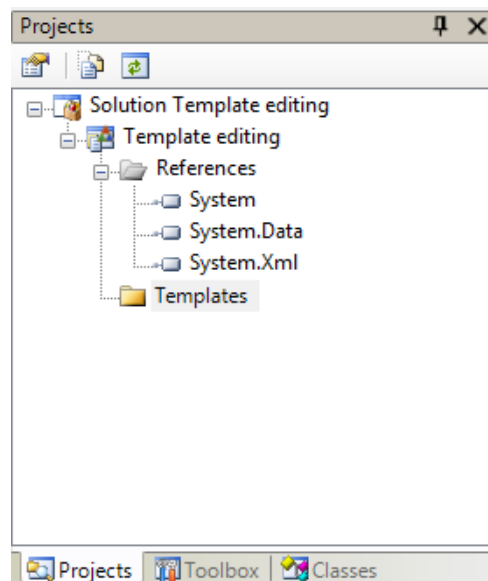


Figure 2. Creating a new folder

Your project pad should now look something like Figure 3.

Figure 3. The project with a Template directory



Now add one or more templates to the project. To add them to the Templates folder, right-click on Templates in the project pad and choose Add | Existing Item as in Figure 4 (if you want to add the template to the project root, do the same but by right-clicking on the project line).




Figure 4. Adding an existing item

A file dialog will appear. Navigate to the directory containing the template file you want to add to your project. You'll probably need to change your file type selector to All Files (*.*) to see the template files. Select the file or files you want to add and click on the Open button (Figure 5).




Figure 5. Adding templates to a project




Figure 6. Creating links to existing files.

The default is Copy, but normally you'll want to use Link rather than copy so you can edit the templates in their usual location.

Now, a word about projects and directories. As you've probably gathered from this example, files in a project don't have to be in the project's application tree. In the above example the following entries are created in my project file:

```
<ItemGroup>
  <Folder Include="Templates" />
  <None Include="..\..\cci\Programs\C6\Template\ABCONTRL.TPW">
    <Link>Templates\ABCONTRL.TPW</Link>
  </None>
  <None Include="..\..\cci\Programs\C6\Template\ABDROPS.TPW">
    <Link>Templates\ABDROPS.TPW</Link>
  </None>
  <None Include="..\..\cci\Programs\C6\Template\ABFILE.TPW">
    <Link>Templates\ABFILE.TPW</Link>
  </None>
  <None Include="..\..\cci\Programs\C6\Template\ABGROUP.TPW">
    <Link>Templates\ABGROUP.TPW</Link>
  </None>
</ItemGroup>
```

I use a non-standard location for C6 on this machine (c:\cci\programs\c6) which is why the above XML may look a little odd, but really all that's happening is the project is storing a relative path pointing to the linked files.

Coloring the templates

With the templates in hand I'm ready to work on color syntax. If I open one of the template files at this stage I don't get any color syntax. That's because there's no default coloring set up for the templates (although this may change once the AppGen is included). Figure 7 shows how to apply coloring rules via a right-click in the editor pad.

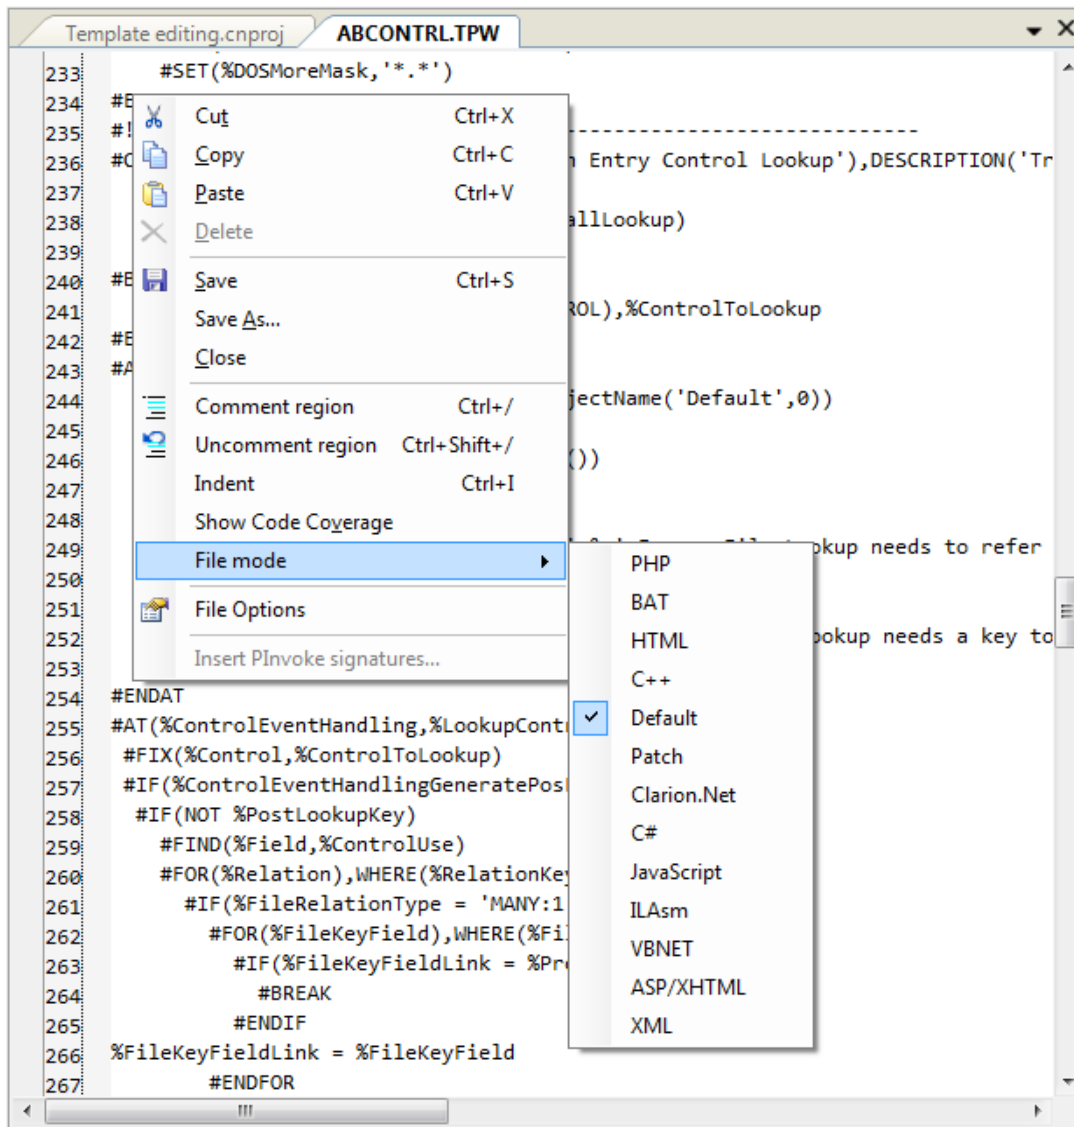


Figure 7. Changing how source is colored

If I select a file mode of Clarion.NET (or Clarion in C7) I'll get some minimal coloring, as shown in Figure 8.

```

86  #!-----
87  #GROUP(%StandardThreadInitialize,%ThreadProcedure,%InitThreadName)
88      #IF(%InitThreadName)
89          #FIND(%ThreadName,%InitThreadName)
90      #ENDIF
91  START(%ThreadProcedure)
92  #!-----
93  #GROUP(%StandardProgressWindow)
94  Progress:Thermometer BYTE
95  ProgressWindow WINDOW('Progress...'),AT(,,142,59),CENTER,TIMER(1),GRAY,DOUBLE
96      PROGRESS,USE(Progress:Thermometer),AT(15,15,111,12),RANGE(0,100)
97      STRING(''),AT(0,3,141,10),USE(?Progress:UserString),CENTER
98      STRING(''),AT(0,30,141,10),USE(?Progress:PctText),CENTER
99      BUTTON('Cancel'),AT(45,42,50,15),USE(?Progress:Cancel)
100     END
101  #!-----
102  #GROUP(%StandardValueAssignment,%AssignTo,%AssignValue)
103  #IF(%AssignValue)
104  %AssignTo = %(%StripPling(%AssignValue))
105  #ENDIF
106  #!-----
107  #GROUP(%StripPling,%Incoming)
108  #IF(SUB(%Incoming,1,1)='!')
109      #RETURN(SUB(%Incoming,2))
110  #ELIF(SUB(%Incoming,1,1)='')
111      #RETURN('EVALUATE(' & %StripPling(SUB(%Incoming, 2)) & ')')
112  #ELIF(UPPER(%Incoming)='TRUE' OR UPPER(%Incoming)='FALSE')
113      #RETURN(%Incoming)
114  #ELIF(UPPER(%Incoming)=''TRUE'' OR UPPER(%Incoming)=''FALSE'')
115      #RETURN(%Incoming)
116  #ELSE
117      #RETURN( '''' & QUOTE(%Incoming) & '''' )
118  #ENDIF
119  #!-----

```

Figure 8. Templates in Clarion.NET/Clarion file mode

To my mind the great thing about having color syntax for templates is to use it to differentiate the template code from the generated code. There's a basic and very easy way to do this, and while it doesn't cover every situation it will go a long ways toward making templates more readable.

Template statements begin with # and symbols begin with %, and I'll show you how to set up two simple rules to deal with those two situations.

From the main menu in the IDE choose Tools | Options.

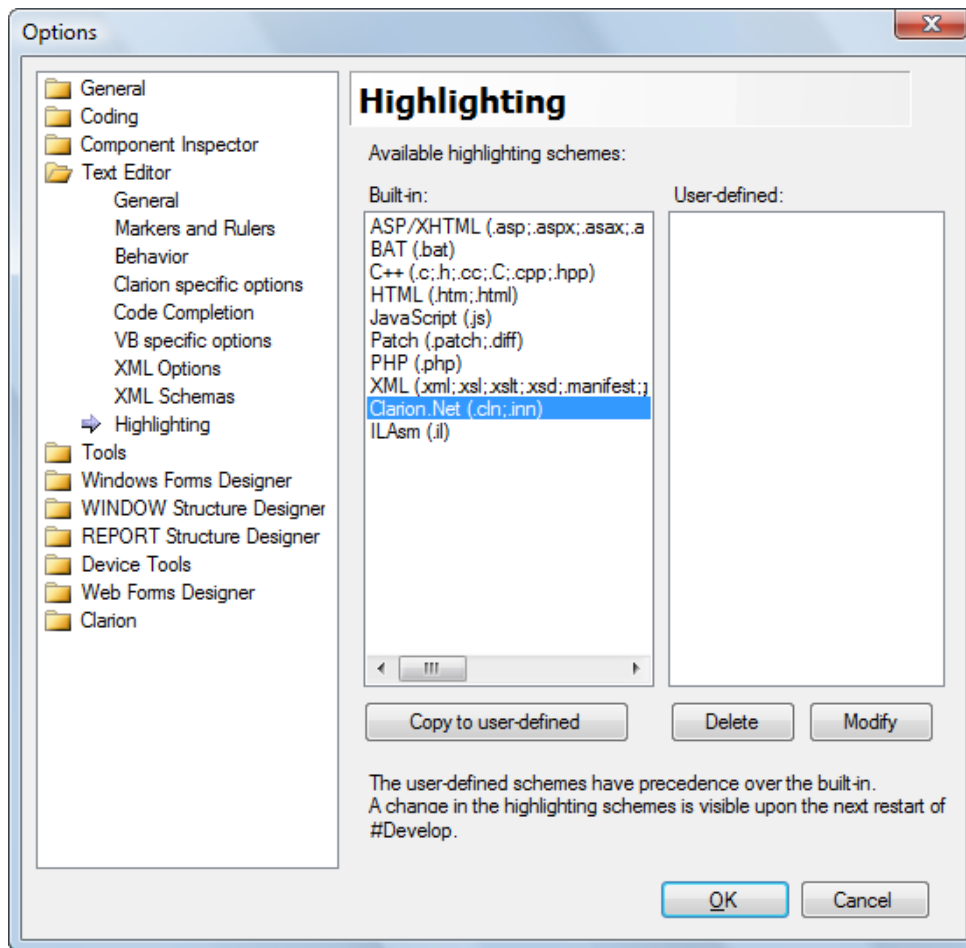


Figure 9. Copying the Clarion highlighting rules

Click on Copy to User Defined, then select the newly copied item and click on Modify. You'll see the window shown in Figure 10.

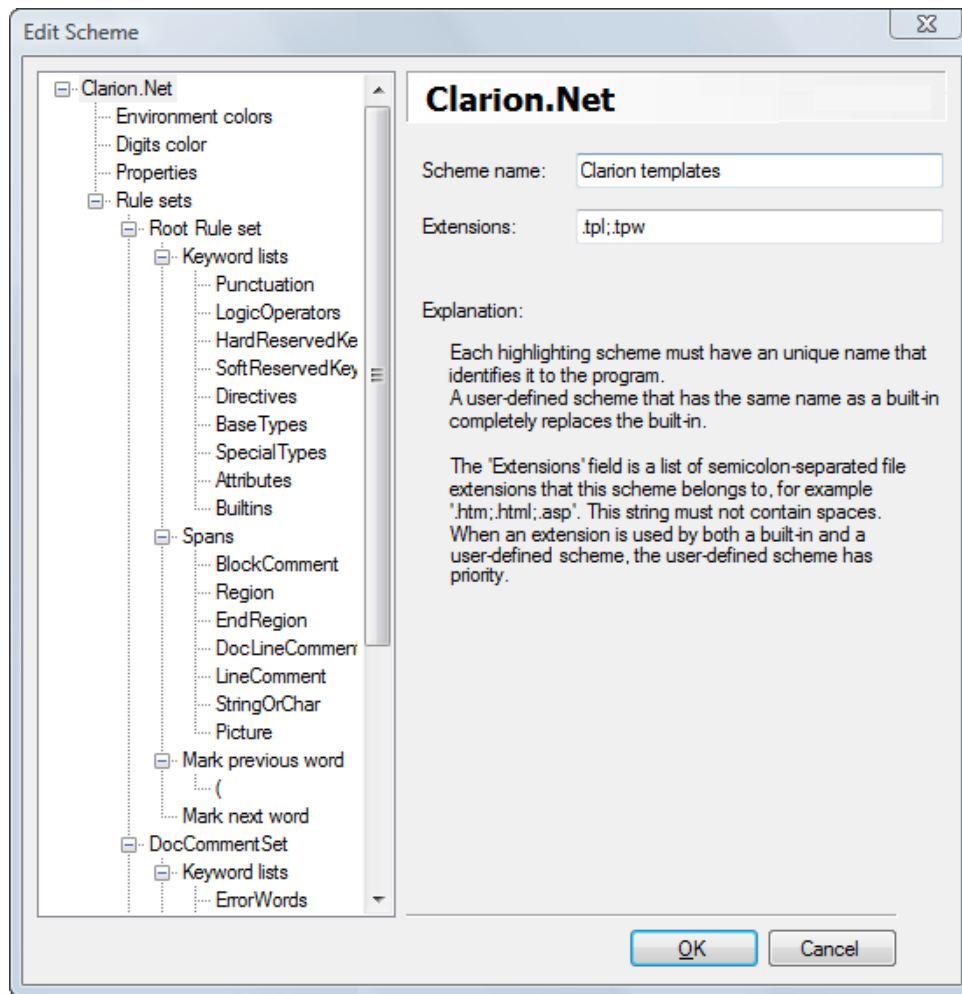


Figure 10. Creating a user-defined rule set

Change the scheme name to Clarion templates and the extensions to .tpl;.tpw to have these settings automatically applied to templates. (As noted in the screen shot you can also replace the built-in scheme by simply not changing the name. I haven't tried this.)

In the tree list on the left side locate Rule sets | Root Rule set | Mark next word.

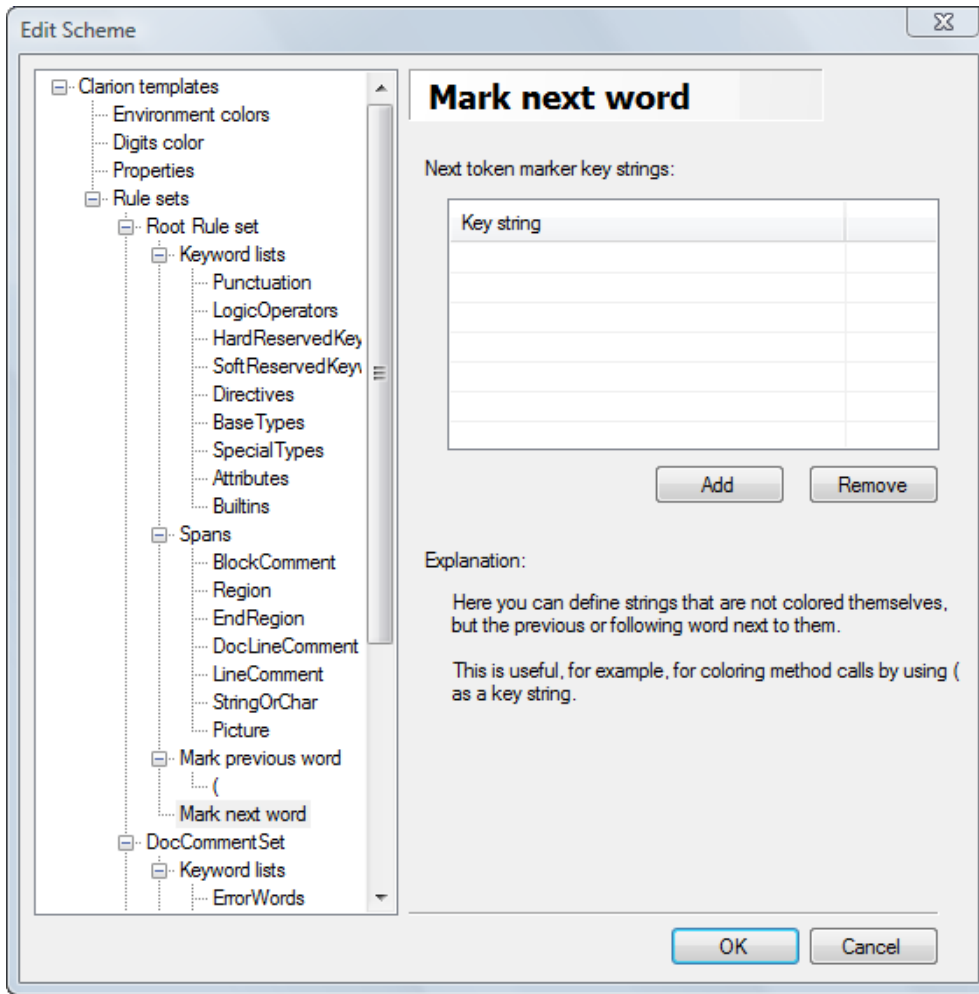



Figure 11. Creating a rule

The list of key strings will (probably) be empty. Click on Add to get the dialog in Figure 12. Type a # character and click OK.



Now here's where it gets just a little bit tricky. As Figure 13 shows the # entry has appeared underneath the Mark next word item. You need to click on that entry to see the options for coloring words beginning with #.

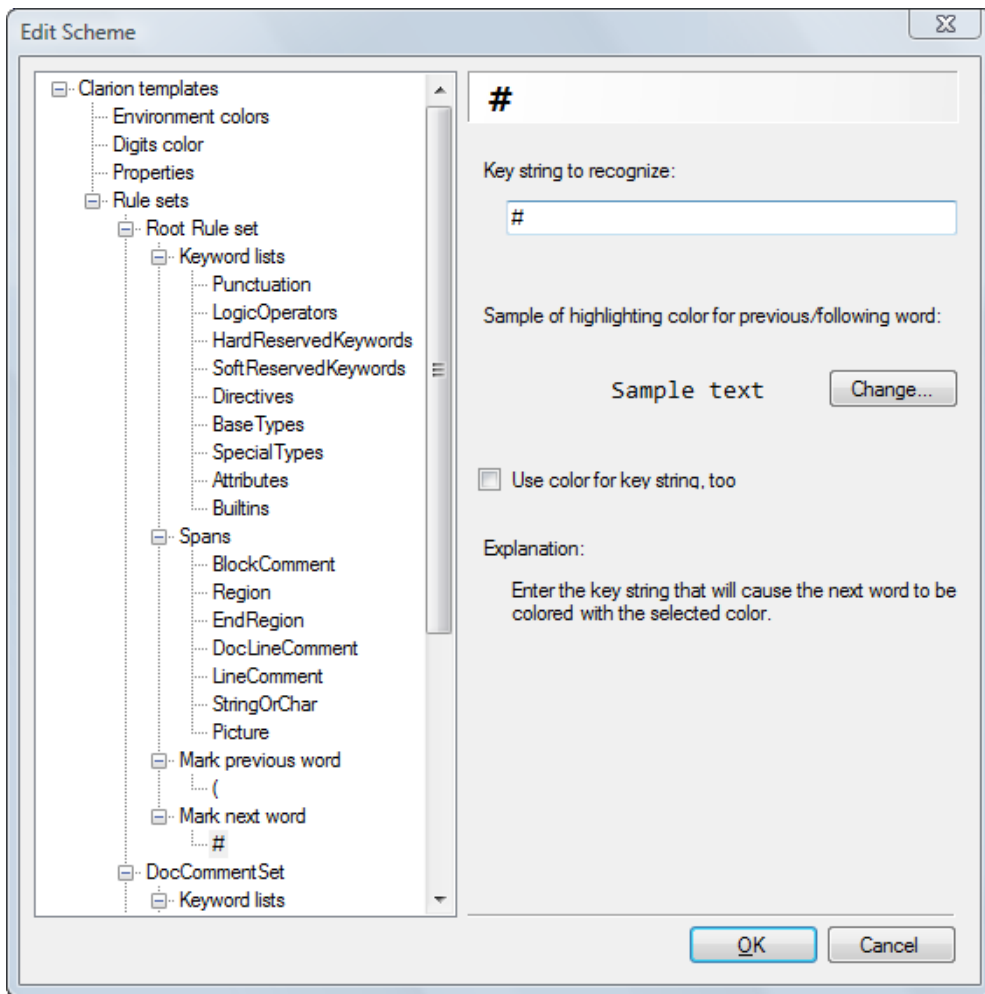


Figure 13. Defining the key string

Click on the Change button to get the options dialog in Figure 14.




Figure 14. Setting the color

Choose a color; I also like to set template statements to bold. Click OK.

There's one more important item. As shown in Figure 15 you need to check the Use color for key string too option.

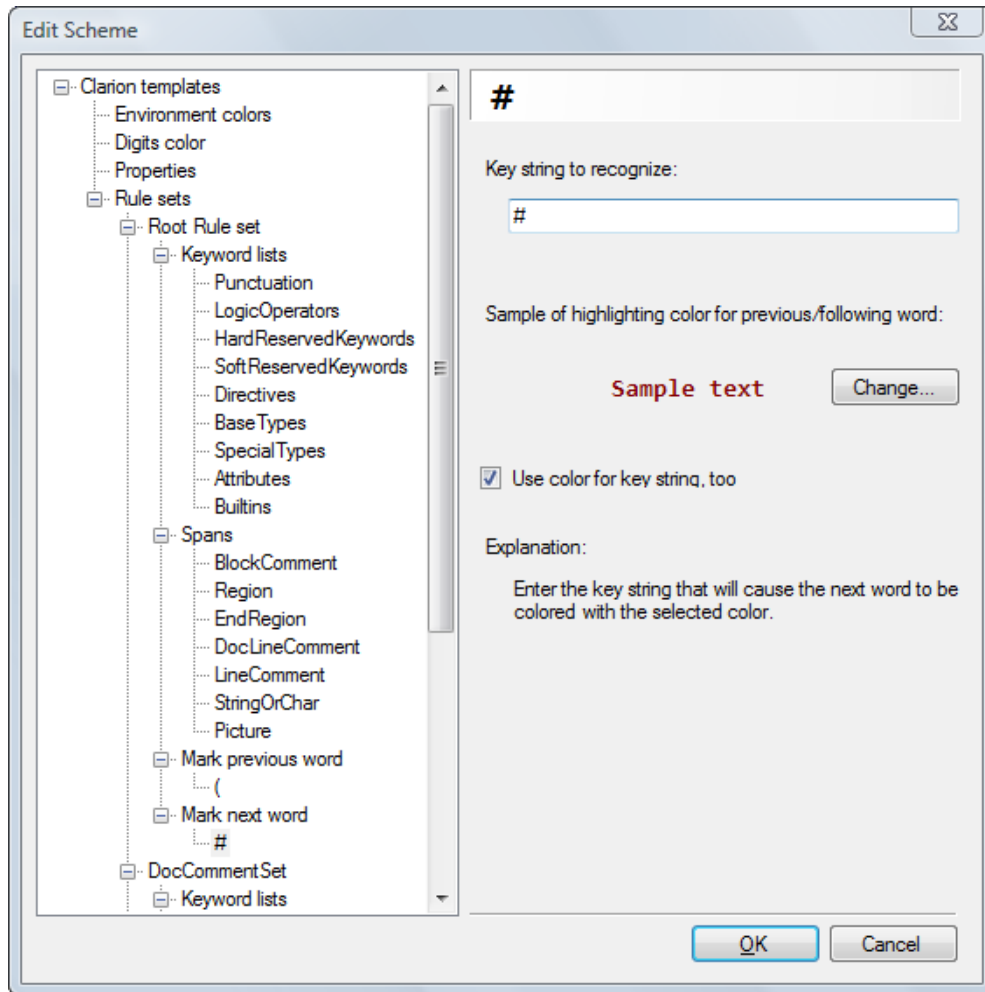


Figure 15. Coloring the key string

Repeat the process with an entry for %. Choose a different color and be sure to again check the Use color for key string too option.

To apply the colors to your source files you can either restart the IDE or just right-click in the editor pad and apply the Clarion Templates file mode. Figure 16 shows the result.

```

86  #!-----
87  #GROUP(%StandardThreadInitialize,%ThreadProcedure,%InitThreadName)
88      #IF(%InitThreadName)
89          #FIND(%ThreadName,%InitThreadName)
90      #ENDIF
91  START(%ThreadProcedure)
92  #!-----
93  #GROUP(%StandardProgressWindow)
94  Progress:Thermometer BYTE
95  ProgressWindow WINDOW('Progress...'),AT(,,142,59),CENTER,TIMER(1),GRAY,DOUBLE
96      PROGRESS,USE(Progress:Thermometer),AT(15,15,111,12),RANGE(0,100)
97      STRING('',AT(0,3,141,10),USE(?Progress:UserString),CENTER
98      STRING('',AT(0,30,141,10),USE(?Progress:PctText),CENTER
99      BUTTON('Cancel'),AT(45,42,50,15),USE(?Progress:Cancel)
100      END
101  #!-----
102  #GROUP(%StandardValueAssignment,%AssignTo,%AssignValue)
103  #IF(%AssignValue)
104  %AssignTo = %(%StripPling(%AssignValue))
105  #ENDIF
106  #!-----
107  #GROUP(%StripPling,%Incoming)
108  #IF(SUB(%Incoming,1,1)='!')
109      #RETURN(SUB(%Incoming,2))
110  #ELSEIF(SUB(%Incoming,1,1)='=')
111      #RETURN('EVALUATE(' & %StripPling(SUB(%Incoming, 2)) & ')')
112  #ELSEIF(UPPER(%Incoming)='TRUE' OR UPPER(%Incoming)='FALSE')
113      #RETURN(%Incoming)
114  #ELSEIF(UPPER(%Incoming)=''TRUE'' OR UPPER(%Incoming)=''FALSE'')
115      #RETURN(%Incoming)
116  #ELSE
117      #RETURN( '''' & QUOTE(%Incoming) & '''' )
118  #ENDIF
119  #!-----

```

Figure 16. Color highlighted template source

Color coding template statements and symbols makes it easier to keep the template code and the generated code separate. I've also found it a direct help in writing symbols; I sometimes forget to include the leading % and the lack of coloring instantly notifies me of the problem.

David Harms is an independent software developer and the editor and publisher of Clarion Magazine. He is also co-author with Ross Santos of *Developing Clarion for Windows Applications*, published by SAMS (1995), and has written or co-written several Java books. David is a member of the American Society of Journalists and Authors (ASJA).

Reader Comments

Posted on Monday, July 21, 2008 by Nardus Swanevelder

Nice article Dave.

One interesting "issue" if you use !~ the IDE handles everything after that as a comment until it gets a ~!

If you use #!~ the IDE still sees it as !~ from a colouring point of view. It doesn't recognise it as a Template comment.

I have used a string like this (#!~~~~~)in my template code for readability purposes and with the new coloring in the IDE all coloring after this line is green. That is if green is your color for comments.

This is not a big issue, I just had to change all those lines to #!*~~~~~

Thanks

Nardus

[Add a comment](#)

Clarion Magazine

OSCON: Orbited Pushes Data To Web Pages

Published 2008-07-25

Orbited is an implementation of Comet, and in programming terms Comet is a web application model where the client, via Javascript, keeps an open connection back to the server. This allows the server to push data back to the client without having to first wait for a request.

Michael Carter and Jacob Rus presented Orbited in the form of a dialog over the confusion between Comet/Orbited and standard TCP socket communication. In essence Orbited is a TCP layer that works over port 80; normal socket communications are often blocked by firewalls/routers. Orbited is also a safe transport layer which can recover from dropped packets.

Currently Orbited is a separate server which runs alongside your web server. In a typical Orbited scenario when the browser requests a page from the web server, the web server communicates with the Orbited server; from that point on the Orbited server communicates directly with the browser via an event handler.

Orbited can be used to present data such as IRC or email (e.g. Google mail) to the browser. One of the topics Carter and Rus argued in their unusually-structured presentation was whether to handle protocols on the server or on the client. In the case of IRC, one choice is to convert IRC to something like JSON for ease of handling in the browser; the other is to push the raw IRC data out to the browser and use Javascript to convert the IRC stream into something displayable. The latter is a more efficient approach since it lightens the server load and leverages computing power on the browser machines.

At present Orbited is only available as a separate server. In part this is because web servers generally spawn a new thread for each client. If you kept all those connections open you could run out of processes. Also Orbited is written in Python to enable rapid development, but Python can't use multiple CPUs because it has a global interpreter lock. If you wanted the current Orbited server to integrate with a web server you'd have to have multiple instances each with its own IP address or port. A multi-threaded integrated server is in testing but isn't expected to be completed until after the Python version has settled down.

Also available is the [js.io](#) library which provides client-level protocol handling. Protocols currently supported or in development include AMQP, IMAP, IRC, LDAP, SMTP, SSH, Stomp, telnet and XMPP. A multi-player browser-base game built on Orbited is available at [unicommand.com](#).

Reader Comments

[Add a comment](#)

Clarion Magazine

OSCON: Open-Sourcing Identity

Published 2008-07-23

Atlassian's Dave O'Flynn announced a new (and as yet unnamed) open source identity management project at OSCON. Atlassian's effort is intended to address the shortcomings of existing identification systems such as LDAP, SAML and OpenID.

O'Flynn began with a discussion of current identity management systems, describing their implementations as painful and rife with annoying problems. He outlined four requirements for such systems:

- Authentication - is are you really you?
- Authorization - what you are allowed to do?
- Info about the user - email, address, etc. Not related to authorization.
- Info about the other users on the system - necessary if you need to interact with those users.

Existing solutions do not meet all of those requirements:

- [LDAP](#) is a standard query language for finding info about people. It's like SQL - it's the same across all databases until you want to do something interesting. You need a complex compatibility layer to handle all the permutations. LDAP supports authentication, authorization, user data, and details about other users. The biggest problem is LDAP is just too hard to work with. It can take a year to learn all the quirks of the different implementations.
- [OpenID](#) has been getting a lot of press. For instance, MySpace is giving all its users an OpenID but oddly you can't use an existing OpenID at MySpace. The big problem with OpenID is that you're asking people to remember a URL. OpenID also doesn't support proper authorization so you still have to store the permissions on your site. OpenID 2 supports attribute exchange, but there's still no way to find out who else is on the system.
- [SAML](#) stands for Security Assertion Markup Language. This is a great idea employing digitally signed XML, but it's too difficult to use. Also any companies that want to use the system need to exchange public keys. SAML isn't useful for discovering other users.
- [Information Cards/Cardspace](#). Microsoft's effort provides some good .NET tools but has similar problems to SAML. There is no authorization or discovery of other users.
- [Acegi/Spring Security](#) - This product is popular with Java developers and provides authentication and basic authorization.


Faced with the limitations of existing identify management solutions Atlassian recently decided to implement and publish its own solution. The product will be a set of REST APIs along with a test suite to allow anyone to create their

own implementation. The data store and the APIs are separable so you can change back ends as needed. In the worst case scenario you'll have to change the URL of the server.

While the authentication portion could be used anywhere identity is needed (such as a publicly available web app) the authorization part of the API is more likely to find use within organizations where credentials can be internally administrated. This API won't be suitable for all situations; the goal is to handle the easy 80%.

Atlassian hopes to have a basic version out for Christmas-ish. Chances are there will be a few revisions of the API, and once the specifications have settled down they'd like to take the API to one of the standards organizations.

Interoperability is important - this system needs to be able to work with existing systems



like SAML, and will not invalidate that investment. The focus is to make identify management easier for app developers. Licensing will be permissive, most likely something like BSD or Apache.

Reader Comments

[Add a comment](#)

Clarion Magazine

OSCON: Open Source Gains Ground In The Enterprise

Published 2008-07-23

An [O'Reilly Radar report](#) released today estimates that 10% of all new IT jobs involve at least some open source software. The report cites explosive growth in open source projects, draws on analysis of job posting data and considers a number of case studies.

Open source adoption data isn't easily obtained; traditionally analysts look to vendor reports but that data isn't available where software can be freely downloaded. Open source itself has experienced rapid growth, however: SourceForge, the largest repository of open source code, has grown from 12,500 projects in 2000 to almost 200,000 projects by the end of 2007, an annually compounded growth rate of 55%. SourceForge anticipates reaching one billion annual downloads by next year.

In partnership with SimplyHired, O'Reilly analyzed job posting data from about half of the Fortune 1000. One in 52 enterprise job postings were classified as "open source", a high number given that 2.3% of all jobs are IT jobs. On further investigation this was found to be due to the presence of some large technology companies (such as Google, Yahoo!, Sun) in the sample set. For typical non-technology companies open source positions represented .2%, or approximately 10% of IT jobs.

Hiring requirements can only suggest software adoption, but recruitment data also indicates the rate of acceptance of individual products. The LAMP stack is still experiencing solid growth as is [Ruby on Rails](#), and Linux remains the most in-demand open source product based on job postings. Content management frameworks such as [Django](#), [Drupal](#) and [Alfresco](#) are prominent among the products showing fastest growth although overall numbers are still very small.

Factors influencing the adoption of open source by enterprise include adaptability to new requirements, code quality and security, and cost. In many parts of the world, the absence of contractual obligations to foreign software companies is an important factor. The O'Reilly report illustrates the interplay of these factors with a variety of case studies and provides guidance for creating an open source action plan.

The report is available in PDF only from [O'Reilly](#) for US\$399.

Reader Comments

[Add a comment](#)

Clarion Magazine

OSCON: Pro PostgreSQL

Published 2008-07-22

In his Pro PostgreSQL tutorial session Robert Treat covered a variety of topics including understanding the nature of the PostgreSQL community, performance optimizations, upgrade paths, availability and hardware. Treat is a database architect at [OmniTI](#) and co-author of the book [Beginning PHP and PostgreSQL 8](#). He also maintains the [phpPgAdmin](#) software package and has been recognized as a major contributor to the PostgreSQL project for his work over the years.



Robert began with an overview of the nature of the PostgreSQL open source community. Unlike commercial databases there is no one corporate backer deciding the future direction of PostgreSQL. Rather, changes to the product happen via an organic process involving a variety of groups.

The focus of the development community is the [pgsql-hackers mailing list](#). Anyone can get involved at this level, but most of these folks don't actually have the rights to commit changes to the source code. Instead they write and test their own code and submit patches. These patches can only be applied to the official source code by committers, a much smaller group of hackers numbering a dozen or so. There is no formal process to become a committer - just contribute a lot of good patches and in time you may be invited to become a committer.

At the top of the community hierarchy is the core team which has seven members. Not all of these people are committers. For instance Mark Fournier focuses on release management, and Josh Berkus (lately of Sun) is the media guy.

Resources

A big part of getting along with PostgreSQL is knowing what resources are available; to some extent you have to figure out for yourself what you need. Your starting point should be [www.postgresql.org](#) (this site is maintained by the PostgreSQL web team).

Other useful sites include:

- [PGfoundry.org](#) (a.k.a. [projects.postgresql.org](#)) is a good place to start looking for ready-made PostgreSQL solutions. Click on the Project tree tab for a central repository.
- [planetpostgresql.org](#) is a blog aggregator. Possibly more important than PGFoundry.
- [archives.postgresql.org](#) contains mailing list archives back to '97, with full text search.
- [IRC chat](#)
- The IRC channel also has a bot called [rtfm_please](#) - to query it on the IRC channel use two commas followed by the search term.

Installing

You should be using a package manager/installer appropriate to your platform. If you're installing on Windows things are pretty simple - just use the Windows installer. Installing on *nix is potentially more complicated: if you're dealing with multiple platforms superuser accounts and command line utilities can vary. Package management only makes things consistent per platform, not across platforms.

The core package includes the database server as well as commonly-used utilities. But there's a whole lot of contrib (user contributed) code as well. On Windows this is part of the install; on other platforms it's usually a separate install. One caveat on contrib modules: upgrades can be a hassle of the contrib code depends on anything that was changed in the core PostgreSQL code. In general the contrib code is very solid. "Trust these more than your own code."

One of the more popular contrib modules is [pgcrypto](#), which is not included in the core due to export restrictions. For help with tuning look at [pgstatstuple](#), [pg_buffercache](#), and [pg_freemap](#).

There are various procedural languages available ([four](#) ship with the core code but there are about a dozen). Some are

not included because they're released under different licenses, or because of dependency management issues, or because they're still pretty new. PL/Perl is probably the best of the bunch.

Logging

Step one after installation is to figure out logging. You'll definitely want this so you can work out exactly what errors you are getting, since the interface you're using to the database probably won't give you detailed information. Usually the logs are in `pg_log` under the data directory (`$PGDATA`) - check `postgresql.conf` if in doubt of the location. On *nix the log files may be in `/var/log/`. If you're able to log in to the console then show `data_directory` will reveal the file locations.

For performance reasons it's a good idea to keep the log files and the data files on separate drives.

Authentication

Authentication is the next hurdle. If you don't see an error in the log (look for FATAL) then check your firewall because you're probably not hitting the database at all. Check your `pg_hba.conf` configuration.

TRUST authentication is okay if you have a secure network or you're just doing development, but md5 is a better way to go as it offers decent security. Treat didn't have much good to say about IDENT which requires that your user name is the same as the PostgreSQL database user name. This is fine if it's what you want, but it can cause a lot of headaches if you're not aware of how IDENT works.

There are other ways to handle authentication but these are the big three.

Upgrades

Treat was candid about the problems with PostgreSQL upgrades. There is no in-place upgrade tool included with the core distribution so you will typically have to handle major version upgrades yourself. Normally "major" means any upgrade in which either the first or the second version number has changed; usually third version number changes are bug fixes/features only and do not involve database changes, just new binaries. All major upgrades involve pain, and skipping releases usually makes things worse. Always, always make a backup first and be sure to read the release notes.

There is a tool called `pg_migrator` which attempts to handle in-place upgrades. Right now it only works for 8.1 -> 8.2, but there is no 8.2 -> 8.3 path at present (and maybe never - the next path may be 8.3->8.4). `pg_migrator` works by rewriting the system catalogs, and there's no way to go back to the previous version. It's pretty new.

The simplest way to upgrade is to use `pg_dump` from the old database and `pg_restore` to the new database. Treat recommends the `-Fc` option for compressed files (on the order of a tenth of the size of the uncompressed file). If you have a large database this is going to be a long, long process, on the order of hours. And you pretty much need to have both versions of PostgreSQL installed at the same time. Use the `pg_dump` from the newer version, not the older version.

Note that `pg_dump` 8.2 has problems with versions prior to 7.3. If you're on a prior version then you need to upgrade to 7.3, run `adddedeps` to fix the system catalog info, then upgrade to 8.3.

There are several ways you can avoid lengthy downtime. One is to replicate your database (using `slony`, `postgres-r` etc.) to the newer version. When you're sure everything's working properly just cut over to the new server. With `slony`, which is a master-slave system, you'll need to set the new server to be the master, but that's quickly done and downtime may be under a minute. But be aware of network load issues while you're doing the initial recreation of the slave (newer format) database.

Configuration

Keep in mind that PostgreSQL ships with very conservative settings to enable it to run on lowly hardware with minimal available memory. That almost certainly means it won't run optimally on your system.

There are around 200 config options, but only about 20 will really matter for most DBAs. The top five are:

- `effective_cache_size` - You probably want to use whatever ram is available after your other apps are up and running. This memory will only be used if required. 75% on a dedicated machine is probably about right.
- `shared_buffers` - This does preallocate memory so don't set it higher than it needs to be. Start at 20% of your ram and tweak downwards. It all depends on the nature of your database traffic. This requires a database server restart so get it right before you into production.
- `default_statistics_target` - The default (10) is probably too low. Set it to 100 as a starting point. The higher you set it the more data sampling it will do. Set it really high (say 400) is probably too high for most situations.
- `work_mem` - This is memory for sorting/hasing inside an SQL query. You may want to set this higher, but this is the max amount allocated for certain tasks but on a per-thread basis, so be careful if you have a lot of concurrent processes.
- `checkpoint_segments` - These segments are chunks of data being written to disk. If you have a lot of write activity you want to increase this number. A value of 30 isn't unheard of. If fast recovery time after a crash matters then you want to keep

this number lower since if you have more of these files the database will take longer to process them on restarting.

Keeping it clean

PostgreSQL's vacuum command cleans up the database by reclaiming unused space, updating table stats, and preventing transaction ID wraparound. In previous versions you had to run vacuum explicitly, but PostgreSQL now has an autovacuum mode which you can successfully use with the default settings. There may be times when you need to run vacuum manually, usually only on tables that get excessive write traffic. Vacuum has a lazy and a full mode - full locks the table, and you should never have to run it regularly.

You may find situations where you need to reindex tables. This area has improved with recent releases but keep it in mind if you run into index performance problems.

Hardware

Don't go cheap on the hardware. Think lots of ECC RAM, redundant disks, RAID (hardware or software, but software RAID can be easier to fix if you run into problems. For hard drives SCSI is best but SATA is often acceptable. If you're on Solaris consider RAID Z. Good hardware solves a lot of problems.

PostgreSQL makes good use of multiple CPUs. You get great payback up to 8 CPUs, or even 16, but by 32 you've probably maxed your bang for buck.

On 8.3 one processor isn't better than another per se. See tweakers.net for benchmarks - they do a lot of testing with PostgreSQL as well as MySQL. The site is mostly in Dutch but there's still lots of good info to be had. (My tip: Google with the search term postgresql site:tweakers.net).

When using multiple disk configurations the usual rules apply - don't mix access types if you can help it. For instance, put write-ahead logs on their own disk (RAID 1 is fine for that) since these are all sequential writes. Put data directory on its own disk. Treat recommended RAID 10 over RAID 5 which he said was too slow (although one attendee suggested that after eight spindles you won't see a difference).

The more spindles the better. More controllers is even better if you want to use tablespaces. Separate tables into tablespaces by the kind of access they need - read vs write.

Postgresql scales up well - just throw more hardware at it. Use connection pooling if necessary but a single machine, properly equipped, can handle 1000+ connections with terabytes of data and 10,000 transactions per second.

Watch out for write-caching controllers - if you need this be sure to use battery backed controllers to avoid situations where the server crashes and you lose transactions.

On *nix systems do not use the NFS operating system; it is not fully safe.

Backups and availability

Treat went over a number of options for ensuring data remains available and recoverable.

On *nix you can do file system snapshots with zfs, lvm. Wrap the snapshot with `pg_start_backup` and `pg_stop_backup` for safety.

You can also exploit PostgreSQL's [point-in-time-recovery](#) (PITR) feature which employs write-ahead log (WAL) files. If you take a snapshot of the database all you have to do to make it current is to replay any WAL files created since the last snapshot, and you don't have to be exact about when the logs start as long as they include the missing data.

If you're using replication with PostgreSQL then you're probably talking master-slave replication, and there's always the possibility that if the master goes down there will be transactions that never made it out to the slave. But you can do a very quick switchover. Another replication option is [bucardo](#), which was open sourced a while ago. It's multi-master that also does master-slave, but it isn't yet available on Windows. Newer yet is [postgres-r](#) which is multi-master synchronous and open sourced just this month. It looks promising, but these are still early days.

There are a variety of other techniques for ensuring availability which rely on either duplicated data or just duplicated server processes. Shared disks let you run multiple servers on one data set, but you still have a single point of failure in the data store. You could try [pgpool](#) which is a statement-based dual master system; all your inserts and updates go to two different machines. There's no synchronization; rather you're relying on the SQL being executed identically on both machines. Be aware that `random()`, `now()`, and sequences could be different so even if all the statements execute the data stores might not be identical. More problematic is you have to run it as a superuser - you can't do normal authentication. And [pgpool](#) itself becomes a single point of failure.

Scalability

You can use replication to scale up your database, although in most cases this works best for scaling read operations since writes would have to be replicated on all machines. You can also use [pgpool-II](#) (not really related to [pgpool](#)) to put

a single database across multiple machines. This product comes from the Japanese PostgreSQL community so the doc can be difficult to access. pgcluster is similar to pgbpool-II and is a synchronous multi-master but scales read operations only.

The core team has announced that at some point they will officially support PITR read-only slaves, and there's a push on for warm standby.

Treat also mentioned some performance improvement strategies including delving into the pg_stat* system tables and exploiting the EXPLAIN statement. Set the log_min_duration_statement to log slow queries. Keep in mind that PostgreSQL supports partial indexing with WHERE clauses and functional/expressional indexes. Full text search is much faster than LIKE expressions and has support for multiple languages, custom dictionaries, special indexing options. Full text search is integrated into 8.3.

Comments

Treat's presentation was professional, thorough and clear. Although he's clearly a big fan of PostgreSQL he didn't shy away from the problems inherent in a community-based open source project, and he was blunt about PostgreSQL's weaknesses particularly when it comes up database upgrades. I'd like to see in-place upgrades (don't hold your breath) and a solid master-master replication option (perhaps postgres-r will fill that niche), but in its current state PostgreSQL is clearly capable of some pretty heavy lifting.

Additional Resources

- [Robert Treat's blog](#)
- Clarion Magazine's [PostgreSQL articles](#)

Reader Comments

[Add a comment](#)

Clarion Magazine

OSCON: Even Faster Web Sites

by Dave Harms

Published 2008-07-31

Steve Souders is the former Chief Performance Yahoo! and is now doing similar performance-related work at Google.

Souders' contribution to web site performance was to turn the standard analysis upside down. Instead of focusing on server-side performance he turned his attention to end-user issues. In any given web page most of the load time happens on the client side, not the server side.

On iGoogle's pages only 9-17% of page load time is taken up by the server, depending on whether the content is cached. Generally client side processing is in excess of 90%. If scalability is the issue, focus on the back end. Otherwise focus on the front end, and you should probably start there anyway. Cut the front end in half and you can save 40-45% of the user experience.

You should at least be able to cut 25% off your load times.

Souders came up with [14 rules](#) for high performance web sites. These are codified in the book "[High Performance Web Sites](#)" and in the [YSlow](#) plugin for [Firebug](#). HPWS volume 2 is in the works, and this talk focused on the first three of ten rules discussed in the book. Souders is also the author of [Cuzillion](#) which is a test web page generator - Cuzillion is short for "cuz there are a zillion pages to check."

Cuzillion lets you create test pages with various elements; you can then drag those elements to different locations and change the elements' properties. Souders demonstrated Cuzillion using a private beta release of [HttpWatch](#) for [Firebug](#). He modified his test page, reloaded it, and viewed the results in time-chart format which made it easy to see how much time was taken up with different page components.

Javascript is often responsible for slow page loading because other file downloads are blocked until the script download is completed. One way around this is to split the initial payload; load only the necessary Javascript with the page and get the rest via the onload event. You can also trick the browser into downloading Javascript files in parallel. There are six techniques for this, including XHR eval, XHR injection, putting the script in an Iframe, putting the script in a DOM element, using script defer (only in IE), and using document.write which is not recommended. All of these techniques vary in how they trigger the browser's "busy" indicators. XHR doesn't trigger any indicators, and other techniques may differ across browsers.

Long-executing inline scripts also block rendering and downloads. You have to relinquish control every 300 ms or so with a setTimeout. In IE you can use the defer attribute for inline scripts as well.

FF blocks parallel downloads while downloading stylesheets. In IE they don't unless the stylesheet is followed by an inline script. And be sure to use link not @import for stylesheets.

Souders concluded with a few announcements and suggestions. Not surprisingly he recommends [Yslow](#) for web site testing. [HTTPWatch](#) is the best web packet sniffer out there, and well worth the price. The current release version is only available for IE but the FireFox version should be out in a month or two so you'll be able to get those cool time charts in Firebug. Don't use the Net panel in Firebug 1.05 as it's very buggy. If you are on IE, Opera or Safari you can use [Firebug Lite](#).

[David Harms](#) is an independent software developer and the editor and publisher of Clarion Magazine. He is also co-author with Ross Santos of *Developing Clarion for Windows Applications*, published by SAMS (1995), and has written or co-written several Java books. David is a member of the American Society of Journalists and Authors ([ASJA](#)).

Reader Comments

[Add a comment](#)

Clarion Magazine

OSCON: Using Cascading Style Sheets

by Dave Harms

Published 2008-07-31

Gavin Doughtie is a front end developer on Google's Picasa Web Albums and a contributor to the [dojo toolkit](#).

Doughtie covered a number of cascading style sheet (CSS) features that can take the place of Javascript code. There are costs associated with Javascript, including download time, parsing, and runtime performance. Other costs include ridicule for using the language (laughter).

CSS fundamentals include flow (the way elements are displayed in a page), float (determining where an element will appear within another element), and positioning (specifying absolute or relative position). CSS isn't always the best solution - for tabular data you're usually still better off using tables.

You can use pseudo-selectors like `:hover` in place of Javascript, but there are often cross-browser issues. IE6 only supports `:hover` for the anchor tag, so you need to use an empty anchor tag and a div inside your inline tag, which you're not supposed to do.

To vertically center a div use:

```
top: 50%;  
margin-top: -50%;
```

For tables set `table-layout` to `fixed`. To draw fancy borders you only need three images: one for the four corners, another for sides, another for top and bottom. At each position only show the portion of the image you need to show.

You can use relative measurements for image resizing. Set the image size in ems and have a container element for the image. Then change the font size for that container. This can be faster than similar Javascript code: the biggest benefit is on IE6, the FireFox, but on Safari image resizing with Javascript runs faster than the CSS approach.

Future capabilities

[WebKit](#) is an open source browser engine that's used by Safari and the iPhone. WebKit provides many advanced capabilities such as rotating images and text and using SVG masks (where you open a window on an underlying image). FireFox 3 has a new rendering engine based on [Cairo](#) and can do a lot of the fancy stuff like WebKit.

WebKit's market penetration is very small, but it's worth noting that it's available on 100% of iPhones.

Doughtie closed by ranting on how IE6's many failings are holding back web development. His advice: just drop IE6 support (if you can).

[David Harms](#) is an independent software developer and the editor and publisher of Clarion Magazine. He is also co-author with

Ross Santos of *Developing Clarion for Windows Applications*, published by SAMS (1995), and has written or co-written several Java books. David is a member of the American Society of Journalists and Authors ([ASJA](#)).

Reader Comments

[Add a comment](#)

Clarion Magazine

OSCON: Secrets of Javascript Libraries

Published 2008-07-24

John Resig is a JavaScript Evangelist for the Mozilla Corporation and the author of the book [Pro Javascript Techniques](#). He's also the creator and lead developer of the [jQuery](#) JavaScript library.

Resig's presentation covered a variety of Javascript programming techniques for cross-browser support but he also touched on some fundamentals. There are three core aspects to Javascript; if you fully understand these you're well on your way to mastering the language:

- Functions
- Closures
- Function prototypes

Resig also listed his three favorite JS libraries: [Prototype](#), [jQuery](#) (naturally) and [base2](#).

According to Resig testing Javascript can be ridiculously easy; begin with an `assert(expression,output)` function which can be implemented any way you like. In typical testing you might also add an event listener on page load; this way you can queue up any assert calls so they are displayed after the page has loaded.

Javascript is single-threaded; as a result you can use `setInterval` like an automanaged queue that will continue to take tests, run them, and run the next test when that test is done. use `setInterval(1)`; on IE `setInterval(0)` is interpreted as `setTimeout`.

Resig suggested a strategy for achieving cross-platform compatibility:

- Pick your browsers
- Know your enemies
- Write your code

You probably don't want to support every browser out there. In terms of cost/benefit, IE6 has the highest cost and second highest benefit. IE7 has the second highest cost and highest benefit. Firefox has very low cost and the third highest benefit. Choose Safari 3 over Opera 9.5.

Yahoo publishes a [graded browser support list](#). This covers 99% of all users.

What really matters is the cost benefit to the users, so you may end up supporting all browsers anyway, meaning the previous version, the current version (the most commonly used version), and the next version.

There are five ways your reusable code can be attacked (military terminology used advisedly):

- Browser bugs
- Regressions (next browser release)
- Bug fixes (next browser release)
- Missing features
- External code and markup

Your main focus will be browser bugs/differences. Your best defense is a very good test suite, and your best offense is feature simulation where your code provides features missing in a particular browser. Make your code resistant to outside influences, and don't break outside Javascript code. One way to achieve this is by strictly enforcing namespace use. If iterating through properties use `hasOwnProperty` to avoid iterating over something added by another library.

If your stylesheets are specified after your Javascript code it isn't guaranteed that the stylesheets will be loaded by the time your code executes, so always put the stylesheets first.

Use object detection to work out whether the browser is missing features, then gracefully degrade, all the way back to a simplified page if necessary. For instance, if `elem.addEventListener` doesn't exist use `elem.attachEvent`. Consider checking for all your requirements at the start, and if they aren't available then go to your fallback position.

Don't assume browser bugs will always be there - a browser bug fix can break your code.

Feature simulation is more sophisticated than object detection; you make sure the API works as expected and you capture bug fixes gracefully. If you go by the published standard then you're ready for the bug fix. Change out your browser detection code to checking in a function for the existence of the feature itself and implementing the workaround or the alternate technique.

Regressions are much more difficult to handle, but object detection can help. Always test your code during the beta process.

Just avoid doing things that cause known problems where there's no real benefit. For example, don't set negative height/width. Limit your code to the things that do work across all known browsers.

Some problems are untestable:

- Has an event handler been bound?
- Will an event fire?
- Do CSS properties like color actually apply?
- Problems that cause a browser crash

And some tests are impractical:

- Performance
- Verifying AJAX requests

Writing cross-browser Javascript code is all about reducing the number of assumptions you make, but you can't achieve perfection.

Function definition

Javascript lets you declare both named and anonymous functions, and these can be assigned to variables and object properties. You can define functions anywhere in the current scope and they will be available everywhere. You can even put a function after the return statement in another function. The code in the enclosing function will terminate before the enclosed function is reached, but that doesn't stop the function from being callable. Order matters when you're assigning, not when you're defining.

Recursive functions work as you'd expect.

When creating anonymous functions assigned to properties you can give those anonymous functions a name; these are "named anonymous" functions. But the scope of that name is only within the anonymous function itself, not outside.

The arguments property contains what is generally an array of passed arguments. It also has a property called callee which refers to the function itself. You can use arguments.callee within a function to call that function recursively.

Functions are also objects, and you can arbitrarily attach properties to both functions and objects. Resig showed an example of Javascript code that created a cache of functions, which can be useful when dealing with events.

When you do assignments the result of the assignment is the value that was put in. If you want to know if the object exists you can use a double ! operator (!!) which effectively translates the object (or its absence) into true (or false).

Resig discussed execution context and the use of this, which defines the current object. That may not be an explicitly declared object: if you define a function globally then the function is actually a property of the global object.

You can modify the context with .call - the first parameter passed to .call is the new context. You can also use .apply; while .call takes individual arguments, .apply takes an array of arguments. You can also use the native array methods with .call to simulate an array.

All functions have a length property which is the number of arguments expected.

Javascript doesn't allow function overloading by default - a second function by the same name will overwrite the first function. Instead you add methods to objects.

Functions can have two different states. Besides normal state they can be instantiated. You can create a function prototype inside a function and then instantiate it with new.

Resig demonstrated a syntax that was a sort of temporary anonymous function. In its simplest form it looks like (function(){ }) () - this is a function that is immediately created, run and then destroyed. It's kind of like block level code - whatever you do here will never be seen outside this code. You can also put a function like this in a parameter list.

Closures

Closures can be a difficult concept to grasp. Javascript functions can contain both variables and other functions; this is a little like a function being an object containing methods and properties. And if the enclosing function (a.k.a. the object containing the inner function and variables) terminates by returning a reference to one of its inner functions then the enclosing function and everything it contains is not destroyed as you might expect. Instead that inner function (to which you have a reference) can still access the variables and functions inside the enclosing function. Closures do come at a cost since the reference needs to be maintained.

Among other things closure is useful for library wrapping; you can hide most of what's in the library and just expose the bits you want.

Reader Comments

[Add a comment](#)

Clarion Magazine

OSCON: A/B Testing

by Dave Harms

Published 2008-07-24

Ben Tilly has been active in the Perl community for many years, and is on The Perl Foundation grant committee. He is currently the Reporting Architect at Pictage, Inc where his responsibilities include A/B testing.

Ben Tilly described A/B testing is the process of creating two versions of a web page (or other code) and evaluating users responses in order to optimize the user's interaction. For many sites this translates as maximizing conversion of visitors to customers.

Some of the points to keep in mind as you test:

- Remove any unneeded form fields
- Add relevant form fields - ask the users the questions they want to be asked.
- Use marketing landing pages - when you bring someone to your site (i.e. via a search engine or affiliate page) take them to a page that's relevant to their interest.
- Keep forms short - consolidate help links and other information on a linked page.
- Interstitial pages can help if they're only shown to people who really need to see them, for instance when you ask a further question based on something specified in a form.

A/B testing is not a substitute for:

- Talking to users
- Actually watching people using your software
- Acceptance tests (checking off use cases or other specifications)\
- Unit tests
- Thinking about what you're doing. A/B testing is a way to test what you already believe to be good; it's a very expensive way to sift through all the possible things you might want to do.

In order to evaluate your A/B testing you need some statistical analysis tools. Tilly's advice is to use the [G-test](#), which was created by Karl Pearson in 1900. The G-test is a close relative of the chi-square and is slightly better.

The G-test only answers yes/no questions and only handles two versions. That means you need to keep your changes small. If you're modifying a bunch of things on the page (form fields at the same time as color, graphics etc.) you're not going to



have a clear idea of what changes make a difference.

With two groups of users and two pages you have a 2x2 matrix with four numbers: GroupA_Yes, GroupA_No, GroupB_Yes, GroupB_No. You keep the test running until the G-test (ideally) shows 99% confidence. If the differences are small and your web site traffic isn't particularly high you may need to run your tests for weeks or even months; a high traffic site may yield results in hours or days, although it's important to note that depending on your users you may get different results depending on the time of day, day of week, or part of the month.

Tilly provides some analysis tools on his [web page](#), but you'll have to work out how to handle the mechanics of testing for your own web site or other software. Build a framework and in particular automate reporting so you can get daily reports without having to do any additional work. You'll

want to track how your tests are coming along so you know when to end them. Write lots of tests, and don't keep failed code around. The point of A/B testing is to eliminate stuff you don't want.

When a user enters the test you'll need to assign them to one group or the other. You want to be able to force a person into one group or another, and you also need a way to end the test and hook in the good code permanently. Easy termination is also useful if there's a bug in the test and it needs to be shut down.

Use ticket numbers in test names and standardize your naming scheme. use descriptive test and version names.

Suggested columns:

- Person id
- Test name
- Version name
- Creation date
- Unique key on person id and test name

An important concept in web site testing is the conversion funnel. The point at which users enter our web site is the widest part of the funnel; at each step the funnel narrows. Start at the top of the funnel and work on improving each step.

Work out what your metrics are. If you have ads you have conflicting metrics - you make money when they click on the ad but then they leave your site.

You could easily end up with 20-40 metrics. If the metric is the thing you're trying to change and you have 99% confidence in your tests, you can believe in the test. ut if you have a set of metrics be careful about trusting in just one metric over the others. If several metrics agree, that's good. Conflicting metrics require a business decision.

Remember you are testing unique people not just page hits and link clicks.

When you change a page some users will resist change and have trouble adapting. If this is an issue you may need to restrict the testing to new users.

Metrics alone won't save you, although they can alert you to the problem. At rent.com a wrongly worded email subject cost the company about 4 million dollars over nine months until someone noticed the subject line didn't match the content.

David Harms is an independent software developer and the editor and publisher of Clarion Magazine. He is also co-author with Ross Santos of *Developing Clarion for Windows Applications*, published by SAMS (1995), and has written or co-written several Java books. David is a member of the American Society of Journalists and Authors ([ASJA](#)).

Reader Comments

[Add a comment](#)

Clarion Magazine

OSCON: More to come

Published 2008-07-25

Blogs on this site

- [»All Blog Entries](#)
- [»Clarion 7 Clarion.NET](#)
- [»Future Articles](#)
- [»News flashes](#)
- [»Nifty Stuff](#)

I hope you've been enjoying this year's OSCON coverage. I have one session left to attend and then I'm done. But I still have a bunch of notes to post, on everything from some interesting things said during keynotes to securing databases to web development. Plus some other news and info and a bunch of pictures I haven't had a chance to review yet. Most if not all of that will have to wait until next week. And I have some of the more regular Clarion fare lined up as well, so there's lots of good stuff in the pipeline.

Clarion Magazine

Ouch

Published 2008-07-25

Blogs on this site

- [»All Blog Entries](#)
- [»Clarion 7 Clarion.NET](#)
- [»Future Articles](#)
- [»News flashes](#)
- [»Nifty Stuff](#)

Sun's Tim Bray is speaking on changes in programming languages. And I have to say this is possibly the most annoying use of background music I have ever heard. I'd rather listen to my daughter playing Build-a-Bear on her DS.

Ah, it turns out the music was written by Terry Riley, in C, in 1964. So there's a reason it's being played.

Just not a good reason.

Clarion Magazine

MS Becomes Apache Platinum Sponsor

Published 2008-07-25

Blogs on this site

- [»All Blog Entries](#)
- [»Clarion 7 Clarion.NET](#)
- [»Future Articles](#)
- [»News flashes](#)
- [»Nifty Stuff](#)

Three years ago MS sponsored one of the OSCON lunches, and some attendees were so offended they refused to eat the lunch. Microsoft has continued to attend and increase its participation in open source. This morning Microsoft's Sam Ramji and the ASF's Justin Erenkrantz announced that Microsoft has become a platinum sponsor of the Apache Software Foundation.

Clarion Magazine

OSCON: Wednesday review

Published 2008-07-23

Blogs on this site

- [»All Blog Entries](#)
- [»Clarion 7 Clarion.NET](#)
- [»Future Articles](#)
- [»News flashes](#)
- [»Nifty Stuff](#)

As expected it's been a busy day. The morning keynotes began at 8:45 a.m. and the last sessions closed down just after 6 p.m. Highlights for me today included Josh Berkus on securing PostgreSQL (lots of food for thought there) and Atlassian's proposed open source identify management API (which put me in mind of the recent ClarionMag articles on LDAP). I'll have an article up on the Atlassian project shortly; I'm still going over my PostgreSQL notes and I'll also have something to say about the keynotes, plus I have a few more pics of the conference.

Clarion Magazine

Wednesday welcome

Published 2008-07-23

Blogs on this site

- [»All Blog Entries](#)
- [»Clarion 7 Clarion.NET](#)
- [»Future Articles](#)
- [»News flashes](#)
- [»Nifty Stuff](#)

In the conference center parkade this morning I spotted an old beater with a couple of stickers in the rear view window. One simply said "Linux", and the other said "Coding is not a crime."

Dude. I've seen code that was worthy of punishment.

It's just about time for the keynote addresses. Meanwhile I have a couple of tutorial reports from yesterday that need some cleaning up - hopefully I'll get a chance to post them before the day is out, but I have a pretty fully schedule alerady.

Ah, we're rolling now with a video parodying video game development.

"When you create sound effects, do they come out of your head?"

"No, they come out of my mouth."

Heh heh.

Clarion Magazine

OSCON gets under way

Published 2008-07-21

Blogs on this site

- »All Blog Entries
- »Clarion 7 Clarion.NET
- »Future Articles
- »News flashes
- »Nifty Stuff

[Australia aside](#), there's been a dearth of Clarion conferences lately. But that doesn't necessarily mean there aren't conferences out there that are of interest to Clarion developers. This week I'm at [OSCON 2008](#) which is O'Reilly's annual open source conference held, as usual, at the [Oregon Convention Center](#) in [Portland, Oregon](#). Portland's something of a hotbed of open source development and attendance is usually around 2500-3000, which is a wee bit more than your typical Clarion conference.

But this is an open source conference, and Clarion is a closed source product. What exactly is ClarionMag doing here?

In my experience there are plenty of Clarion developers already using open source products and tools. For the most part that means databases such as [MySQL](#), [Firebird](#), [PostgreSQL](#) and [Ingres](#), and web tools, of which there are a great many ([Apache](#) and [PHP](#) are popular choices).

At OSCON I'm focusing on the database and web tracks. Today and Tuesday are tutorial sessions; the conference proper begins on Wednesday. This morning I took in Robert Treat's informative PostgreSQL session (more on that coming up in the mag); this afternoon is a tough choice between test-driven development and web site A/B testing. Tomorrow's schedule sports some web-related sessions including one on Javascript libraries (if you're doing web development and not yet using Javascript just wait a bit - you will).

So far the mainstream buzz at OSCON is all on the mobile side, where Linux-based operating systems are gaining market share. At the moment that doesn't have a lot of bearing on the work most Clarion developers do, so I'll leave that coverage to the [big dogs](#). Meanwhile there's plenty of database- and web-related goodness to come.

Clarion Magazine

ClarionMag is at OSCON 2008

Published 2008-07-19

Blogs on this site

- [»All Blog Entries](#)
- [»Clarion 7 Clarion.NET](#)
- [»Future Articles](#)
- [»News flashes](#)
- [»Nifty Stuff](#)

Next week I'll be in Portland Oregon for [OSCON 2008](#), O'Reilly's open source conference which typically draws around 2500 developers.

You may be wondering what I'm doing at an open source conference, given that Clarion is a closed source platform. In fact many Clarion developers use open source tools, particularly databases such as MySQL, PostgreSQL and Firebird. Quite a few Clarion developers also use open source web tools such as PHP. And while open source isn't quite as common in the Windows world as it is among Linux users, it's gaining ground.

At OSCON I'll be paying particular attention to the database and web tracks but I'll be reporting on anything I think is of interest to ClarionMag readers. Stay tuned!

Clarion Magazine

How to corrupt an APP

Published 2008-07-02

Blogs on this site

- »All Blog Entries
- »Clarion 7 Clarion.NET
- »Future Articles
- »News flashes
- »Nifty Stuff

Sometimes software can be in use for years before someone stumbles across a major bug. The other day [Lee White](#) needed to add an image resource to a C6 application, and to help pick out the file he needed from a directory with a lot of images he began the file name with --, as in --yada--.bmp.

The app saved and compiled fine, but the image file corrupted the project. Result: next time Lee opened the app he got a "Syntax error in project file" and the .APP file was unusable.

The problem won't happen if you have a single leading dash in the file name; you need two dashes. Lee suspects the parser is getting confused because two dashes are used to denote comments:

```
#compile "junk_BC0.CLW" -- GENERATED
#compile "JUNK_BC.CLW" -- GENERATED
#compile "junk.clw" -- GENERATED
#compile "junk001.clw" -- GENERATED
```

The solution is to hex edit the APP file, find the offending file name in the project data and change the -- to some other string. After that your app should be usable again. Just remember to make a backup copy *before* you do the hex edit.