

RZ 3653 (# 99663) 03/20/06
Computer Science 7 pages

Research Report

Automatic User Interface Composition by a Semantic Web Client

Daniela Bourges-Waldegg, Marcel Graf, Christian Hoertnagl

IBM Research GmbH
Zurich Research Laboratory
8803 Rüschlikon
Switzerland
{dbw, grm, hoe}@zurich.ibm.com

LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties). Some reports are available at <http://domino.watson.ibm.com/library/Cyberdig.nsf/home>.

IBM Research
Almaden · Austin · Beijing · Delhi · Haifa · T.J. Watson · Tokyo · Zurich

Automatic User Interface Composition by a Semantic Web Client

Daniela Bourges-Waldegg, Marcel Graf, Christian Hoertnagl
*IBM Zurich Research Laboratory, Säeumerstrasse 4
CH-8803 Rüschlikon, Switzerland
{dbw, grm, hoe}@zurich.ibm.com*

The presented demo is a simple web-based application for managing a wish list of products, in this case books. It uses an RDF-based model comprising books, authors and keywords and benefits from a generic programming model with facets that automatically attach context-sensitive commands based on what data and metadata arrives to the application; a first prototype is currently realized using a JSP (JavaServer Pages) tag library. Applying the framework to another example domain is straightforward, because model and views are well isolated from the generic parts. Knowledge workers testing this proof-of-concept application will make generous use of mouse right-clicks for activating dynamically composed popup menus. Summary points:

- RDF serves as unifying technology for both persisting the model and driving the user interface. The Jena semantic web platform mediates all manipulations of the model both while in memory and while stored across servers on the network, in that the demo application is currently programmed against the Jena API.
- By default, commands attach to information regardless of its presentation context.
- Because information is semantically tagged, it is straightforward to e.g. sort names correctly by surnames or to link to other services such as Amazon.com¹.
- The JSP tag library spreads work between the client (JavaScript) and server (servlets).

On the server side, this prototype implementation relies on the Apache Tomcat servlet container and Jena 2.3 libraries; the only extra piece necessary on the thin clients is a small JavaScript library that is automatically shipped with the content. The following seven screenshots comprise a representative sample demo walk through.

1. Main demo application portal

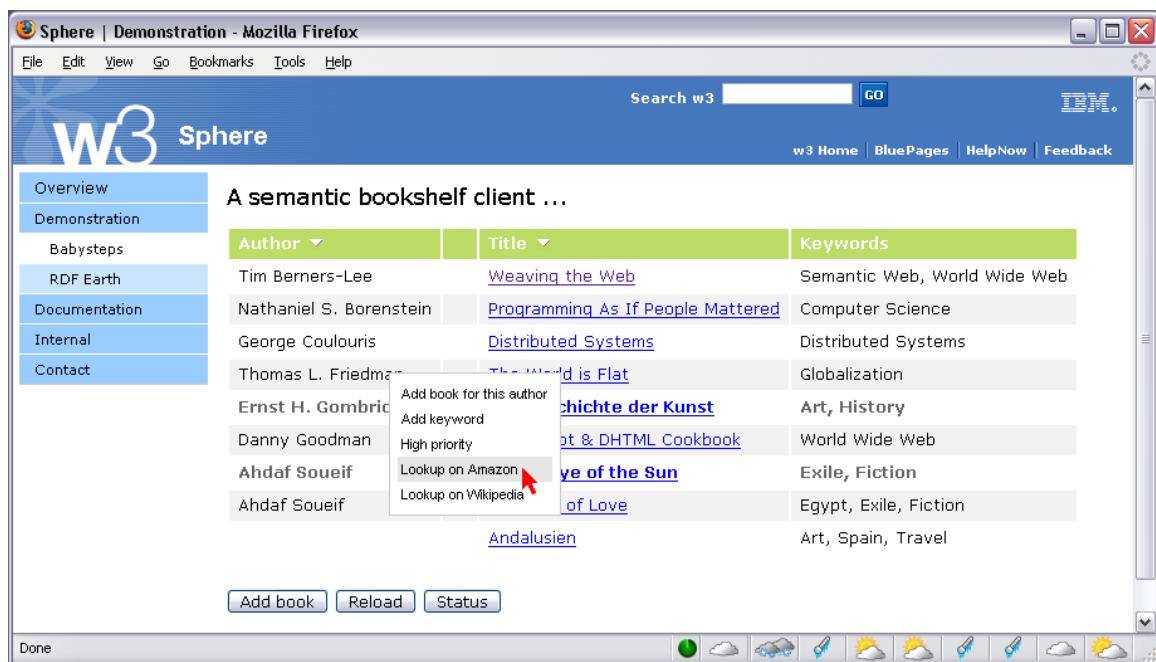
The main application portal looks at first like a conventional web application. However, the server-side data is captured in an RDF graph, and the content displayed at the client preserves this semantic information. The exploitation of such information to tailor the user interface is what is specific to the presented thin client. When a user performs graphical interactions (e.g. right-clicking on items, following hyperlinks, sorting items, etc.) this information is used to offer functionality that is best suited to the specific kind of data.

In the example screenshot (Screenshot 1), this leads to the presentation of a popup menu for performing application-specific functions in relation to books and their

¹ Amazon.com is a registered trademark of Amazon.com, Inc.

associated keywords – in the particular case shown, the popup menu, displayed after right-clicking on an author, offers in addition author-related functionality. This and other popup menus are created dynamically, and there is no piece of code or other description that explicitly binds the specific listed items together, except for the implicit fact that they all become relevant (and are hence automatically offered) in a particular usage situation.

The framework’s exploitation of semantic information in the underlying model is also visible in relation to its sorting rules: author names are sorted by surnames-then-given-names, whereas looking up author information in Wikipedia² (command entry visible in popup menu) requires a given-name-then-surname use of the same information. The underlying semantic model makes it straightforward to distinguish between given names, surnames and other types of strings, and to convert between different human-friendly notations of abstract information.



Screenshot 1: main application portal, showing a dynamically generated pop-up menu that offers commands relevant in the particular semantic context

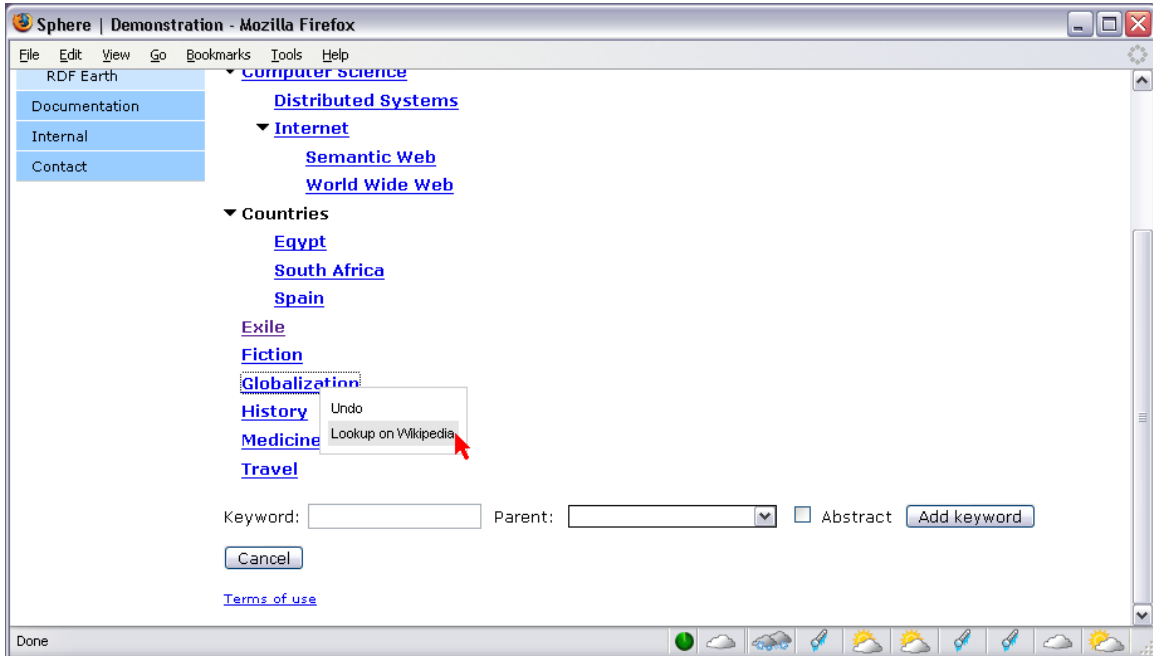
2. Facets encapsulate self-contained user interface behavior

The next screenshot shows another web page that belongs to the same application: it can be accessed by choosing the “Add keyword” entry from a popup menu, such as the one shown before. Since the popup menus and other parts of the user interface are constructed dynamically by the framework, it is already clear that this entry can appear in multiple places, and this gives the user convenience in manipulating information elements (in this case keywords), in principle regardless of their contexts of appearance.

For instance, consider that the popup entry “Lookup on Wikipedia” appears on both Screenshots 1 and 2, although there is a semantic difference: Screenshot 1 was captured

² Wikipedia is a registered trademark of the non-for-profit Wikimedia Foundation,

after right-clicking on an author name: This Wikipedia lookup will therefore reveal information about the author. If we had right-clicked on a keyword (to the right) instead, a different choice of appropriate entries would have been offered, and Wikipedia lookup would have revealed information about the keyword term. This is also the case with the options offered in the popup menu displayed in Screenshot 2 .

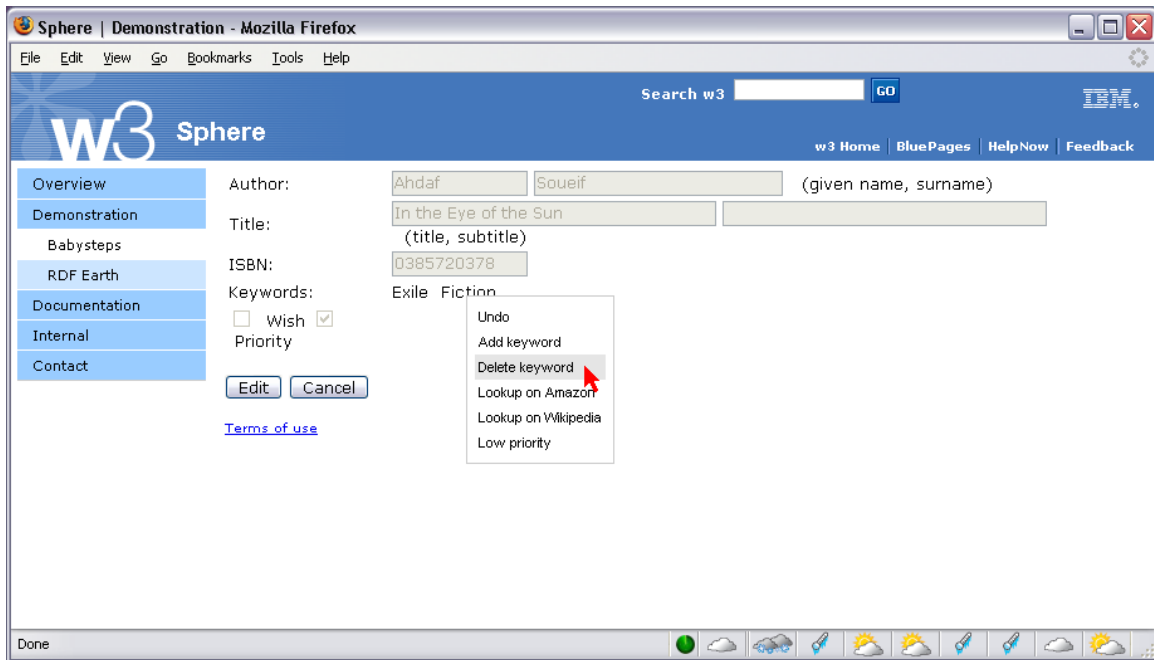


Screenshot 2: add-keyword page, showing popup menu associated to a keyword; commands have different semantics depending on their context of invocation

3. Different user interface manipulations can trigger facets

Adding and deleting keywords, as well as all other commands that have side-effects, can be undone by selecting the generic command “Undo”.

Not only popup menu entries trigger command executions by facets: submitting an electronic form, following hyperlinks (e.g for collapsing the tree view in Screenshot 2), and other graphical manipulations by a user lead to similar effects, and they employ the same basic abstractions in the underlying framework. For instance, changing the interface style such that book details (Screenshot 3) can be accessed by issuing commands from popup menus instead of by following hyperlinks (e.g. book titles in Screenshot 1) requires only changing a Java method signature and a JSP library tag in one place each. Supporting both styles similarly requires adding a Java method (signature) and a JSP library tag.



Screenshot 3: book details are displayed after following a hyperlink; this behavior can be also attached to a popup menu command

4. The demo application interacts with other services.

The next Screenshot shows a Wikipedia article as use case for integration of other services into the demo. In this particular case, a simple HTTP redirect occurred as a result of the selection of a popup menu entry e.g. in Screenshot 2. However, interactions between the demo application and other existing applications could be much richer; in the particular sample domain e.g. Amazon Web Services (AWS) could be used to find the average rating of a selected book by those Amazon.com reviewers whose correlation with local ratings (see “High priority” in Screenshot 1) are above a certain numeric threshold.

It is a fundamental property of the underlying RDF technology, that models can be distributed across servers (e.g. the user providing some ratings, and more arriving from AWS). Important scaling effects can be achieved from expressing metadata (e.g. indication of what is a positive rating as opposed to a positive value in another semantic domain) in equivalent terms, so the necessary data mixing and matching can be automated in software. The OWL (Web Ontology Language) layer that builds on top of RDF helps to coalesce or align ontologies if necessary (e.g. this demo application uses binary ratings, whereas other sites use 1-to-5 stars, but the semantic mapping between the two is straightforward).

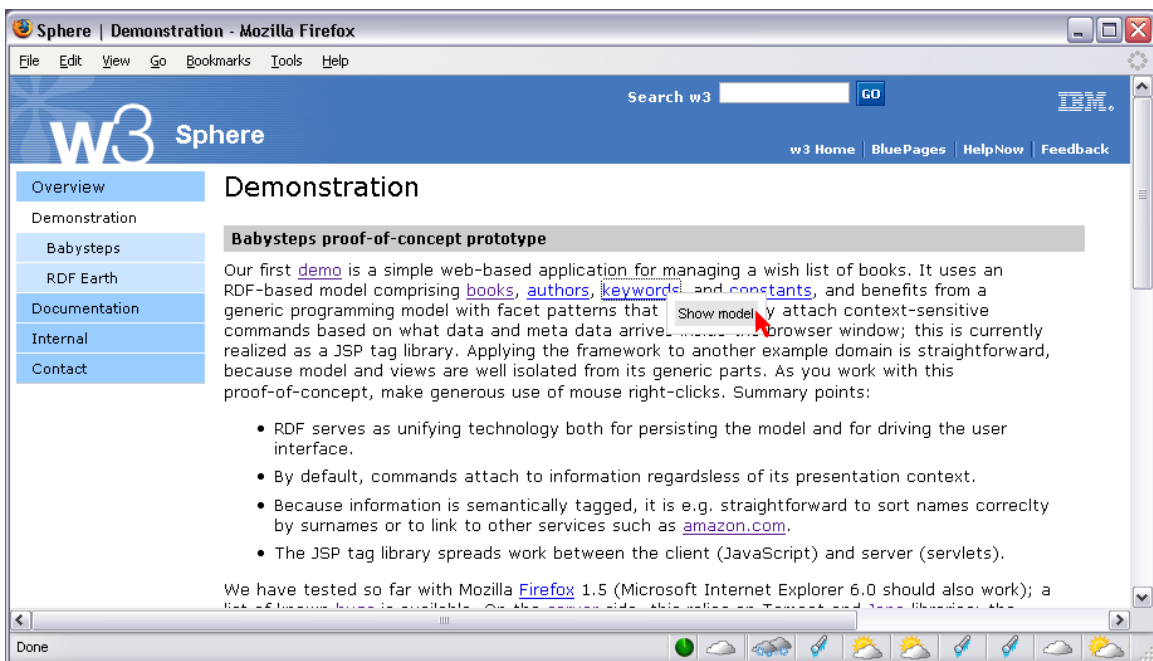
The multiple opportunities to combine lateral sources of information that are afforded by Semantic Web base technologies, leads to a greater amount of overall (partly inferred) information and hence a greater variety of possible manipulations in a user interface. The presented demo application and its underlying framework are representative of a more general approach that composes aspects of user interfaces from dynamic properties of underlying models, but because of the aforementioned scaling this approach makes particular sense for models that encode unified semantic information, as is the case here.



Screenshot 4: linking to an external service

5. The framework can enrich existing pages with legacy content

The next Screenshot demonstrates the feasibility of accessing facets also from legacy web-pages. The only change that is required concerns introduction of JSP library tags at those places where verbatim strings alone cannot convey the semantic category of data to which certain facets should attach.



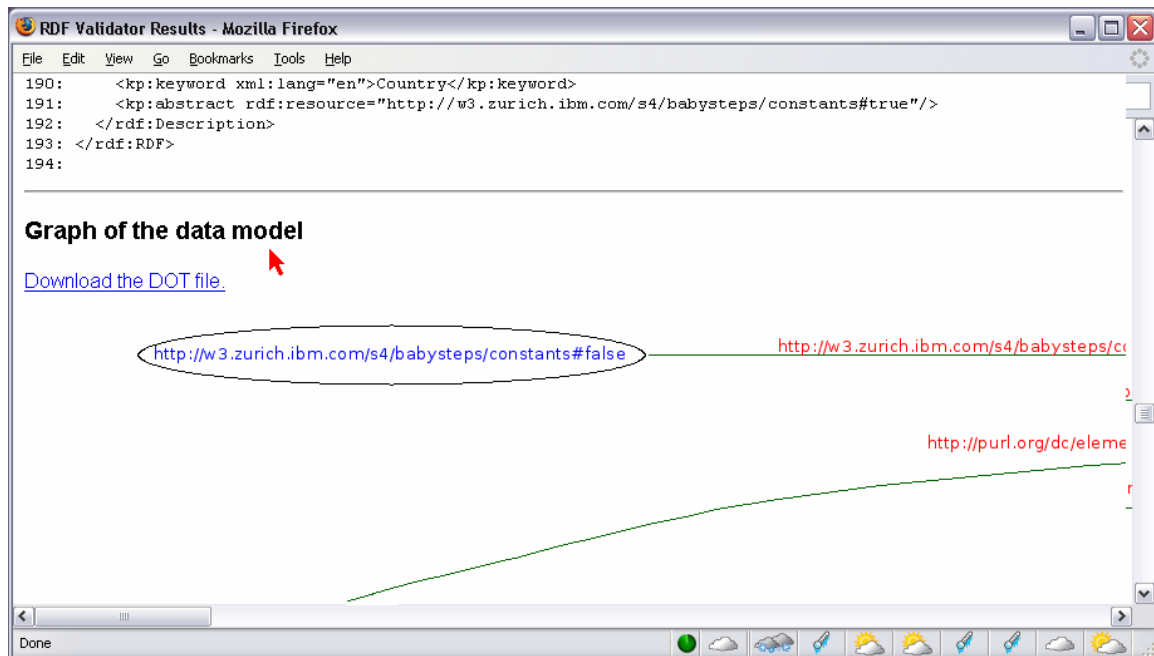
Screenshot 5: adding semantic behavior to a legacy web page

In this first prototype, the implementation relies more on the server-side (servlets and JSP tag library) and less on the client-side (small JavaScript library for rendering popup menus, etc.). In the future, we will evolve this and other prototypes in such a way that more functionality will be shifted to (thin) clients, and that the same dynamic user interface can be achieved without support from particular servers (RDF/A under consideration).

6. W3C Resource Description Framework (RDF) defines model

The demo application is also integrated with the W3C Validation Service (accessible e.g. via “Show model” in Screenshot 5 or Screenshot 7), as shown in Screenshot 6. Invocation of this service, which is deployed remotely from our demo on the W3C web server, confirms that the underlying model, as examined and manipulated by the Jena API, fully conforms to Semantic Web specifications.

Next to triple notation and RDF/XML serialization of our sample model, it also renders a convenient graphical depiction as directed labeled graph. The resulting image is too big to fit on an article page, much less inside a window the size of the other screenshots, and it is best inspected by testing the actual demo live. However, even the small corner that is visible in Screenshot 6 indicates that the demo makes use of the Dublin Core Metadata Initiative (DCMI) element reference set to describe certain resources (titles, creators, etc.).



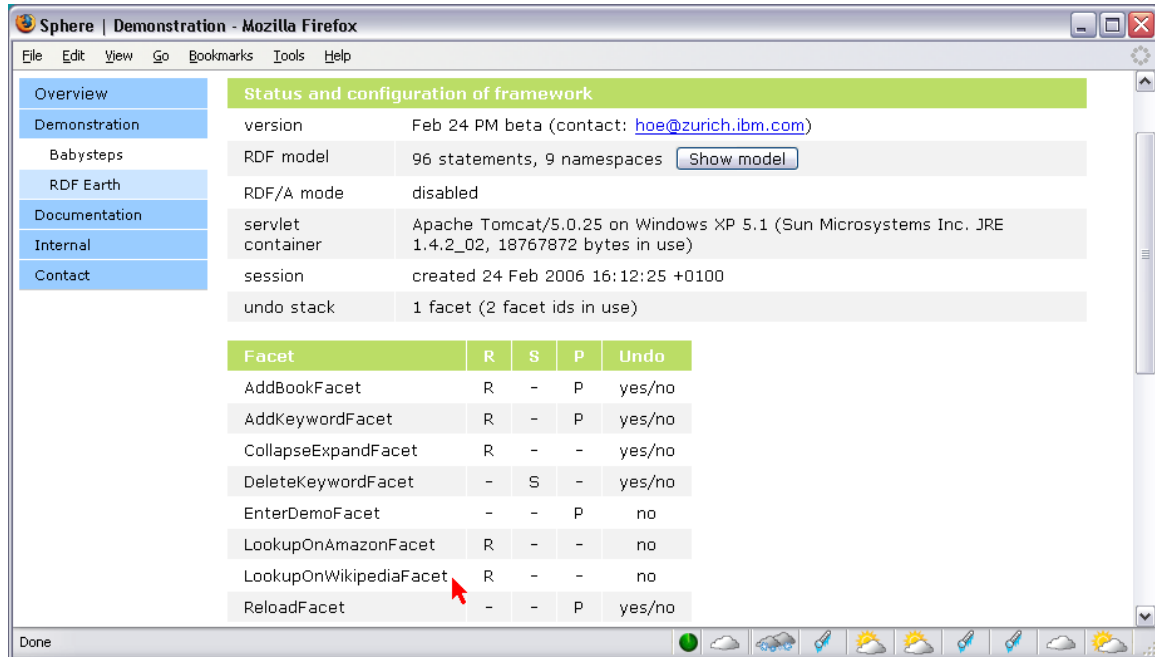
Screenshot 6: integration with W3C RDF validation service

7. The framework manages an open set of facets

Screenshot 7 gives a glimpse of what specific facets contribute to the current demo application. The three columns with binary markers encode trigger conditions that the

listed facets are aware of, and that e.g. concern structural properties of the model graph. Additional trigger conditions are possible.

Further details of the implementation are beyond the scope of this demo write-up and will be documented in another conference submission.



The screenshot shows a web browser window titled "Sphere | Demonstration - Mozilla Firefox". The browser's menu bar includes "File", "Edit", "View", "Go", "Bookmarks", "Tools", and "Help". On the left side, there is a navigation menu with the following items: "Overview", "Demonstration", "Babysteps", "RDF Earth", "Documentation", "Internal", and "Contact". The main content area is titled "Status and configuration of framework" and contains the following information:

- version: Feb 24 PM beta (contact: hoe@zurich.ibm.com)
- RDF model: 96 statements, 9 namespaces (with a "Show model" button)
- RDF/A mode: disabled
- servlet container: Apache Tomcat/5.0.25 on Windows XP 5.1 (Sun Microsystems Inc. JRE 1.4.2_02, 18767872 bytes in use)
- session: created 24 Feb 2006 16:12:25 +0100
- undo stack: 1 facet (2 facet ids in use)

Below this information is a table with the following columns: "Facet", "R", "S", "P", and "Undo". The table contains the following rows:

Facet	R	S	P	Undo
AddBookFacet	R	-	P	yes/no
AddKeywordFacet	R	-	P	yes/no
CollapseExpandFacet	R	-	-	yes/no
DeleteKeywordFacet	-	S	-	yes/no
EnterDemoFacet	-	-	P	no
LookupOnAmazonFacet	R	-	-	no
LookupOnWikipediaFacet	R	-	-	no
ReloadFacet	-	-	P	yes/no

A red arrow points to the "LookupOnWikipediaFacet" row. The browser's status bar at the bottom shows "Done" and a series of weather icons.

Screenshot 7: enumerating facets and their triggers

(end of demo walk through; demo can be shown live or captured in video; references, etc., not included in this demo write-up)