

Creating and Maintaining Metadata Vocabularies for Network-based Photo-sharing

Scenera Research LLC

Cary, North Carolina

www.sceneraresearch.com

Metadata vocabulary development allows mainstream Web users to create photo-sharing metadata vocabularies according to their individual and group needs. It has a form-driven user interface that allows people to create and define metadata vocabularies without requiring them to know or understand the underlying schema languages and syntax. Once a vocabulary is defined, it is stored in the Internet Metadata Library where users can organize it, search for it, share it, and apply it to photo image files and other object types as needed.

Introduction

While the differences in vocabulary between languages such as English, Arabic, and Chinese are clear, each individual language's vocabulary has less obvious differences that are just as important for the accurate communication of written and spoken information. For instance, English has British and American spellings for words that have the same or similar meanings. English also has different meanings for words spelled the same way, different pronunciations for words spelled the same way, and different spellings for words pronounced the same way. Through education and experience we understand and accept the most common inconsistencies; however most of us cannot handle all of the nuances of our vocabulary, much less the nuances of the vocabularies for all of the languages and dialects that exist. Instead, we have ways of cataloging and indexing information to find what we need when we need it. Libraries, dictionaries, thesauri, and encyclopedias all have indexing systems that help us deal with the particulars of language and vocabulary.

Library indexing was originally developed to organize and catalog books and documents written in all languages. Many library indexing vocabularies use hierarchical cataloging schemes that are similar to the yellow pages of a telephone book, but other types of systems are also used, such as the essentially numeric Dewey Decimal and Library of Congress systems for books, and the Machine Readable Cataloging (MARC) system for digital media. These systems have become interoperable with the systems and standards that are used to index and catalog books in many different languages worldwide.

Catalogs and indexes for vocabularies and languages are also used to help locate, track, and manage information on the Internet. One of the fastest growing and most pervasive sources of Internet information, and currently one of the most complex and disjointed, is metadata. Most Web users know metadata as information about image files that they create on their digital cameras and upload to photo image sharing Web sites. In this context, metadata typically describes an image's technical format, its contents, and the processes used to create the image. Embedding metadata in an image is a valuable

way to gather, index, and store detailed properties so the image can be readily cataloged, indexed, and retrieved.

Metadata uses vocabularies that are defined by metadata schemas. The schemas allow the users of a system to agree on the meanings, spellings, rules, and other attributes of the terms used in metadata indexing, much like the systems we use for spoken languages. Open and proprietary vocabulary definition tools do exist, but few widely-used standards are currently in place to allow the range of interoperability that is needed across the Internet today. One such system is defined by the Dublin Core schema, which has fifteen properties that define a vocabulary that can be used to catalog and index metadata openly across different file types and spoken languages. Dublin Core is an excellent baseline and standard, but requires significant interface development for implementation, which can be difficult for the typical Internet user to manage. It also requires broadening and depth to capture all of the properties of data that Internet users really need.

The Metadata Object Description Schema (MODS), developed by the Library of Congress' Network Development and MARC Standards Office, has overcome many of the problems encountered with Dublin Core. It uses nineteen top properties to define a vocabulary, plus 64 lower properties. It also has the flexibility for extension that is required for successful growth with the Internet. MODS can be used for all types of digital objects, and, like Dublin Core, it is not proprietary. Its use in cataloging digital document files, multimedia files, and other objects has grown widely since its 2002 introduction. However, because of its sophistication, it is also difficult and time-consuming for the average Internet user to learn and keep up with as it evolves.

An excellent non-proprietary vocabulary that was designed specifically for digital cameras and photographs is defined by the Exchangeable Image File Format (EXIF), and it is used in digital cameras more than any other standard. However, like the Dublin Core schema, the EXIF vocabulary is limited in the breadth of properties that it defines for metadata. Although it is used by some photo-sharing Web sites, it requires serious training to be expanded by average Web users to meet their ever-growing needs. Photo-sharing sites that don't use EXIF typically use tags to define the content of photographs. Although they are simpler to use, tags are often inaccurate and do not describe the full contents of files.

We developed a vocabulary definition system for photo-sharing that allows inexperienced users to develop custom metadata vocabularies to suit their needs, modify and merge existing vocabularies, associate vocabularies with specific metadata, and share their vocabularies with other users. It is a user-friendly, network-based, interoperable system that can be viewed and operated through standard Web browsers. While it can be used for any type of object associated with metadata, here we have focused on applying it to digital photographic image files for network photo-sharing.

Image Metadata

There are three major categories of metadata associated with images: descriptive, structural, and administrative.

- Descriptive metadata helps in finding images by describing and identifying them so they can be searched, retrieved, and managed. It includes a file's creator, creation date, title, keywords, and other basic identification information.
- Structural metadata describes the internal structure of image files and the relationships between the inner parts, and allows the contents to be searched and presented.
- Administrative metadata helps in securely maintaining information in storage. It assists in the management, tracking, relocation, restructuring, and re-use of files in storage, and typically includes information on creation, quality control, copyrights, history, and preservation.

These categories make up the top layer of a hierarchy that has many properties beneath it, such as subject, shutter speed, aperture, pixels, exposure time, flash, and focal length.

Image file formats that contain metadata properties include EXIF, Joint Photographic Experts Group 2000 (JPEG 2000), Tagged Image File Format (TIFF), Portable Network Graphics (PNG), FlashPix, and Portable Document Format (PDF). Some formats, such as Graphics Interchange Format (GIF) and standard JPEG, offer free-text comment fields, but, like tagging, they lack the accuracy and precision that a more structured metadata vocabulary can provide for photograph indexing, storage, retrieval, and sharing.

Some Web-based photo hosting and photo-sharing sites, such as Flickr, Facebook, and Twango, support the metadata properties associated with photographic image files. However, the metadata vocabularies defined by these properties are restricted. To make up for limitations, some camera manufacturers produce their own metadata properties. Although some of these metadata handling systems are accommodating, they are proprietary and require specialized software to parse, interpret, and store the information.

Some photo-sharing sites support users' specification of custom metadata values through extended properties provided by the site. But what those sites allow is limited, particularly when they don't let users define the new properties or they limit users to a fixed number of custom ones. And when users are given the freedom to define metadata properties, it rarely extends beyond support for adding new keyword-value pairs, as allowed by EXIF. The methods that photo-sharing sites use to store and transmit metadata are also proprietary, restricting their use beyond a particular host site. In fact, searches are often limited to a subset of the limited metadata properties supported by a site.

The technology needed to broaden the range of metadata that is available to users exists today, and it can be configured to allow inexperienced digital image file users to easily develop their own metadata vocabularies.

Metadata Vocabulary Management

The metadata vocabulary management process we designed to create and maintain the vocabularies includes a metadata repository, a database for user and group accounts, and a Web server (Figure 1). Users of Internet photo-sharing sites register themselves with the server to become members of the service where they specify, assign, and manage their metadata and create and manage their metadata

vocabularies. Members can search the server for metadata they collect and vocabularies they define. With permission, they can also use the metadata and vocabularies defined by others.

In our configuration, the metadata vocabularies and the schemas that define the vocabularies and their data properties are stored and indexed by a repository called the Internet metadata library, or IML, which is described in more detail below.

Account information, contact information, and the preferences of each registered user are stored in user and group account databases. User and group administrators specify the metadata vocabularies that are supported for the images associated with their accounts, and users and groups specify the vocabularies they use and share with other users and groups.

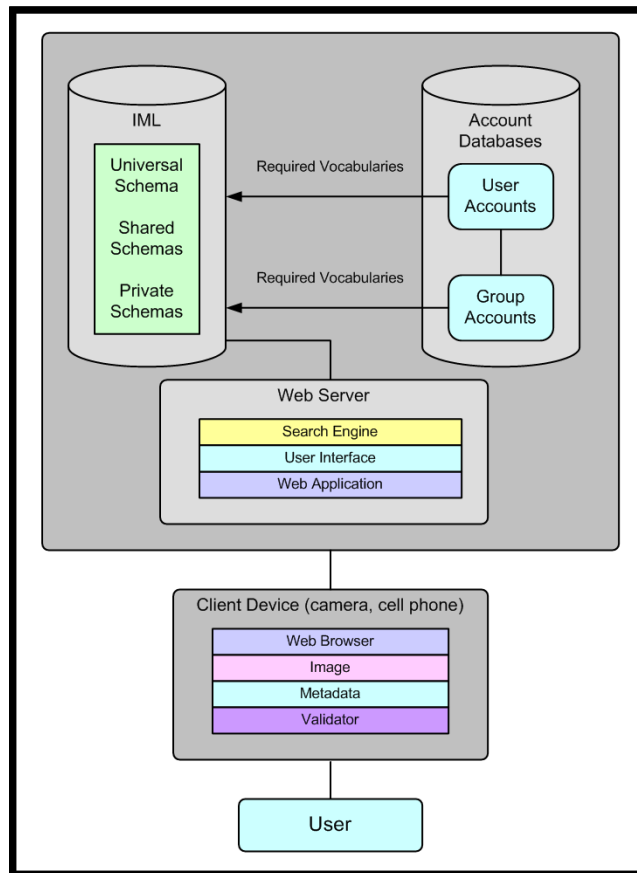


Figure 1. Online metadata vocabulary management, which includes an Internet metadata library, a database for user and group accounts, and a Web server connected to a photo image client device and its user.

The account databases are set up so that certain metadata properties and vocabularies can be required before they are associated with photo image files. For instance, a vocabulary that carries metadata about the owner of the account, such as her name and email address, is often a requirement.

User groups may share common policies, including permission settings, user interface options, required and optional metadata vocabularies, subscription lists, and event/notification policies. For security, roles

and access control lists manage and restrict user access. They are supported by Java Authentication and Activation Services (JAAS).

The Web server displays a form-driven user interface that allows users to easily define vocabularies through their standard Web browsers. The Web server application hides the details of the grammar and syntax of the vocabulary schema specifications and enforces their rules. It also makes vocabularies readily available so users can access them to create new vocabularies, to share them, and to associate metadata with image files.

Metadata Repository

The repository we use for metadata vocabulary management is the Internet metadata library (IML) (1). The IML provides indexed metadata, vocabulary, and schema storage and retrieval according to users' needs. It also integrates metadata into its other services, such as image presentation, searches, indexing, and grouping. Users can access these metadata services directly from any object that is accessible by a Web browser.

The configuration of the IML in its most basic form is shown in Figure 2. It includes a client device with a browser that displays retrieved image files, such as a PC, cell phone, or camera. The metadata associated with the files includes a Service URL that references a Service Profile Application on a server, and the Service Profile Application manages profiles stored in a database. Each profile in the database lists the services available to an image file based on its attributes, such as its identity, content type, provider, or owner, and those listed services are each defined in a corresponding Service Descriptor. Through the IML, Web users can store and retrieve image files, use metadata to search for images on the Internet, and update metadata in image files on the Web. It also stores the components needed to develop metadata vocabularies.

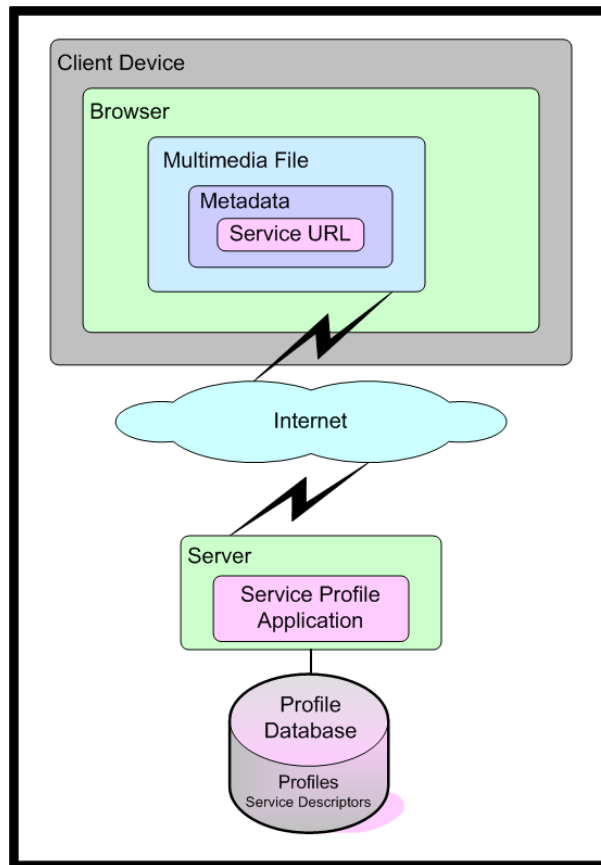


Figure 2. The basic Internet metadata library, which is a repository we use to store and index the vocabularies and schemas used in metadata vocabulary management.

Metadata Vocabulary Schemas

A schema language is used to create a metadata schema that specifies the rules, syntax, grammar, and vocabulary of a metadata language. Each schema specifies the metadata properties and constraints, which together define the vocabulary. Users and groups define their own vocabularies through the site's user interface, which includes the universal vocabulary plus the vocabularies that they create or borrow from others.

There are three types of schemas used in metadata vocabulary management: a universal schema, shared schemas (schemas shared by users), and private schemas (schemas owned by users), and they are defined using Resource Description Framework (RDF). Developed by the World Wide Web Consortium, RDF provides a schema language that specifies the schemas that are used to define metadata vocabularies, and it allows interoperability between applications that exchange metadata. RDF also allows descriptions of Web objects, defined as any object with a uniform resource identifier (URI) as its address, to be made available in machine-readable form. RDF is an application of extensible markup language (XML) and is specific to file descriptions. It supports a class system, where a class specifies a set of properties and the constraints applied to the values of those properties. A collection of values associated with the properties has a property type and value that is specified in the associated

schema, and the schemas are uniquely identified by assigning each a URI. RDF utilizes the namespace facility of XML to point to the URI; thus, the schema is accessed at the URI that is identified by the namespace.

Recently, some camera manufacturers and photo-sharing sites have begun to store the metadata they support in XML and RDF format, which goes a long way toward making the information useful to a wider range of applications. But although RDF enables metadata to be understood by many applications and enables it to be infinitely extendable, its syntax can be complicated to learn for those who aren't well versed in using computer markup languages. For instance, this is a portion of the RDF syntax that describes a report:

```
<Description about = http://sceneralabs.com/report.html>
  <DC:Title>Specifying and Assigning Metadata
  </DC:Title>
  <DC:Creator>Paul Morris<DC:Creator>
  <DC>Date>2008-09-01</DC>Date>
  <DC:Subject>Metadata, RDF</DC:Subject>
</Description>
</RDF>
```

People who are not comfortable with XML or with defining the RDF syntax are likely to find it difficult to create metadata vocabularies to suit their particular needs. Our system makes metadata vocabulary development much easier for them.

Metadata Vocabulary Development

Our user-friendly metadata vocabulary creation process is illustrated in Figure 3. Once a user logs on the metadata vocabulary management server, the Web server application displays Web pages where users can create vocabularies, add references to vocabularies from other places on the Web, and add references to external metadata vocabularies.

Users can create new vocabularies from scratch or develop them from existing ones by editing or merging vocabularies. To assist in vocabulary development, the Internet metadata library can provide an online repository of vocabularies that can be used as starting points for creating them.

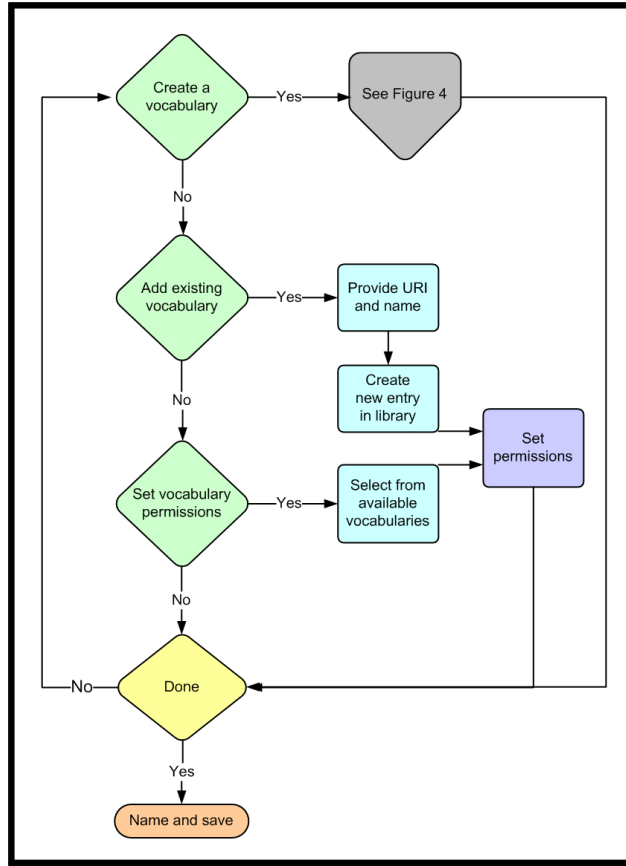


Figure 3. The process of creating metadata vocabularies from existing ones using the metadata vocabulary developer.

Adding Vocabulary Properties

When a user chooses the option to add a property to a metadata vocabulary, the Web application gives two options: to create a new property or to reuse properties from an existing vocabulary (Figure 4). New properties are created for the vocabulary by entering the property name and specifying the constraints for the possible values it holds. For each new property that a user creates, the Web server application prompts the user for the name of the property and the user selects from a list of property types, such as String, List, Boolean, or Numeric. Depending on the type, the Web server application then prompts the person to provide additional constraints, such as the list of possible values for a list type or the upper and lower limits for a numeric range.

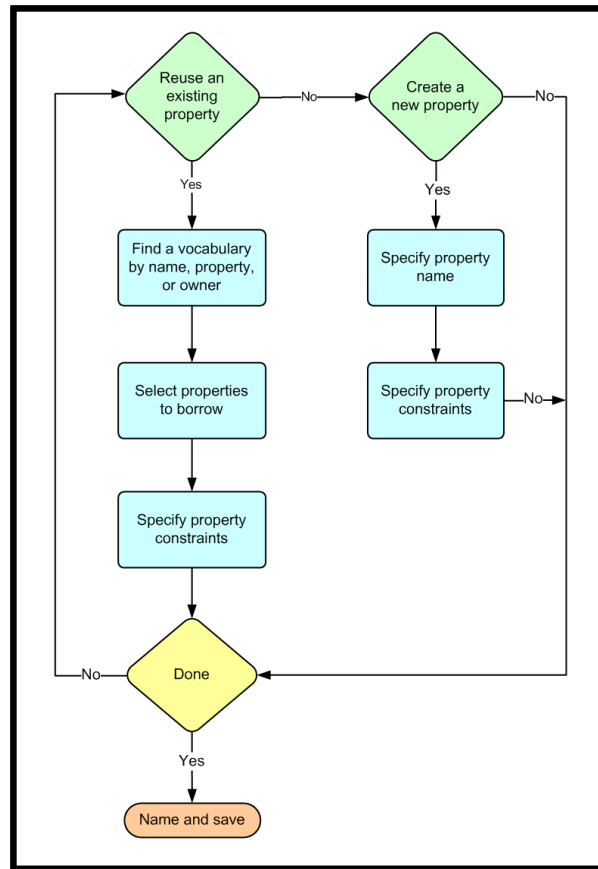


Figure 4. The process of creating metadata vocabulary properties or borrowing them from existing vocabularies using the Internet metadata library.

The Web server application builds the user interface, generating forms where the user can enter values for the assigned metadata properties. Extensible Stylesheet Language Transformation (XSLT) transforms the RDF schema into XHTML forms to display in the user's Web browser. The user then enters data values in the forms as needed for the vocabulary properties. As fields are entered for each property, the Web server application uses its defined constraints to determine the correct form element to use for data input, such as a text field, selection list, or choice list.

The Web server application validates the entered values based on the chosen vocabulary constraints. When the data are proved valid and all of a property's values are entered, the Web server application associates the metadata with the object and saves it. The process of creating new properties is repeated as many times as needed. When all entries are complete, the user names and saves the new vocabulary in the library.

Creating New Metadata Vocabularies

When a user opts to create a new metadata vocabulary, the form-driven user interface generated by the Web server application gives the option of creating vocabularies from scratch or creating them from existing ones by editing, expanding, and merging multiple vocabularies.

Creating Vocabularies from Scratch

Creating new vocabularies from scratch is very similar to creating new properties. First, a user chooses to create it, and then she creates the first property by naming it, adding its values, and specifying its constraints. The user goes through that cycle until all needed properties are defined and the vocabulary is complete. The user then names and saves the new vocabulary in the library, as shown in Figure 4.

Reusing Vocabularies

When creating a new vocabulary by reusing existing ones, a search facility opens and allows the user to locate the existing vocabulary based on its URI, vocabulary name, names of properties contained in it, its owner, or the intended target type, such as photo image. Those vocabularies may be found in the library or on the Web. This search tool makes vocabularies easy to browse and will continue to include additional metadata libraries as more of them become available on the Internet.

Once a user selects an existing vocabulary, she enters new property values using the same procedures described for creating new properties until the new vocabulary is complete, named, and saved as shown in Figure 4.

Reusing Vocabularies: A User Example

Sam wants to create a metadata vocabulary that describes the photograph images of his three dogs. He decides to begin by borrowing properties from the Dublin Core, the general metadata vocabulary. He logs on the Internet metadata library and uses the vocabulary search tool to locate the Dublin Core by entering its URI.

When the Dublin Core is located and Sam selects it, the Web server application displays all of its properties. Sam selects the properties he wants to borrow, including the Subject, Title, and Description, and then he searches for another vocabulary. This time he searches for a library using the property “Breed,” and a library with several dog-specific properties opens. From it, he selects the Breed, Vaccination, and Training properties.

Sam repeats this borrowing process from other vocabularies, and as the properties are borrowed the Web server application lets him specify additional constraints. For example, the “Subject” property in the Dublin Core is a string of keywords, and he selects to restrict the Subject string to the names of his three dogs, plus a general name for photographs that hold more than one of them: Spot, Rover, Sam, Jr., and My Mutts. His new vocabulary is designed specifically for the photographs of his dogs and the Web application displays Sam’s personal list of pet names to choose from. From now on, his well-defined properties make it very easy for him to find a particular photograph or specific information about his dogs.

Merging Vocabularies

Merging vocabularies is similar to creating and reusing them. When a user uploads photograph image files to be shared, the Web application checks the user account and relevant group accounts, and retrieves vocabularies that are specific to photo images. The user selects the vocabularies that hold the properties she wants, and then selects the option to merge them.

The merging process performs two primary functions. First, it removes duplicate properties that result from borrowing them from multiple vocabularies. Second, because the same property may have different sets of constraints specified by the different vocabularies, it ensures that the predefined policies for conflicting constraints are enforced. The process of merging metadata vocabularies and assigning the final vocabulary to a photo image is illustrated in Figure 5.

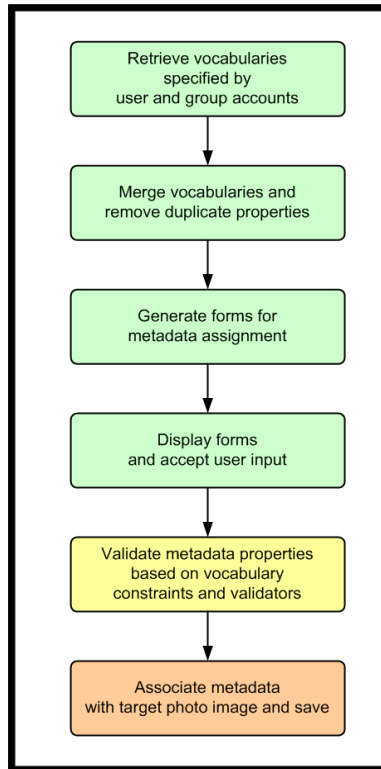


Figure 5. The process of merging metadata vocabularies and assigning the final vocabulary to a digital image.

Merged metadata vocabularies usually borrow properties from each other, so the Web server application merges the required metadata before it prompts users to enter new values for properties. Once the vocabularies are merged, users can enter new properties and adjust the constraints using the same procedures described above for creating new vocabularies.

Applying Constraints

Document Content Description (DCD) is the vocabulary that RDF uses to specify structural and content constraints. It is designed specifically for XML documents and covers the names of properties, the contents of data, and the values of properties (2). The constraints it defines are limited and, although properties can be added, customization requires in-depth knowledge of XML. Our system gives users the option of easily associating a vocabulary validator before a newly-defined metadata vocabulary is stored, as shown in Figure 6.

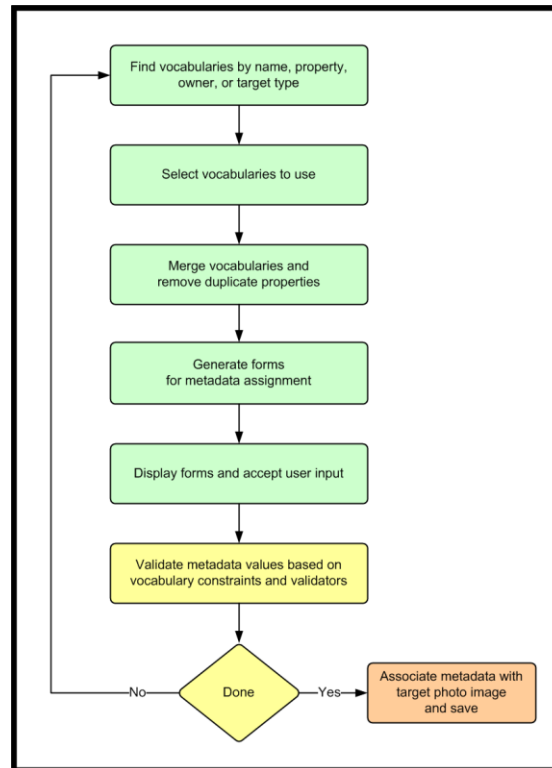


Figure 6. The process of allowing the user to associate optional metadata with an image that already exists in the photo-sharing site.

The vocabulary validator is a software program that supports an API provided by the system. It runs on either the server or the client computers and uses a plug-in interface provided by the Web application. This plug-in defines an interface that is compatible either with Java’s Enterprise JavaBeans (EJB) technology or with a remote method invocation technology, such as Remote Method Invocation (RMI) or Internet Inter-ORB Protocol (IIOP), which makes the location of the validator unknown to the Web server application that invokes it. After it ensures that the constraints specified by the RDF schema are enforced, the Web server application passes the metadata to the vocabulary validator for further constraint enforcement.

Constraint conflicts are also resolved. For example, if two required vocabularies both borrow the property “Subject” from the Dublin Core and both have additional constraints on the values of “Subject,” a conflict occurs. When it is in its default setting, the system ensures that the most restrictive constraints are applied. Users can select different constraint values through manual input or by opting to always choose the constraints of one vocabulary over another.

System Application and Expansion

With its underlying application of Dublin Core and RDF standards, metadata vocabulary management can be used across different networks and Web sites, depending on defined permissions and assignments. This means that virtually any Internet or intranet system, such as the World Wide Web and networks for corporations, government agencies, schools, non-profits, or other organizations, can be

used to build photo-sharing sites that can be accessed and used by users with all levels of Web experience. The information captured from the image files can be tailored to meet the specific needs of corporations, regulatory bodies, individuals, and peripherals as desired, and any of that information can be shared or secured.

The vocabulary generator can be made universally accessible across the World Wide Web using a service access tool, such as the Service Access Framework (3). It can also be broadened for use with virtually any metadata specification language. Further work in this area will include using multiple vocabulary generators that lead to one repository. That one repository, such as the Internet metadata library, will use schemas from each of the multiple generators.

Still images aren't the only types of objects that can be used with this system; any object that can have metadata associated with it can be used. That includes multiple image files, such as time-lapse images, burst images, and panoramic images, plus non-image files, such as audio files, movies, and text documents. When used with a tool like the Internet metadata library, metadata vocabulary customization can be used to store, organize, search for, retrieve, share, and secure nearly any type of object sought and handled by network users and administrators.

With expansion and application of metadata standards across networks, the use of metadata vocabularies and libraries will continue to broaden. What is needed is the on-going acceptance of existing object and metadata standards, and maintaining a flexible system that allows everyday users to feel comfortable handling them. All of the necessary tools exist today. All they require is implementation.

Acknowledgments

The author would like to thank Julie Tomlinson for her writing and editing.

References

1. "[An Internet Metadata Library](#)." Scenera Research LLC, 2008. 14pp.
2. "[Document Content Description for XML](#)." World Wide Web Consortium, 1998.
3. "[The Service Access Framework](#)." Scenera Research LLC, 2008. 11pp.

Contact

Contact Scenera Research at info@sceneraresearch.com.