



Escenic Content Engine
Integration Guide
5.2.7.2







Copyright © 2010-2011 Vizrt. All rights reserved.

No part of this software, documentation or publication may be reproduced, transcribed, stored in a retrieval system, translated into any language, computer language, or transmitted in any form or by any means, electronically, mechanically, magnetically, optically, chemically, photocopied, manually, or otherwise, without prior written permission from Vizrt.

Vizrt specifically retains title to all Vizrt software. This software is supplied under a license agreement and may only be installed, used or copied in accordance to that agreement.

Disclaimer

Vizrt provides this publication “as is” without warranty of any kind, either expressed or implied.

This publication may contain technical inaccuracies or typographical errors. While every precaution has been taken in the preparation of this document to ensure that it contains accurate and up-to-date information, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained in this document.

Vizrt’s policy is one of continual development, so the content of this document is periodically subject to be modified without notice. These changes will be incorporated in new editions of the publication. Vizrt may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

Vizrt may have patents or pending patent applications covering subject matters in this document. The furnishing of this document does not give you any license to these patents.

Technical Support

For technical support and the latest news of upgrades, documentation, and related products, visit the Vizrt web site at www.vizrt.com.

Last Updated

21.12.2011





Table of Contents

1 Introduction	7
1.1 Conventions Used in This Manual	8
1.2 How The Web Service Works	8
1.3 Change Logs	10
1.3.1 Accessing Change Logs From the "Start" Feed	11
1.3.2 Accessing a Section's Change Log	12
1.4 Searching	12
2 Supported MIME Types	15
2.1 application/atom+xml	15
2.2 application/vnd.vizrt.payload+xml	16
2.3 application/vnd.escenic.content+xml	16
2.4 application/opensearchdescription+xml	16
3 Supported Link Relations	17
3.1 Standard Link Relations	17
3.1.1 down	17
3.1.2 edit	17
3.1.3 next	18
3.1.4 previous	18
3.1.5 related	19
3.1.6 search	19
3.2 Proprietary Link Relations	20
3.2.1 http://www.vizrt.com/types/relation/changelog	20
3.2.2 http://www.vizrt.com/types/relation/content-items	21
3.2.3 http://www.vizrt.com/types/relation/home-section	21
3.2.4 http://www.vizrt.com/types/relation/inboxes	21
3.2.5 http://www.vizrt.com/types/relation/lists	22
3.2.6 http://www.vizrt.com/types/relation/list-pool	22
3.2.7 http://www.vizrt.com/types/relation/pages	22
3.2.8 http://www.vizrt.com/types/relation/protection-domains	22
3.2.9 http://www.vizrt.com/types/relation/publication	23
3.2.10 http://www.vizrt.com/types/relation/section	23
3.2.11 http://www.vizrt.com/types/relation/lock	23
3.2.12 http://www.vizrt.com/types/relation/preview	23
4 Standard Element Usage	25



4.1 Content Item State Representations	26
5 Viz Data Format	29
5.1 payload	29
5.2 field	30
5.3 value	30
5.4 list	30
6 How to...	31
6.1 Get Started	31
6.2 Navigate The Section Hierarchy	32
6.3 Search For Content	33
6.4 Change a Content Item	34
6.5 Preview a Content Item	36
6.6 Follow Content Item Relations	37
6.7 Lock a Content Item	39
6.7.1 Get the Resource to Lock	40
6.7.2 Get Public Locks	40
6.7.3 Post Lock Request	41
6.7.4 Put Updated Resource	42
6.7.5 Delete Private Lock	44
7 Advanced Features	45
7.1 Using Change Logs Effectively	45
7.2 Optimistic Concurrency	46

1 Introduction

This manual describes the Escenic Content Engine's open web service, and provides basic advice on how you can use it to integrate the Content Engine with other applications and web services. Client applications can communicate with the web service by means of a REST API. This manual assumes that you are already familiar with the characteristics of REST (Representational State Transfer) APIs, and the principles on which they are based. If you are not, you can find out more about REST here:

- http://en.wikipedia.org/wiki/Representational_State_Transfer

The Content Engine's web service uses the ATOM publishing protocol, and has the following characteristics:

- All operations are performed using HTTP **GET**, **PUT**, **POST** and **DELETE** commands.
- All resources provided by the web service are identified by URIs.
- Given the initial URI of the web service, all the resources it can provide are then discoverable via hypertext links.
- Resources returned by the web service fall into three categories (as defined by the ATOM publishing protocol):
 - **Feed resources** (that is, ATOM feeds containing a list of resource links)
 - **Entry resources** (single ATOM entry elements containing, for example, content items)
 - **Media resources** (binary resources such as images or video clips)
- Wherever possible, Standard ATOM constructs are used in the entry resources returned by the web service. However, some extensions are used. These fall into two categories:
 - Some RDF (Dublin Core) constructs are used where suitable standard ATOM constructs are not available. See [chapter 4](#) for details.
 - A small number of proprietary extensions are used where necessary.

Here is a complete list of the ATOM standards on which the web service is based:

- RFC 4287 — Atom Syndication Format
- RFC 5023 — Atom Publishing Protocol
- RFC 4685 — Atom Threading Extensions
- RFC 5005 — Feed Paging and Archiving
- OpenSearch — OpenSearch search descriptions are used to describe how to search in an Atom collection.
- draft-snell-atompub-tombstones-06 (moving forward to RFC soon) — Atom Tombstones
- draft-divilly-atom-hierarchy-03 — Atom Hierarchy Relations

Since the web service API is largely based on these open, documented standards, the reference information in this guide is limited to the proprietary extensions needed to provide a comprehensive service.

1.1 Conventions Used in This Manual

The Content Engine's REST API is completely language-independent: it is based entirely on the use of URIs, and all resources returned by the web service either consist of XML documents or are binary objects described in accompanying XML documents. You can therefore write client applications using more or less any programming language you choose (although some languages provide better tools for communicating via the HTTP protocol and handling XML documents than others).

In order not to require knowledge of a particular programming language all of the examples shown in this manual are based on the use of `curl`. `curl` (see <http://curl.haxx.se/>) is a command line utility for transferring data using URI syntax, and supports all the HTTP operations needed to communicate with the Content Engine web service. It is a freely-available open source application and is available for all major computing platforms.

1.2 How The Web Service Works

This section contains a brief general description of how a client application is expected to make use of the Content Engine's web service.

Assuming the client has no a priori knowledge of a particular web service or what it can provide, it can initiate contact by sending an HTTP `GET` request to the web service's "start" URI. This is usually

```
http://host-ip-address/webservice/escenic/section/ROOT/subsections
```

where *host-ip-address* is the host name or IP address of the Content Engine. This "start address" (plus a user name and password) is the minimum information a client needs to be able to interact at some level with the web service. The following `curl` command, therefore, will initiate contact with a Content Engine web service:

```
curl -u user:password -X GET http://host-ip-address/webservice/escenic/section/ROOT/subsections
```

The web service responds to such a request by returning a **feed resource**: an ATOM feed document that looks something like this:

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/section/ROOT/subsections</id>
  <link rel="self"
        href="http://host-ip-address/webservice/escenic/section/ROOT/subsections"
        type="application/atom+xml"/>
  <updated>2010-06-23T16:51:06.721Z</updated>
  <title type="text">Publication root sections you are authorized to.</title>
  <link rel="changelog"
        href="http://host-ip-address/webservice/escenic/changelog"
        type="application/atom+xml"/>
</feed>
```



```

<entry>
  <id>http://host-ip-address/websevice/escenic/section/1</id>
  <title type="text">frontpage</title>
  <updated>2010-06-22T10:16:46.309Z</updated>
  <category term="directory" scheme="http://www.vizrt.com/types"/>
  <link rel="edit-media"
    type="application/vnd.escenic.content+xml;type=com.escenic.section"
    href="http://host-ip-address/websevice/escenic/section/1"/>
  <link rel="down"
    href="http://host-ip-address/websevice/escenic/section/1/subsections"
    type="application/atom+xml"/>
  <link rel="http://www.vizrt.com/types/relation/inboxes"
    href="http://host-ip-address/websevice/escenic/section/1/inboxes"
    type="application/atom+xml"/>
  <link rel="http://www.vizrt.com/types/relation/lists"
    href="http://host-ip-address/websevice/escenic/section/1/lists"
    type="application/atom+xml"/>
  <link rel="http://www.vizrt.com/types/relation/pages"
    href="http://host-ip-address/websevice/escenic/section/1/pages"
    type="application/atom+xml"/>
  <link rel="http://www.vizrt.com/types/relation/content-items"
    href="http://host-ip-address/websevice/escenic/section/1/content-items"
    type="application/atom+xml"/>
  <link rel="http://www.vizrt.com/types/relation/changelog"
    href="http://host-ip-address/websevice/escenic/changelog/section/1"
    type="application/atom+xml"/>
</entry>
</feed>

```

The document contains a series of **entry** elements representing the root sections of all the Content Engine's publications. Each entry contains a sequence of link elements representing the resources to which the requesting user has access. In the example, above, for example, the requesting user has access rights to many resources in the publication "Escenic Times", but not to "frontpage Test".

A **link** element has three attributes:

href

The URI of a resource.

type

This attribute contains a MIME type id identifying the type of data the resource contains. The MIME type **application/atom+xml**, for example, indicates that the resource is another ATOM feed document. All the types used by the Content Engine web service are described in [chapter 2](#).

rel

This attribute describes the meaning of the resource in terms of its relationship to the current resource. All the link relations used by the Content Engine web service are described in [chapter 3](#).

These three attributes provide the client application with sufficient information to determine:

- Which resource (if any) it needs to access next
- The URL needed to access the resource
- What kind of data it can expect to find in the resource.

The client can construct a new **GET** request from the information returned in this document and obtain a new feed resource. It can, for example obtain

a feed containing links to all the top-level sections in the Escenic Times publication by submitting the following request:

```
curl -u user:password -X GET http://host-ip-address/webservice/escenic/section/22/subsections
```

By recursively retrieving feed resources in this way, it is possible to navigate down through the section tree, then from a section to a list of in-boxes, for example, from there to a specific in-box. At various points along this route you can also retrieve feed resources that list links to actual content items (The content items in a section, for example, or the content items in a list or in-box). By following these links you can retrieve entry resources that represent actual content items, and media resources such as image files.

The web service allows entry and media resources to be modified by users with sufficient privileges. This is done by submitting requests using HTTP **PUT**, **POST** and **DELETE** operations:

- The client can update an entry or media resource by submitting a **PUT** request containing a modified resource and the URI of the resource to be updated.
- The client can create a new entry or media resource by submitting a **POST** request containing a new resource and the URI of the Atom collection (i.e. feed resource) to which it is to be added.
- The client can delete an entry or media resource by submitting a **DELETE** request containing the URI of the resource to be deleted.

1.3 Change Logs

Most of the resources returned by the web service in response to HTTP **GET** requests correspond to standard Content Engine structures that are well-known and described elsewhere: sections, content items, section pages, lists and inboxes. **Change logs**, however, are new structures that requires some explanation.

A change log is a record of all the changes made to a set of content items. The entries in a change log are ordered by change date and are in reverse chronological order (that is, the most recent change is listed first). All change logs are maintained in real time: every time a content item is created, deleted or modified, the change is inserted at the front of all appropriate change logs.

Change logs are maintained for **all** the content items managed by the Content Engine. Special change logs are also maintained for all of a Content Engine's person records (that is, content items containing information about all registered users of the Content Engine - editors, authors, publication administrators and registered site users).

The Content Engine maintains a number of different change logs, in order to ensure that users can only access change log entries for content items to which they have read access. Specifically, the Content Engine maintains one content item change log for each of its **protection domains**. A protection domain is a set of sections governed by an **explicitly defined** set of access permissions. What this means in practice is that there is always at least



one protection domain for each of a Content Engine's publications, since there will always be an explicit set of access permissions associated with each publication's root section. A publication will have additional protection domains for each subsection that does not simply inherit access controls from its parent section, but has explicitly defined access controls of its own.

In addition to a change log for each protection domain, the Content Engine maintains one person record change log for each publication.

A user only has access to a change log if he/she has access to the relevant protection domain.

A change log is returned as an Atom **paged collection**. The first entry in the first Atom feed returned represents the most recently changed content item in the protection domain represented by the change log. The next entry represents the second most recent change, and so on. You can get the next page, a feed containing an older set of changes by following this feed's **next** link.

Change logs can be accessed in two ways:

- Via a link in the web service's "start" feed resource
- Via a link in a section's "content items" feed resource

1.3.1 Accessing Change Logs From the "Start" Feed

The highlighted link in the following "start feed":

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/section/ROOT/subsections</id>
  <link rel="self"
    href="http://host-ip-address/webservice/escenic/section/ROOT/subsections"
    type="application/atom+xml"/>
  <updated>2010-05-31T12:45:38.144Z</updated>
  <title type="text">Publication root sections you are authorized to.</title>
  <link rel="http://www.vizrt.com/types/relation/protection-domains"
    href="http://host-ip-address/webservice/escenic/changelog"
    type="application/atom+xml"/>
  ...
</feed>
```

returns a feed resource containing a list of all the change logs the requesting user is allowed to access:

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <id>http://host-ip-address/webservice/escenic/changelog</id>
  <title type="text">Protection domains accessible to demo_admin</title>
  <updated>2010-06-23T17:54:09.173Z</updated>
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <link rel="self"
    href="http://host-ip-address/webservice/escenic/changelog"
    type="application/atom+xml"/>
  <entry>
    <id>http://host-ip-address/webservice/escenic/publication/demo</id>
    <title type="text">demo</title>
    <link rel="http://www.vizrt.com/types/relation/changelog"
      href="http://host-ip-address/webservice/escenic/changelog/publication/1"
      type="application/atom+xml"/>
  </entry>
```

```

<entry>
  <id>http://host-ip-address/webservice/escenic/section/3</id>
  <title type="text">topicRoot</title>
  <link rel="http://www.vizrt.com/types/relation/changelog"
    href="http://host-ip-address/webservice/escenic/changelog/section/3"
    type="application/atom+xml"/>
</entry>
<entry>
  <id>http://host-ip-address/webservice/escenic/section/4</id>
  <title type="text">New Articles</title>
  <link rel="http://www.vizrt.com/types/relation/changelog"
    href="http://host-ip-address/webservice/escenic/changelog/section/4"
    type="application/atom+xml"/>
</entry>
...
</feed>

```

One of the change log links in this feed (the one with "publication" in its `id` attribute) is the person record change log. All the others correspond to protection domains.

1.3.2 Accessing a Section's Change Log

You can access the change log that contains all the changes in a specific section from the section's feed resource:

```

<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/section/1/subsections</id>
  <link rel="self"
    href="http://host-ip-address/webservice/escenic/section/1/subsections"
    type="application/atom+xml"/>
  <updated>2010-06-23T17:35:45.180Z</updated>
  <title type="text">Subsections for Section with id=1</title>
  <entry>
    <id>http://host-ip-address/webservice/escenic/section/4</id>
    <title type="text">New Articles</title>
    <updated>2010-06-22T10:16:46.637Z</updated>
    <category term="directory" scheme="http://www.vizrt.com/types"/>
    ...
    <link rel="http://www.vizrt.com/types/relation/changelog"
      href="http://host-ip-address/webservice/escenic/changelog/section/4"
      type="application/atom+xml"/>
  </entry>
</feed>

```

Note that the change log accessed in this way is guaranteed to contain all the changes made to the content items in the section, but it may also contain changes made in other sections, since it is the change log for the protection domain to which the section belongs.

1.4 Searching

The Content Engine web service includes a search service based on **OpenSearch**. OpenSearch is a standard for:

- Describing how to search in Atom collections
- Annotating search results returned in Atom feeds



For details see <http://www.opensearch.org/Home>. Using OpenSearch helps to ensure that the search service complies with REST principles, and can be searched by clients with no prior knowledge of the Content Engine or its internal structures.

Every <http://www.vizrt.com/types/relation/subsection> entry feed returned by the web service contains an <http://www.vizrt.com/types/relation/content-items> link:

```
...
  <link rel="http://www.vizrt.com/types/relation/content-items"
        href="http://host-ip-address/websevice/escenic/section/22/content-items"
        type="application/atom+xml"/>
...
</feed>
```

This link returns an empty feed resource representing the collection of all content items in the section. It also, however, contains a **search** link:

```
<link rel="search"
      href="http://host-ip-address/websevice/open-search/escenic/22/content-search-description.xml"
      type="application/opensearchdescription+xml"/>
```

This link returns an OpenSearch document containing URI templates that can be used to search for content items in the section:

```
<OpenSearchDescription xmlns="http://a9.com/-/spec/opensearch/1.1/">
  <ShortDescription>Escenic search</ShortDescription>
  <Description/>
  <Url type="application/atom+xml"
        template="http://host-ip-address/websevice/publication/dev.nightly/search/escenic/22/
{searchTerms}/?pw={startPage?}&amp;c={count?}" />
  <Url type="text/html"
        template="http://host-ip-address/websevice/publication/dev.nightly/search/escenic/22/
{searchTerms}/?pw={startPage?}&amp;c={count?}&amp;format=html" />
  <Contact>http://www.escenic.com</Contact>
  <Tags/>
  <LongName>Escenic Content Engine Search</LongName>
  <Image height="16" width="16" type="image/x-icon">http://host-ip-address/websevice/images/
ece.ico</Image>
  <Query role="example" searchTerms="cat" />
  <Developer>Escenic AS</Developer>
  <Attribution/>
  <SyndicationRight>private</SyndicationRight>
  <AdultContent>>false</AdultContent>
  <OutputEncoding>UTF-8</OutputEncoding>
  <InputEncoding>UTF-8</InputEncoding>
</OpenSearchDescription>
```

One of the templates returns results in a paged Atom feed, the other returns the results as HTML. The client can use one of them to construct and submit a search request:

```
curl -u stale:stale -X GET http://host-ip-address/websevice/publication/pub-name/search/escenic/22/
Obama/?pw=1&c=10
```

If the Atom feed template (`type="application/atom+xml"`) is used, then a paged Atom feed is returned in which the actual results are wrapped in an **OpenSearch** element and presented in elements belonging to a proprietary, **undocumented and deprecated** Vizrt format identified by the namespace <http://www.escenic.com/2007/content-engine>. The use of this format is temporary and you should not invest any effort in attempting to make use of it: it will be replaced with a properly documented and supported format in a future version of the Content Engine.





2 Supported MIME Types

Every `link` element in the feed resources returned by the Content Engine web service has a `type` attribute. This attribute contains a MIME type identifier that identifies the type of data in the resource referenced by the link. This chapter lists all the `type` URIs that client applications can expect to find in feed resources returned by the Content Engine web service, and descriptions of their meaning.

The MIME type URIs found in Content Engine feed resources fall into two categories:

- Standard MIME type identifiers defined in the IANA MIME type registry (<http://www.iana.org/assignments/media-types/>). The descriptions of these relations consist of a reference to the appropriate specification and a brief description of how it is used by the Content Engine web service.
- Proprietary MIME types defined by Vizrt.

Note that the MIME type in a link is only intended to provide client applications with a hint about what kind of data to expect. It is not a contract and cannot be relied upon to always be correct. The resource referenced by the link may, for example, no longer exist. The MIME type's primary purpose is to provide information that the client application can use when selecting the link to follow. Client applications should therefore always use the HTTP header of the resource returned by the link to determine what kind of resource it actually is and how it should be processed.

2.1 **application/atom+xml**

This standard MIME type indicates that the referenced resource is an Atom feed, collection or entry.

The Content Engine web service uses this MIME type for:

- Subsection feeds
- Content items feeds
- Search result feeds
- Change log feeds
- Section lists feeds
- Section in-boxes feeds
- Section pages feeds
- Content item entries

2.2 **application/vnd.vizrt.payload+xml**

This proprietary MIME type indicates that the referenced resource is a Viz Data Format (VDF) document. For a description of VDF, see [chapter 5](#).

2.3 **application/vnd.escenic.content+xml**

This proprietary MIME type indicates that the referenced resource has content in an undocumented, deprecated format identified by the namespace `http://www.escenic.com/2007/content-engine`.

The Content Engine uses this MIME type for various kinds of user content. The MIME type id will usually be followed by one of the following type specifications in order to indicate precisely what kind of user content is to be expected:

com.escenic.section

The VDF document contains an Escenic section definition.

com.escenic.inbox

The VDF document contains an Escenic inbox definition.

com.escenic.list

The VDF document contains an Escenic inbox definition.

com.escenic.page

The VDF document contains an Escenic inbox definition.

The use of this format is temporary and you should not invest any effort in attempting to make use of it: it will be replaced with a properly documented and supported format in a future version of the Content Engine.

2.4 **application/opensearchdescription+xml**

This standard MIME type indicates that the referenced resource is an OpenSearch search description document.

The Content Engine uses this MIME type to describe the URL format that must be used to search for content in a section. For more information about searching, see [section 1.4](#); for an example of how to search, see [section 6.3](#).

3 Supported Link Relations

Every `link` element in the feed resources returned by the Content Engine web service has a `rel` attribute. This attribute contains a **link relation**, a name that identifies the relationship between the current resource and the resource referenced by the link. This section lists all the link relations that client applications can expect to find in feed resources returned by the Content Engine web service, and descriptions of their meaning.

The link relations found in Content Engine feed resources fall into two categories:

- Standard link relations defined in one of the ATOM specifications. The descriptions of these relations consist of a reference to the appropriate specification and a brief account of how they are used by the Content Engine web service.
- Proprietary relations defined by Vizrt. The descriptions of these relations consist of a brief account of their meaning and use.

3.1 Standard Link Relations

3.1.1 **down**

This link relation is defined in <http://tools.ietf.org/html/draft-divilly-atom-hierarchy-03>.

It is used by the Content Engine web service to represent hierarchical section relationships. A section's subsection's, for example, are represented by **down** links. You can navigate down through a publication's section hierarchy by simply following **down** links.

The following example shows a link of this type:

```
<entry>
  ...
  <link rel="down"
        href="http://host-ip-address/webservice/escenic/section/3461/subsections"
        type="application/atom+xml"/>
  ...
</entry>
```

3.1.2 **edit**

This link relation is defined in [RFC 5023](#).

It is used by the Content Engine web service to represent an entry containing an actual content item that can be edited by the client application and re-submitted.

The following example shows a link of this type in a change log entry:

```
<entry xmlns:app="http://www.w3.org/2007/app"
```

```

xmlns:metadata="http://xmlns.escenic.com/2010/atom-metadata"
xmlns:dcterms="http://purl.org/dc/terms/">
...
<link href="http://host-ip-address/webservice/escenic/content/4" rel="edit"/>
...
</entry>

```

3.1.3 next

This link relation is defined in [IETF RFC-5005](http://tools.ietf.org/rfc/rfc5005/).

Some of the ATOM collections returned by the web service (for example, search results and change log resources) may be **paged** - divided into a series of linked feed resources. A **next** link points to the next page (the next feed resource in the series).

The following example shows a change log resource containing a link of this type:

```

<feed xmlns="http://www.w3.org/2005/Atom">
...
  <link rel="self"
    href="http://host-ip-address/webservice/escenic/changelog/publication/12/before/9425007?
count=10"/>
  <link rel="previous"
    href="http://host-ip-address/webservice/escenic/changelog/publication/12/after/9425006?
count=10"/>
  <link rel="next"
    href="http://host-ip-address/webservice/escenic/changelog/publication/12/before/9424970?
count=10"/>
...
</feed>

```

Note that many Atom feeds are ordered by date, with the newest entry first. This means that the **next** page of a change log will contain links to **older** resources, not **newer** ones.

3.1.4 previous

This link relation is defined in [IETF RFC-5005](http://tools.ietf.org/rfc/rfc5005/).

Some of the ATOM collections returned by the web service (for example, search results and change log resources) may be **paged** - divided into a series of linked feed resources. A **previous** link points to the previous page (the previous feed resource in the series).

The following example shows a change log resource containing a link of this type:

```

<feed xmlns="http://www.w3.org/2005/Atom">
...
  <link rel="self"
    href="http://host-ip-address/webservice/escenic/changelog/publication/12/before/9425007?
count=10"/>
  <link rel="previous"
    href="http://host-ip-address/webservice/escenic/changelog/publication/12/after/9425006?
count=10"/>
  <link rel="next"
    href="http://host-ip-address/webservice/escenic/changelog/publication/12/before/9424970?
count=10"/>
...
</feed>

```



Note that many Atom feeds are ordered by date, with the newest entry first. This means that the **previous** page of a change log will contain links to **newer** resources, not **older** ones.

3.1.5 related

This link relation is defined in [IETF RFC-4287](http://tools.ietf.org/rfc/rfc4287/).

It is used by the Content Engine web service to represent content item relations. You can navigate from a content item to its related content items by following **related** links.

The following example shows a link of this type:

```
<entry xmlns="http://www.w3.org/2005/Atom">
  <content type="application/vnd.vizrt.payload+xml">
    ...this content item in VDF format...
  </content>
  <link href="http://host-ip-address/webservice/escenic/content/1513"
        rel="related" type="application/atom+xml; type=entry"
        metadata:group="STORYREL">
    ...related content item's summary fields in VDF format...
  </link>
  <link href="http://host-ip-address/webservice/escenic/content/5391"
        rel="related" type="application/atom+xml; type=entry"
        metadata:group="STORYREL">
    ...related content item's summary fields in VDF format...
  </link>
  ...
</entry>
```

3.1.6 search

This link relation is defined in <http://www.opensearch.org/Specifications/OpenSearch/1.1>.

It identifies a link to an **OpenSearch** specification. When a **search** link appears in a feed resource returned by the web service, it indicates that the collection is searchable. The resource referenced by the **search** link is an OpenSearch document describing the format of the search requests that may be submitted to search the collection. Links of this type occur in feed resources representing content items collections.

The following example shows a feed resource containing a link of this type:

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/section/22/content-items</id>
  <link rel="self"
        href="http://host-ip-address/webservice/escenic/section/22/content-items"
        type="application/atom+xml"/>
  <updated>2010-06-01T07:08:10.427Z</updated>
  <author/>
  <title type="text">Content items for section with dbid="22"</title>
  <link rel="search"
        href="http://host-ip-address/webservice/open-search/escenic/22/content-search-
description.xml"
        type="application/opensearchdescription+xml"/>
</feed>
```

3.2 Proprietary Link Relations

3.2.1 <http://www.vizrt.com/types/relation/changelog>

This is a proprietary link relation. It identifies a link to an Escenic **change log** (see [section 1.3](#)). A change log is a paged collection (that is, a chained series of feed resources) containing links to all the changes made in a particular protection domain, ordered by modification date. For a detailed description of what a change log is, see [section 1.3](#). Links of this type can be found in

- The protection domain feed resource returned from a `http://www.vizrt.com/types/relation/protection-domains` link (see [section 3.2.8](#)).
- The sub-section feed resources returned from `down` links (see [section 3.1.1](#)).

The following example shows a protection domain feed resource containing links of this type:

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <id>http://host-ip-address/webservice/escenic/changelog</id>
  <title type="text">Protection domains accessible to demo_admin</title>
  <updated>2010-06-23T17:54:09.173Z</updated>
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <link rel="self"
    href="http://host-ip-address/webservice/escenic/changelog"
    type="application/atom+xml"/>
  <entry>
    <id>http://host-ip-address/webservice/escenic/publication/demo</id>
    <title type="text">demo</title>
    <link rel="http://www.vizrt.com/types/relation/changelog"
      href="http://host-ip-address/webservice/escenic/changelog/publication/1"
      type="application/atom+xml"/>
  </entry>
  <entry>
    <id>http://host-ip-address/webservice/escenic/section/3</id>
    <title type="text">topicRoot</title>
    <link rel="http://www.vizrt.com/types/relation/changelog"
      href="http://host-ip-address/webservice/escenic/changelog/section/3"
      type="application/atom+xml"/>
  </entry>
  <entry>
    <id>http://host-ip-address/webservice/escenic/section/4</id>
    <title type="text">New Articles</title>
    <link rel="http://www.vizrt.com/types/relation/changelog"
      href="http://host-ip-address/webservice/escenic/changelog/section/4"
      type="application/atom+xml"/>
  </entry>
  ...
</feed>
```

 A change log can contain either:

- Changes to ordinary content items, in which case the URI in the `href` attribute will end with `section/id`, or
 - Changes to person records (users, authors, registered site visitors and so on), in which case the URI in the `href` attribute will end with `publication/id`
-



3.2.2 <http://www.vizrt.com/types/relation/content-items>

This link relation identifies a link to a collection of Escenic content items. Links of this type occur in feed resources representing publication sections.

Currently, the feed resource returned by a link of this type is always empty (that is, the feed contains no **entry** elements). In future, it is the intention that this will not be the case, such links will return paged collections. In the mean time, however, a feed resource of this type is still useful:

- It provides a target for HTTP **POST** operations, allowing content items to be added to the section it represents
- It contain a **search** link, enabling the section to be searched (see [section 3.1.6](#)).

The following example shows a link of this type:

```
<entry>
...
  <link rel="http://www.vizrt.com/types/relation/content-items"
        href="http://host-ip-address/webservice/escenic/section/22/content-items"
        type="application/atom+xml"/>
...
</entry>
```

3.2.3 <http://www.vizrt.com/types/relation/home-section>

This link relation identifies a link to a content item's home section. Links of this type occur in entry resources representing content items.

The following example shows a link of this type:

```
<entry>
...
  <link rel="http://www.vizrt.com/type/relation/home-section"
        href="http://host-ip-address/webservice/escenic/section/4"
        title="New Articles"
        type="application/atom+xml; type=entry"/>
...
</entry>
```

3.2.4 <http://www.vizrt.com/types/relation/inboxes>

This link relation identifies a link to a collection of Escenic in-boxes. Links of this type occur in feed resources representing publication sections.

The following example shows a link of this type:

```
<entry>
...
  <link rel="http://www.vizrt.com/types/relation/inboxes"
        href="http://host-ip-address/webservice/escenic/section/22/inboxes"
        type="application/atom+xml"/>
...
</entry>
```

3.2.5 <http://www.vizrt.com/types/relation/lists>

This link relation identifies a link to a collection of Escenic lists. Links of this type occur in feed resources representing publication sections.

The following example shows a link of this type:

```
<entry>
  ...
  <link rel="http://www.vizrt.com/types/relation/lists"
        href="http://host-ip-address/webservice/escenic/section/22/lists"
        type="application/atom+xml"/>
  ...
</entry>
```

3.2.6 <http://www.vizrt.com/types/relation/list-pool>

This link relation identifies a link to a list pool - a collection containing links to all the content items in a list or inbox. Links of this type occur in feed resources representing a section's lists or inboxes.

The following example shows a link of this type:

```
<entry>
  ...
  <link href="http://host-ip-address/webservice/content/com.escenic.list-pool/423"
        rel="http://www.vizrt.com/types/relation/list-pool"
        type="application/xhtml+xml"/>
  ...
</entry>
```

3.2.7 <http://www.vizrt.com/types/relation/pages>

This link relation identifies a link to a collection of Escenic section pages. Links of this type occur in feed resources representing publication sections.

 Currently, links with this relation will not work: they will return an **HTTP 404** error.

The following example shows a link of this type:

```
<entry>
  ...
  <link rel="http://www.vizrt.com/types/relation/pages"
        href="http://host-ip-address/webservice/escenic/section/22/pages"
        type="application/atom+xml"/>
  ...
</entry>
```

3.2.8 <http://www.vizrt.com/types/relation/protection-domains>

This link relation identifies a link to a feed resource listing all the **protection domains** the requesting user has access to. One link of this type is included in the web service's "start" feed resource.

The following example shows the "start" feed resource containing a link of this type:

```
<feed xmlns="http://www.w3.org/2005/Atom">
```



```

...
<title type="text">Publication root sections you are authorized to.</title>
<link rel="http://www.vizrt.com/types/relation/protection-domains"
      href="http://host-ip-address/webservice/escenic/changelog"
      type="application/atom+xml"/>
...
</feed>

```

3.2.9 <http://www.vizrt.com/types/relation/publication>

This link relation identifies a link to the publication to which a content item belongs. Links of this type occur in entry resources representing content items.

The following example shows a link of this type:

```

<entry>
...
  <link rel="http://www.vizrt.com/type/relation/publication"
        href="http://host-ip-address/webservice/escenic/publication/demo"
        title="demo"
        type="application/atom+xml; type=entry"/>
...
</entry>

```

3.2.10 <http://www.vizrt.com/types/relation/section>

This link relation identifies a link to one of the sections in which a content item appears. Links of this type occur in entry resources representing content items.

The following example shows a link of this type:

```

<entry>
...
  <link rel="http://www.vizrt.com/type/relation/section"
        href="http://host-ip-address/webservice/escenic/section/4"
        title="New Articles"
        type="application/atom+xml; type=entry"/>
...
</entry>

```

3.2.11 <http://www.vizrt.com/types/relation/lock>

This link relation identifies a link that can be used to lock a content item or a specific fragment of the content item. Links of this type occur in entry resources representing content items.

The following example shows a link of this type:

```

<entry>
...
  <link rel="http://www.vizrt.com/types/relation/lock"
        href="http://host-ip-address/webservice/escenic/lock/article/4"/>
...
</entry>

```

3.2.12 <http://www.vizrt.com/types/relation/preview>

This relation identifies a link to a preview of a content item. Links of this type occur in entry resources representing content items.

The following example shows a link of this type:

```
<entry>
  ...
  < link href="http://host-ip-address/webservice/escenic/content/preview"
        rel="http://www.vizrt.com/types/relation/preview"/>
  ...
</entry>
```

4 Standard Element Usage

When the Content Engine's web service returns an Atom feed containing a content item, it maps various content item properties and fields to standard Atom elements. In addition, some elements from other well-known standards are used where suitable standard Atom elements are not available.

This section contains descriptions of the mappings and RDF extensions used by the web service.

When the web service returns content items in an Atom feed (for example, in a change log feed), the following Atom elements are used to hold specific content item properties or field values:

title

This Atom element holds the content item's title field (that is, the field that is nominated as the title field in the `content-type` resource - see). Since the content item title is included in the Atom entry in this way, the title field is omitted from the entry's VDF payload.

updated

This Atom element holds the content item's `Published` property.

published

This Atom element holds the content item's `First Published` property.

The following elements from the Atom Publishing Protocol (namespace URI <http://www.w3.org/2007/app>) and RDF Dublin Core (namespace URI <http://purl.org/dc/terms/>) are also used:

app:edited

This Atom Publishing Protocol element holds the content item's `Last Modified` property.

dcterms:created

This RDF Dublin Core element holds the content item's `Created` property.

dcterms:identifier

This RDF Dublin Core element holds the content item's `Id` property.

app:control

This Atom Publishing Protocol element holds sub-elements containing state information about the content item. For examples of how this element is used together with other state-related elements, see [section 4.1](#).

app:draft

This Atom Publishing Protocol element indicates whether or not the content item is in a draft state or not. For examples of how this element is used together with other state-related elements, see [section 4.1](#).

Metadata items for which no suitable standard elements could be found are included in the entries in the following proprietary namespaces:

`http://xmlns.escenic.com/2010/atom-metadata`

This namespace, which is usually given the prefix `metadata`, contains the following elements:

`metadata:creator`

The content item's `Creator` property. The content model of this element is an Atom person construct.

`metadata:reference`

The content item's `Reference` property.

`metadata:publication`

The content item's `Publication` property.

`metadata:deleted`

This identifies whether a content item is deleted or not. For examples of how this element is used together with other state-related elements, see [section 4.1](#).

`metadata:group`

The name of a content item relation type group. It is used to classify `rel="relation"` links in entry resources. For an example of how this element is used, see [section 6.6](#).

`http://www.vizrt.com/atom-ext`

This namespace, which is usually given the prefix `vaext`, contains the following elements:

`vaext:state`

This element is used to represent additional content item states that cannot be represented by standard Atom or Atom publishing states. It appears as a child of the `app:control` element. For examples of how this element is used together with other state-related elements, see [section 4.1](#).

4.1 Content Item State Representations

The Content Engine state concept is represented by a combination of several Atom feed elements. The following examples show the combinations used to represent each state:

Draft

```
<app:control>
  <app:draft>yes</app:draft>
</app:control>
```

Submitted

```
<app:control>
  <app:draft>yes</app:draft>
  <vaext:state>submitted</vaext:state>
</app:control>
```



Approved

```
<app:control>
  <app:draft>yes</app:draft>
  <vaext:state>approved</vaext:state>
</app:control>
```

Published

This state is represented by the **absence** of an `app:control` element, as specified in the Atom Publishing Protocol specification.

Deleted

```
<metadata:deleted>yes</metadata:deleted>
<app:control>
  <app:draft>yes</app:draft>
</app:control>
```



5 Viz Data Format

Viz Data Format is a proprietary format used to package Escenic content for delivery by the Content Engine's web service.

The VDF format in general allows for a document to contain either a **model**, which describes the structure of the data contained in another VDF document, or a **payload**, which contains actual data. This provides applications using the format with a means of understanding and handling new types of data.

The Content Engine's web service, however, currently only makes use of the payload format, which means that client applications must know the structure of the data in any content items downloaded from the web service. For a Content Engine web service client this is not really a problem, since the structure of all content items is also defined in a Content Engine **content-type** resource (see **Escenic Content Engine Resource Reference, chapter 2**).

The following example shows the general structure of the VDF documents returned by the web service:

```
<vdf:payload xmlns:vdf="http://www.vizrt.com/types"
  model="http://nightly.dev.escenic.com/webservice/content-descriptions/news">
  <vdf:field name="SUMMARY">
    <vdf:value>
      <div xmlns="http://www.w3.org/1999/xhtml">
        <p>Some XHTML content</p>
      </div>
    </vdf:value>
  </vdf:field>
  <vdf:field name="BODY">
    <vdf:value>
      <div xmlns="http://www.w3.org/1999/xhtml">
        <p>Some more XHTML content</p>
      </div>
    </vdf:value>
  </vdf:field>
  <vdf:field name="SCOREBOOST">
    <vdf:value>C</vdf:value>
  </vdf:field>
  <vdf:field name="OPTIONS">
    <vdf:list>
      <vdf:payload>
        <vdf:field name="OPTIONS">
          <vdf:value>C</vdf:value>
        </vdf:field>
      </vdf:payload>
    </vdf:list>
  </vdf:field>
</vdf:payload>
```

The root element is always a **payload** element. The individual elements are described in the following sections.

5.1 payload

The root element of the document.

Syntax

```
<payload
  model="uri"
  >
  <field>...</field>+
</payload>
```

Attributes

model="uri"

Normally, this would be the URI of a VDF model document, describing the structure of the data in this payload. In the case of the Content Engine web service, the URI is simply an identifier, and does not reference a resource of any kind.

5.2 field

A data field: it corresponds to a content item field or part of a complex content item field.

Syntax

```
<field
  name="..."
  >
  (<value>...</value>|<list>...</list>)
</field>
```

Attributes

name="..."

The name of the field (`body`, for example).

5.3 value

The content of a data field. It may consist either of a simple value (text) or a single element from a foreign namespace (typically an XHTML `div` element). The type of content depends on how the field is defined in the `content-type` resource for the publication to which the content item belongs.

Syntax

```
<value>
  (ANY-FOREIGN-ELEMENT*|text)
</value>
```

5.4 list

The content of a data field that contains multiple values. A `list` contains a `payload`, that can in turn contain the field's sub-fields.

Syntax

```
<list>
  <payload>...</payload>
</list>
```



6 How to...

This section contains examples showing the kinds of HTTP requests a client program can submit in order to achieve various common objectives, and the kinds of documents returned by the web service in response to these requests.

6.1 Get Started

No matter what your client program's purpose, the first thing it must always do is send an HTTP **GET** request to the web service's "start" URL:

```
curl -u user:password -X GET http://host-ip-address/webservice/escenic/section/ROOT/subsections
```

It must then parse the Atom feed returned by the web service and extract the URLs it needs:

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/section/ROOT/subsections</id>
  <link rel="self"
    href="http://host-ip-address/webservice/escenic/section/ROOT/subsections"
    type="application/atom+xml"/>
  <updated>2010-06-23T16:51:06.721Z</updated>
  <title type="text">Publication root sections you are authorized to.</title>
  <link rel="changelog"
    href="http://host-ip-address/webservice/escenic/changelog"
    type="application/atom+xml"/>
  <entry>
    <id>http://host-ip-address/webservice/escenic/section/1</id>
    <title type="text">frontpage</title>
    <updated>2010-06-22T10:16:46.309Z</updated>
    <category term="directory" scheme="http://www.vizrt.com/types"/>
    <link rel="edit-media"
      type="application/vnd.escenic.content+xml;type=com.escenic.section"
      href="http://host-ip-address/webservice/escenic/section/1"/>
    <link rel="down"
      href="http://host-ip-address/webservice/escenic/section/1/subsections"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/inboxes"
      href="http://host-ip-address/webservice/escenic/section/1/inboxes"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/lists"
      href="http://host-ip-address/webservice/escenic/section/1/lists"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/pages"
      href="http://host-ip-address/webservice/escenic/section/1/pages"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/content-items"
      href="http://host-ip-address/webservice/escenic/section/1/content-items"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/changelog"
      href="http://host-ip-address/webservice/escenic/changelog/section/1"
      type="application/atom+xml"/>
  </entry>
</feed>
```

Each entry represents the root section of a publication, and the links in each entry represent the operations available for that publication. (There is only one entry in the example feed above: this means that the user submitting the request only has access to one publication at this installation.)

6.2 Navigate The Section Hierarchy

A client application can navigate the section hierarchy by recursively following all links with the `rel` attribute `down` (see [section 3.1.1](#)).

Root section 1 in the "start" Atom feed shown in [section 6.1](#), for example, has a single `down` link that returns a new feed containing entries for all of section 1's subsections. So submitting the request:

```
curl -u user:password -X GET http://host-ip-address/webservice/escenic/section/1/subsections
```

returns a new Atom feed like this:

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/section/1/subsections</id>
  <link rel="self"
    href="http://host-ip-address/webservice/escenic/section/1/subsections"
    type="application/atom+xml"/>
  <updated>2010-06-23T17:35:45.180Z</updated>
  <title type="text">Subsections for Section with id=1</title>
  <entry>
    <id>http://host-ip-address/webservice/escenic/section/4</id>
    <title type="text">New Articles</title>
    <updated>2010-06-22T10:16:46.637Z</updated>
    <category term="directory" scheme="http://www.vizrt.com/types"/>
    <link rel="edit-media"
      type="application/vnd.escenic.content+xml;type=com.escenic.section"
      href="http://host-ip-address/webservice/escenic/section/4"/>
    <link rel="down"
      href="http://host-ip-address/webservice/escenic/section/4/subsections"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/inboxes"
      href="http://host-ip-address/webservice/escenic/section/4/inboxes"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/lists"
      href="http://host-ip-address/webservice/escenic/section/4/lists"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/pages"
      href="http://host-ip-address/webservice/escenic/section/4/pages"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/content-items"
      href="http://host-ip-address/webservice/escenic/section/4/content-items"
      type="application/atom+xml"/>
    <link rel="http://www.vizrt.com/types/relation/changelog"
      href="http://host-ip-address/webservice/escenic/changelog/section/4"
      type="application/atom+xml"/>
  </entry>
  ...
</feed>
```

Each entry in this feed contains a similar `down` link leading to further subsections that the client can follow. If a section has no subsections, then following the link will result in an empty feed like this:

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/section/4/subsections</id>
  <link
    rel="self"
    href="http://host-ip-address/webservice/escenic/section/4/subsections"
    type="application/atom+xml"/>
  <updated>2010-06-23T18:15:53.779Z</updated>
  <title type="text">Subsections for Section with id=4</title>
  <link href="http://host-ip-address/webservice/escenic/section/1/subsections"/>
</feed>
```



A client can therefore use these links to traverse the entire section tree of a publication.

6.3 Search For Content

Every entry in a subsection feed contains a link with the `rel` attribute `http://www.vizrt.com/types/relation/content-items`. This relationship indicates that the link points to a collection of all the content items belonging to the section. If a client application follows one of these links:

```
curl -u user:password -X GET http://host-ip-address/webservice/escenic/section/22/content-items
```

what is returned is an empty feed that contains no actual content item entries, but does contain a link with the `rel` attribute `search`.

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/section/22/content-items</id>
  <link rel="self"
        href="http://host-ip-address/webservice/escenic/section/22/content-items"
        type="application/atom+xml"/>
  <updated>2010-06-03T12:19:57.058Z</updated>
  <author/>
  <title type="text">Content items for section with dbid="22"</title>
  <link rel="search"
        href="http://host-ip-address/webservice/open-search/escenic/22/content-search-
description.xml"
        type="application/opensearchdescription+xml"/>
</feed>
```

If the client application then follows this link:

```
curl -u user:password -X GET http://host-ip-address/webservice/open-search/escenic/22/content-search-
description.xml
```

The web service returns an OpenSearch document describing the URL format required to search through the section's content items:

```
<OpenSearchDescription xmlns="http://a9.com/-/spec/opensearch/1.1/">
  <ShortDescription>Escenic search</ShortDescription>
  <Description/>
  <Url type="application/atom+xml"
        template="http://host-ip-address/webservice/publication/dev.nightly/search/escenic/22/
{searchTerms}/?pw={startPage?}&amp;c={count?}" />
  <Url type="text/html"
        template="http://host-ip-address/webservice/publication/dev.nightly/search/escenic/22/
{searchTerms}/?pw={startPage?}&amp;c={count?}&amp;format=html" />
  <Contact>http://www.escenic.com</Contact>
  <Tags/>
  <LongName>Escenic Content Engine Search</LongName>
  <Image height="16" width="16" type="image/x-icon">http://host-ip-address/webservice/images/
ece.ico</Image>
  <Query role="example" searchTerms="cat" />
  <Developer>Escenic AS</Developer>
  <Attribution/>
  <SyndicationRight>private</SyndicationRight>
  <AdultContent>false</AdultContent>
  <OutputEncoding>UTF-8</OutputEncoding>
  <InputEncoding>UTF-8</InputEncoding>
</OpenSearchDescription>
```

From this information the client can construct a URL that submits a query and specifies how the results are to be returned. Note that two different query templates are supplied: one that returns an Atom feed suitable for further

processing, and one that returns HTML suitable for display. To obtain an Atom feed, for example, the client might submit the following request:

```
curl -u user:password -X GET
'http://host-ip-address/webservice/publication/dev.nightly/search/escenic/22/Obama/?pw=1&count=10'
```

6.4 Change a Content Item

Once your client application has located a content item you are interested in, it can retrieve it, modify it and submit the changed version. The example in this section shows how to make the simplest of changes: a simple change to the title of a content item.

The first objective is to retrieve the content item:

```
curl -u user:password -X GET http://host-ip-address/webservice/escenic/content/43
```

The content item is returned as an Atom entry resource:

```
<entry xmlns="http://www.w3.org/2005/Atom"
  xmlns:app="http://www.w3.org/2007/app"
  xmlns:metadata="http://xmlns.escenic.com/2010/atom-metadata"
  xmlns:dcterms="http://purl.org/dc/terms/">
  <id>http://host-ip-address/webservice/escenic/content/43</id>
  <title type="text">Test</title>
  <app:edited>2010-06-23T09:09:50.654Z</app:edited>
  <dcterms:created>2010-06-22T10:22:20.000Z</dcterms:created>
  <author>
    <name>demo Administrator</name>
    <uri>http://host-ip-address/webservice/escenic/person/2</uri>
  </author>
  <dcterms:identifier>4</dcterms:identifier>
  <metadata:reference source="ece-auto-gen" sourceid="6d7203c9-27d5-4fce-b14a-a466ead83875"/>
  <link rel="http://www.vizrt.com/types/relation/home-section"
    href="http://host-ip-address/webservice/escenic/section/4"
    title="New Articles"
    type="application/atom+xml; type=entry"/>
  <link href="http://wrk-ermo:12345/demo/incoming/article4.ece"
    rel="alternate"/>
  <link href="http://host-ip-address/webservice/escenic/lock/article/43"
    rel="http://www.vizrt.com/types/relation/lock"/>
  <link rel="http://www.vizrt.com/types/relation/publication"
    href="http://host-ip-address/webservice/escenic/publication/demo"
    title="demo"
    type="application/atom+xml; type=entry"/>
  <metadata:creator>
    <name>demo Administrator</name>
  </metadata:creator>
  <metadata:publication href="http://host-ip-address/webservice/escenic/publication/demo">
    <link rel="http://www.vizrt.com/types/relation/home-section"
      href="http://host-ip-address/webservice/escenic/section/4"
      title="New Articles"
      type="application/atom+xml; type=entry"/>
    <link rel="http://www.vizrt.com/types/relation/section"
      href="http://host-ip-address/webservice/escenic/section/4"
      title="New Articles"
      type="application/atom+xml; type=entry"/>
  </metadata:publication>
  <link href="http://host-ip-address/webservice/escenic/content/43" rel="edit"/>
  <link href="http://host-ip-address/webservice/escenic/content/43" rel="self"/>
  <content type="application/vnd.vizrt.payload+xml">
    <vdf:payload xmlns:vdf="http://www.vizrt.com/types"
      model="http://host-ip-address/webservice/content-descriptions/another">
      <vdf:field name="BODY">
        <vdf:value>
          <div xmlns="http://www.w3.org/1999/xhtml">
            <p>This is a test</p>
          </div>
        </vdf:value>
      </vdf:field>
    </vdf:payload>
  </content>
</entry>
```



```
</vdf:payload>
</content>
</entry>
```

Your client application can now modify the title of the content item (highlighted in the above listing). If you are following this example using `curl` from the command line, you can simply copy the returned entry into a text editor and manually change the content of the `title` element:

```
<title type="text">Edited Title</title>
```

and save the results to a file (called `my-edited-item.xml`, for example).

In order to submit the updated content item, you must `PUT` the file you have saved to the same URI from which it was retrieved:

```
curl --include -u user:password -X PUT -H "If-Match: *" -H "Content-Type: application/atom+xml" \
> http://host-ip-address/webservice/escenic/content/4 --upload-file my-edited-article.xml
HTTP/1.1 100 Continue

HTTP/1.1 200 OK
Date: Thu, 24 Jun 2010 09:23:24 GMT
Content-Type: application/atom+xml;type=entry
Content-Length: 0
Server: Jetty(6.1.19)
```

In order for the `PUT` operation to work, you must specify two HTTP headers as shown above:

Content-Type: application/atom+xml

You must specify the content type of the data you are uploading.

If-Match: *

The `If-Match` header value `*` is used here for reasons of simplicity. It is acceptable to use it for test and demonstration purposes, but should **never** be used in a production system. For more information about this header and what it does, see [section 7.2](#).

If you're using `curl`, it's a good idea to specify `--include` with `PUT` operations: `curl` will then output the response header returned from the web service as shown above, and you can verify whether or not the operation was successful:

- A response code in the 2xx range indicates success.
- A response code in the 4xx range means that you made an invalid edit and the server won't accept your modification.
- A response code in the 5xx range means there is a server error.

Submitting a change in this way may not work if the resource you are attempting to modify is already locked by another client application. For more information about this and about how locking works, see [section 6.7](#).

6.5 Preview a Content Item

Before your client application submits a modified content item you may want to allow the user to preview the modified version. The example in this section shows how to preview a simple change to the title of a content item.

The first objective is to retrieve the content item:

```
curl -u user:password -X GET http://host-ip-address/webservice/escenic/content/43
```

The content item is returned as an Atom entry resource:

```
<entry xmlns="http://www.w3.org/2005/Atom"
  xmlns:app="http://www.w3.org/2007/app"
  xmlns:metadata="http://xmlns.escenic.com/2010/atom-metadata"
  xmlns:dcterms="http://purl.org/dc/terms/">
  <id>http://host-ip-address/webservice/escenic/content/43</id>
  <title type="text">Test</title>
  <app:edited>2010-06-23T09:09:50.654Z</app:edited>
  <dcterms:created>2010-06-22T10:22:20.000Z</dcterms:created>
  <author>
    <name>demo Administrator</name>
    <uri>http://host-ip-address/webservice/escenic/person/2</uri>
  </author>
  <dcterms:identifier>4</dcterms:identifier>
  <metadata:reference source="ece-auto-gen" sourceid="6d7203c9-27d5-4fce-b14a-a466ead83875"/>
  <link rel="http://www.vizrt.com/types/relation/home-section"
    href="http://host-ip-address/webservice/escenic/section/4"
    title="New Articles"
    type="application/atom+xml; type=entry"/>
  <link href="http://wrk-ermo:12345/demo/incoming/article4.ece"
    rel="alternate"/>
  <link href="http://host-ip-address/webservice/escenic/content/preview" rel="http://www.vizrt.com/
types/relation/preview"/>
  <link href="http://host-ip-address/webservice/escenic/lock/article/43"
    rel="http://www.vizrt.com/types/relation/lock"/>
  <link rel="http://www.vizrt.com/types/relation/publication"
    href="http://host-ip-address/webservice/escenic/publication/demo"
    title="demo"
    type="application/atom+xml; type=entry"/>
  <metadata:creator>
    <name>demo Administrator</name>
  </metadata:creator>
  <metadata:creator> o
  <metadata:publication href="http://host-ip-address/webservice/escenic/publication/demo">
    <link rel="http://www.vizrt.com/types/relation/home-section"
      href="http://host-ip-address/webservice/escenic/section/4"
      title="New Articles"
      type="application/atom+xml; type=entry"/>
    <link rel="http://www.vizrt.com/types/relation/section"
      href="http://host-ip-address/webservice/escenic/section/4"
      title="New Articles"
      type="application/atom+xml; type=entry"/>
  </metadata:publication>
  <link href="http://host-ip-address/webservice/escenic/content/43" rel="edit"/>
  <link href="http://host-ip-address/webservice/escenic/content/43" rel="self"/>
  <content type="application/vnd.vizrt.payload+xml">
    <vdf:payload xmlns:vdf="http://www.vizrt.com/types"
      model="http://host-ip-address/webservice/content-descriptions/another">
      <vdf:field name="BODY">
        <vdf:value>
          <div xmlns="http://www.w3.org/1999/xhtml">
            <p>This is a test</p>
          </div>
        </vdf:value>
      </vdf:field>
    </vdf:payload>
  </content>
</entry>
```

Your client application can now modify the title of the content item (highlighted in the above listing). If you are following this example using `curl`



from the command line, you can simply copy the returned entry into a text editor and manually change the content of the `title` element:

```
<title type="text">Edited Title</title>
```

and save the results to a file (called `my-edited-item.xml`, for example).

In order to submit the updated content item, you must `POST` the file you have saved to the preview URI identified by the `http://www.vizrt.com/types/relation/preview` relation:

```
curl --include -u user:password -X POST -H "Content-Type: application/atom+xml" \
> http://host-ip-address/webservice/escenic/content/preview --upload-file my-edited-article.xml
HTTP/1.1 100 Continue

HTTP/1.1 201 Created
Date: Mon, 13 Dec 2010 09:23:24 GMT
Location: http://host-ip-address/incoming/article1234.ece?token=-1723563771
Server: Jetty(6.1.19)
```

In order for the `POST` operation to work, you must specify one HTTP header as shown above:

Content-Type: application/atom+xml

You must specify the content type of the data you are uploading.

- A 201 response code indicates success. The `location` header contains the URL of the content item preview.
- A response code in the 4xx range means that you made an invalid edit and the server won't accept your modification.
- A response code in the 5xx range means there is a server error.

6.6 Follow Content Item Relations

This section shows how a content item's relations are included in Atom entry resources, and how you can navigate from a content item to its related content items.

If a retrieved content item has related content items, then they are always included in the returned Atom entry resource. They appear as `link` elements, with the `rel` attribute set to `related`. The following example shows the Atom entry resource for a content item with two related content items (highlighted in **bold**):

```
<entry xmlns="http://www.w3.org/2005/Atom" xmlns:vdf="http://www.vizrt.com/types"
  xmlns:ns2="http://www.w3.org/1999/xhtml" xmlns:ns="http://www.w3.org/2005/Atom"
  xmlns:metadata="http://xmlns.escenic.com/2010/atom-metadata" xmlns:dcterms="http://purl.org/
dc/terms/"
  xmlns:app="http://www.w3.org/2007/app">
  <id>http://host-ip-address/webservice/escenic/content/5403</id>
  <title type="text">UK and US terror alerts</title>
  <updated>2010-10-03T22:07:41.000Z</updated>
  <app:edited>2011-01-14T14:24:16.000Z</app:edited>
  <published>2010-10-03T22:07:41.000Z</published>
  <dcterms:created>2010-10-03T22:07:41.000Z</dcterms:created>
  <author>
    <name>Mark mark</name>
    <uri>http://host-ip-address/webservice/escenic/person/103</uri>
  </author>
  <dcterms:identifier>5403</dcterms:identifier>
```

```

<metadata:reference source="ece-auto-gen"
    sourceid="38f4f7a2-c2bb-4026-a79f-3bde3c3592df" />
<link href="http://host-ip-address/webservice/escenic/section/44"
    rel="http://www.vizrt.com/types/relation/home-section" title="News"
    type="application/atom+xml; type=entry" />
<link href="http://ecedemo5/wf-escenic-times/news/article5403.ece"
    rel="alternate" />
<link href="http://host-ip-address/webservice/escenic/publication/wf-escenic-times"
    rel="http://www.vizrt.com/types/relation/publication"
    title="wf-escenic-times" type="application/atom+xml; type=entry" />
<metadata:creator>
  <name>Mark mark</name>
</metadata:creator>
<metadata:publication href="http://host-ip-address/webservice/escenic/publication/wf-escenic-
times">
  <link href="http://host-ip-address/webservice/escenic/section/44"
    rel="http://www.vizrt.com/types/relation/home-section" title="News"
    type="application/atom+xml; type=entry" />

  <link href="http://host-ip-address/webservice/escenic/section/44"
    rel="http://www.vizrt.com/types/relation/section" title="News"
    type="application/atom+xml; type=entry" />
</metadata:publication>
<link href="http://host-ip-address/webservice/escenic/content/5403" rel="edit" />
<link href="http://host-ip-address/webservice/escenic/content/5403" rel="self" />
<link href="http://host-ip-address/webservice/escenic/lock/article/5403"
    rel="http://www.vizrt.com/types/relation/lock" />
<content type="application/vnd.vizrt.payload+xml">
  <vdf:payload model="http://host-ip-address/webservice/content-descriptions/news">
    ...content...
  </vdf:payload>
</content>
<link href="http://host-ip-address/webservice/escenic/content/1513"
    rel="related" type="application/atom+xml; type=entry"
    metadata:group="STORYREL">
  <vdf:payload model="http://host-ip-address/webservice/content-summaries/news">
    <vdf:field name="TITLE">
      <vdf:value>Obama rejects 'Biden option' on Afghanistan</vdf:value>
    </vdf:field>
    <vdf:field name="LEADTEXT">
      <vdf:value>President Obama ....</vdf:value>
    </vdf:field>
  </vdf:payload>
</link>
<link href="http://host-ip-address/webservice/escenic/content/5391"
    rel="related" type="application/atom+xml; type=entry"
    metadata:group="STORYREL">
  <vdf:payload model="http://host-ip-address/webservice/content-summaries/news">
    <vdf:field name="TITLE">
      <vdf:value>Travel alert issued for U.S. citizens in Europe</vdf:value>
    </vdf:field>

    <vdf:field name="LEADTEXT">
      <vdf:value>The United States ...</vdf:value>
    </vdf:field>
  </vdf:payload>
</link>
</entry>

```

The `rel="related"` attribute setting indicates that the link represents one of the content item's relations. The `metadata:group` attribute specifies the relation type group that the relation belongs to. Relation type groups are defined in the `content-type` publication resource (see **Escenic Content Engine Resource Reference, section 2.21**) and appear as separate relation drop zones in Content Studio.

Note that each relation `link` element also contains the relation's summary fields packaged in a VDF `payload` element. This makes it possible for interactive clients to display the relation summaries along with a content item, in the same way as Content Studio.



A client can retrieve related content items by following the links, for example:

```
curl -u user:password -X GET http://host-ip-address/webservice/escenic/content/1513
```

6.7 Lock a Content Item

A lock is a piece of advice to client applications that some other client is currently modifying all or part of a content item. The correct use of locks ensures that no two clients will simultaneously update the same information. A content item can be locked with two different types of lock:

Resource lock

Locks the entire resource (content item). No other locks can be set on this resource.

Fragment lock

Locks only part of the resource (one field of a content item). Several fragment locks can be set on the same content item as long as they lock different fields. A resource lock, however, cannot be set on a content item locked in this way.

An interactive client such as Content Studio is expected to lock a content item with the appropriate type of lock as soon as a user starts making a modification, and unlock it again when the user either saves the change or explicitly abandons it (for example by clicking on a **Cancel** button). Similarly, an automated client should lock a content item as soon as it "intends" to carry out an update, both to inform other clients of its intentions and to avoid conflicts when it actually performs the update.

The following terms are used when talking about locks:

Acquire

Locking all or part of a content item is described as **acquiring a lock**.

Release

Unlocking all or part of a content item is described as **releasing a lock**. The term **forced release** is used to describe release of a lock owned by another client.

Lock URI

Locks are exposed by the Content Engine's web service as URIs called **lock URIs**. There are three types of lock URI:

lock collection URI

This URI is returned in a content item resource entry, and is the start point for locking a content item.

Public lock URI

These URIs appear in entries returned from a lock collection URI, and simply indicate that a content item is locked.

Private lock URI

This URI is returned to the client that acquires a lock, and is effectively the key to the lock.

Updating a content item involves the following steps:

1. **GET** the resource you want to lock.
2. Optionally, **GET** the resource's public locks (if any) from the resource's lock collection URI.
3. If the resource's existing public locks do not conflict with the lock you require, **POST** a lock request to the lock collection URI. If successful, the web service returns a private lock URI.
4. **PUT** the modified version of the resource. In order for the **PUT** request to succeed you must include the private lock URI in its HTTP header.
5. **DELETE** the private lock URI in order to release your lock.

These steps are described in greater detail in the following sections.

6.7.1 Get the Resource to Lock

Your client application must first **GET** the resource to be locked in order to discover its lock collection URI:

```
curl -u user:password http://host-ip-address/webservice/escenic/content/4
```

The returned resource entry contains a link with the relation type [section 3.2.11](#). This link contains the lock collection URI:

```
<entry>
  ...
  <link href="http://host-ip-address/webservice/escenic/lock/article/4"
        rel="http://www.vizrt.com/types/relation/lock"/>
  ...
</entry>
```

6.7.2 Get Public Locks

Before the client attempts to lock the resource it can optionally check to see whether or not a conflicting lock is already set. This is done by sending a **GET** request to the lock collection URI. For example:

```
curl -u user:password http://host-ip-address/webservice/escenic/lock/article/4
```

The following example shows a response indicating that the resource is already locked:

```
<feed xmlns="http://www.w3.org/2005/Atom">
  <author>
    <name>Escenic Content Engine</name>
  </author>
  <id>http://host-ip-address/webservice/escenic/lock/article/4</id>
  <link rel="self" href="http://host-ip-address/webservice/escenic/lock/article/4" type="application/atom+xml"/>
  <updated>2010-12-04T08:05:08.474Z</updated>
  <entry xmlns:metadata="http://xmlns.escenic.com/2010/atom-metadata">
    <id>http://host-ip-address/webservice/escenic/lock/article/4/public/72</id>
    <title type="text">Lock of the fragment SUMMARY</title>
    <link href="http://host-ip-address/webservice/escenic/lock/article/4/public/72" rel="self"/>
    <metadata:fragment>SUMMARY</metadata:fragment>
    <author>
```



```

    <name>testpub Administrator</name>
  </author>
  <summary type="text">Lock created by: lpt-sai/10.211.10.176</summary>
  <content type="text">The fragment 'SUMMARY' of the resource is locked by testpub Administrator</
content>
</entry>
</feed>

```

The returned feed contains an entry with the public lock URI `http://host-ip-address/webservice/escenic/lock/article/4/public/72`. The `metadata:fragment` element, however, indicates that this is a fragment lock that only locks the content item's `SUMMARY` field. If the client does not intend to lock the same fragment or the whole resource, then it can still acquire a lock.

This step is optional: a client can POST a lock request without first checking the existing locks, and in many cases it may be most efficient to do so. An interactive clients will, however, often need to `GET` this list of public locks in order to present the information in its user interface.

6.7.3 Post Lock Request

To acquire a lock the client must send a `POST` request to the resource's lock collection URI. The posted file must contain a valid Atom entry element like this:

```

<entry xmlns="http://www.w3.org/2005/Atom" xmlns:metadata="http://xmlns.escenic.com/2010/atom-
metadata">
  <summary type="text">Lock created by: User1 from WS</summary>
  <metadata:fragment>TITLE</metadata:fragment>
</entry>

```

The `metadata:fragment` element is only required when requesting a fragment lock. Omit it to request a resource lock.

The request must contain a `Content-Type` header the value `application/atom+xml;type=entry`. A valid request might look like this:

```

curl --include -u user:password -X POST -H "Content-Type: application/atom+xml;type=entry" \
  http://host-ip-address/webservice/escenic/lock/article/4 \
  --upload-file atom-entry-file

```

where `atom-entry-file` is the path of a file containing an Atom entry like the one shown above.

If the lock request succeeds, then the web service returns a `201 Created` response in which `Location` contains the resource's private lock URI:

```

HTTP/1.1 201 Created
Server: Apache-Coyote/1.1
Location: http://host-ip-address/webservice/escenic/lock/article/4/private/9
Content-Length: 0
Date: Sat, 04 Dec 2010 02:35:38 GMT

```

If the lock request fails because the requested resource is already locked someone else has already locked the resource then the web service returns a `409 CONFLICT` response.

Once the client has acquired the private lock URI, it can proceed with the update as described in [section 6.7.4](#). If, however, the client needs to access

the lock's public URI for some reason, it can do so by sending a **GET** request to the private lock URI:

```
curl --include -u user:password http://host-ip-address/webservice/escenic/lock/article/4/private/9
```

This results in a redirect to the public lock URI:

```
HTTP/1.1 303 See Other
Server: Apache-Coyote/1.1
Location: http://host-ip-address/webservice/escenic/lock/article/4/public/39
Content-Length: 0
Date: Sat, 04 Dec 2010 02:38:53 GMT
```

Sending a **GET** request to the public lock URI:

```
curl -u user:password http://host-ip-address/webservice/escenic/lock/article/4/public/39
```

then returns an entry containing information about the lock:

```
<?xml version="1.0"?>
<entry xmlns="http://www.w3.org/2005/Atom" xmlns:metadata="http://xmlns.escenic.com/2010/atom-metadata">
  <id>http://host-ip-address/webservice/escenic/lock/article/4/public/39</id>
  <title type="text">Lock of the fragment TITLE</title>
  <link href="http://host-ip-address/webservice/escenic/lock/article/4/public/39" rel="self"/>
  <metadata:fragment>TITLE</metadata:fragment>
  <author>
    <uri>http://host-ip-address/webservice/escenic/person/1</uri>
    <name>publ Administrator</name>
  </author>
  <summary type="text"/>
  <content type="text">The fragment 'TITLE' of the resource is locked by publ Administrator</content>
</entry>
```

6.7.4 Put Updated Resource

Once the client has a private lock URI, it can **PUT** an update request. The request must contain the following headers:

- **Content-Type**, set to **application/atom+xml; type=entry**
- **X-Escenic-Locks**, containing the private lock URI. If you want to request several fragment locks at once, you can do so by repeating this header.
- **If-Match**, set to the appropriate **E-Tag** value. See [section 7.2](#) for more information about **E-Tag** and the use of the **If-Match** header.

Here is an example update request:

```
curl --include -u user:password -X PUT \
-H "Content-Type: application/atom+xml;type=entry" \
-H "X-Escenic-Locks: http://host-ip-address/webservice/escenic/lock/article/4/private/9" \
-H "If-Match: "129694559e911ee4c6d04212982" \
http://host-ip-address/webservice/escenic/content/4 --upload-file atom-entry-file
```

where *atom-entry-file* is the path of a file containing the updated resource:

```
<?xml version="1.0"?>
<entry xmlns="http://www.w3.org/2005/Atom" xmlns:app="http://www.w3.org/2007/app"
  xmlns:metadata="http://xmlns.escenic.com/2010/atom-metadata" xmlns:dcterms="http://purl.org/dc/terms/">
  <id>http://host-ip-address/webservice/escenic/content/4</id>
  <title type="text">My Updated Title</title>
  <updated>2010-12-01T10:13:00.000Z</updated>
  <app:edited>2010-12-03T10:13:06.000Z</app:edited>
  <published>2010-12-01T10:13:00.000Z</published>
  <dcterms:created>2010-12-01T10:13:00.000Z</dcterms:created>
  <author>
```



```

    <name>Marcus Tullius Cicero</name>
    <uri>http://host-ip-address/webservice/escenic/person/2</uri>
  </author>
  <dcterms:identifier>4</dcterms:identifier>
  <metadata:reference source="ece-auto-gen" sourceid="0b8a67bc-7b2d-493b-b6b5-d659f1beff81"/>
  <link rel="http://www.vizrt.com/types/relation/home-section" href="http://host-ip-address/
webservice/escenic/section/3" title="Examples" type="application/atom+xml; type=entry"/>
  <link href="http://host-ip-address/publ/Examples/article4.ece" rel="alternate"/>
  <link rel="http://www.vizrt.com/types/relation/publication" href="http://host-ip-address/
webservice/escenic/publication/publ" title="publ" type="application/atom+xml; type=entry"/>
  <metadata:creator/>
  <metadata:publication href="http://host-ip-address/webservice/escenic/publication/publ">
    <link rel="http://www.vizrt.com/types/relation/home-section" href="http://host-ip-address/
webservice/escenic/section/3" title="Examples" type="application/atom+xml; type=entry"/>
    <link rel="http://www.vizrt.com/types/relation/section" href="http://host-ip-address/webservice/
escenic/section/3" title="Examples" type="application/atom+xml; type=entry"/>
  </metadata:publication>
  <link href="http://host-ip-address/webservice/escenic/content/4" rel="edit"/>
  <link href="http://host-ip-address/webservice/escenic/content/4" rel="self"/>
  <link href="http://host-ip-address/webservice/escenic/lock/article/4" rel="http://www.vizrt.com/
types/relation/lock"/>
  <content type="application/vnd.vizrt.payload+xml">
    <vdf:payload xmlns:vdf="http://www.vizrt.com/types" model="http://host-ip-address/webservice/
content-descriptions/news">
      <vdf:field name="SUMMARY">
        <vdf:value>More text for Latin lovers</vdf:value>
      </vdf:field>
      <vdf:field name="BODY">
        <vdf:value>
          <div xmlns="http://www.w3.org/1999/xhtml">
            <p>Simple Body</p>
          </div>
        </vdf:value>
      </vdf:field>
    </vdf:payload>
  </content>
</entry>

```

If the `PUT` request succeeds, then the web service returns a 200 `OK` response:

```

HTTP/1.1 100 Continue

HTTP/1.1 200 OK
Server: Apache-Coyote/1.1
Content-Type: application/atom+xml;type=entry
Content-Length: 0
Date: Sat, 04 Dec 2010 02:53:02 GMT

```

In most cases the fragment identifier used in lock requests corresponds directly to the name of a content item field. The fragment identifier `TITLE` used in the above example, however, is a special case. It corresponds to whatever content item field has been designated as the title in the publication's `content-type` resource. This field is never returned in an entry resource's VDF payload: its content is instead returned as the entry's `title` element. In order to update this field, therefore, the client must actually update the entry resource's `title` element as shown in the above example.

For information about how a field is designated as content item's title, see **Escenic Content Engine Resource Reference, section 5.20**.

6.7.4.1 Updating a Resource Without a Lock

A client can in fact attempt to `PUT` an updated resource without first acquiring a lock by simply omitting the `X-Escenic-Locks` header from the request. The web service then attempts to:

1. Acquire a resource lock on the client's behalf.
2. Perform the update.
3. Release the lock.

If it cannot do so because the resource (or a fragment of the resource) is already locked, then it returns a **409 CONFLICT** response.

Note that in this case the web service always treats the request as a resource lock request, even if a fragment lock would be sufficient to carry out the submitted update.

6.7.5 Delete Private Lock

Once the resource has been updated (or if the update is to be abandoned for some reason) the client can release the lock by submitting a **DELETE** request to the private lock URI:

```
curl --include -u username:password -X DELETE http://host-ip-address/webservice/escenic/lock/article/4/private/9
```

If the **DELETE** request succeeds, then the web service returns a **204 No Content** response:

```
HTTP/1.1 204 No Content
Server: Apache-Coyote/1.1
Date: Sat, 04 Dec 2010 02:57:09 GMT
```

It is possible to force the release of a lock held by another client by submitting a **DELETE** request to the public lock URI:

```
curl --include -u user:password -X DELETE http://host-ip-address/webservice/escenic/lock/article/4/public/39
```

If the **DELETE** request succeeds, then the web service returns a **204 No Content** response:

```
HTTP/1.1 204 No Content
Server: Apache-Coyote/1.1
Date: Sat, 04 Dec 2010 03:09:38 GMT
```

If a client fails to delete a private lock URI, the lock will be released by the web service after 24 hours. To keep a resource locked for longer than this, a client must release its lock and immediately acquire a new one.

.....
Any attempt to update a resource using a deleted private lock URI will fail.
.....

7 Advanced Features

This chapter discusses use of some of the Content Engine web service's more advanced features.

7.1 Using Change Logs Effectively

The change log resources (described in [section 1.3](#)) expose all the content items in a protection domain, ordered by their last date of modification.

As with most Atom feeds, the usual way to use them is to poll the original feed URI (for example `http://host-ip-address/web-service/escenic/changelog/section/nnn`) at regular intervals, in order to see what's changed. If the number of changes exceeds the page size, the client can follow the **next** links until it reaches the page containing the last entry it processed.

Although this approach is a common way to use Atom feeds, it is less than optimal if the polling frequency is high, or if many clients are involved, as may be the case with the Content Engine web service. The first page of the change log, an Atom feed of several kilobytes, must be re-rendered every time the original feed URI is polled. Since most of the Content Engine's caches are not used due to authentication, the server is required to serve the same page very frequently, which has a detrimental effect on performance, and does not scale well with multiple clients or higher polling frequencies.

For this reason, the change log feeds include links that go right to the "edge" of the list, allowing clients to poll the actual start of the list instead of just the first page. This "edge" manifests itself as a **previous** link on the first page. It may initially seem illogical to have a **previous** link, implying as it does that there is a page **before** the first page. It is, however, very useful: there is a page before the first page, but it is empty unless a new change has been added to the log since the last poll. Once a client has polled the original feed once, it can subsequently poll the **previous** link. If no changes have been added, then the web service will return an empty feed, consuming very few server resources.

When new content items are added to the change log (or existing content items are modified and therefore moved to the top of the change log), polling the **previous** link will no longer return an empty feed: it will contain entries for all the new changes. The client can deal with these changes, and then continue polling **this** feed's **previous** link.

The consequence of using this approach is that most clients will be polling an empty feed most of the time, and therefore consuming a relatively small amount of server resources: the server is only required to actually render content items when there are changes to communicate to the clients. It also makes client application code simpler, since the client does not need to remember the last entry it processed, and compare it with each entry to

determine whether it is new: all entries in the feed returned by the **previous** link are guaranteed to be new, and do not need to be checked.

7.2 Optimistic Concurrency

The Content Engine relies on a technique called **optimistic concurrency** to deal with concurrent updates to content items. Clients are expected to co-ordinate updates between themselves using optimistic concurrency mechanisms. For background information about this technique, see <http://www.w3.org/1999/04/Editing/>.

Clients are required to **either** make active use of optimistic concurrency when updating content, **or** explicitly bypass it (not recommended), so you need to understand this section in order to make a client that will update content items.

Optimistic concurrency works as follows:

1. Each client that wants to make a modification to a content item downloads a copy of it. Each item is accompanied by an HTTP header called **Etag**. An **Etag** header is a unique identifier generated by the web service, and looks something like this:


```
Etag: "129694559e911ee4c6d04212982"
```
2. Each client remembers the **Etag** associated with the content item it is modifying.
3. After modifying the content item, the client performs a **PUT** operation to return it to the web service, including its **Etag** in an **If-Match** header. Note that an **Etag** value **includes its enclosing quote marks**. A correct **PUT** operation executed with **curl** would therefore look something like this:

```
curl --include -u user:password -X PUT -H 'If-Match: "129694559e911ee4c6d04212982"' \
> -H "Content-Type: application/atom+xml" http://host-ip-address/webservice/escenic/content/4 \
> --upload-file my-edited-article.xml
```

4. The web service checks the supplied **Etag** to see if it matches the current copy of the content item. For the first client to return changes, this will be the case, and the **PUT** operation will succeed. For all other clients that have been working on the content item, the **Etag** will not match and the web service will return a **412 PRECONDITION FAILED** response.

When a client receives a **412 PRECONDITION FAILED** response it should download the latest copy of the content item (with its new **Etag**), re-apply its changes to the new copy, and then **PUT** it back, hopefully succeeding this time. If more than two clients are working on the same copy, however, this might not be the case. In the worst case a client might need to retry many times. Clients should be able to handle a certain number of retries, back off after a few retries, and lower the rate of retry, and possibly even request human intervention.

A client can circumvent the optimistic concurrency mechanism by sending the header "**If-Match: ***" in its **PUT** operations. This header specifies that



the operation should succeed no matter which version is current. The web service will then accept the **PUT** request without performing any concurrency checks. The consequence of doing so is that any changes made to the content item by other clients will be overwritten. Your client should therefore **never** do this unless you know the consequences and are certain that it will not cause problems.