King's Research Portal

DOI:

[10.1080/19386389.2012.661689](https://doi.org/10.1080/19386389.2012.661689)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Gartner, R. (2012). The digital object in context: using CERIF with METS. *Journal of Library Metadata*, 12(1), 39-51. <https://doi.org/10.1080/19386389.2012.661689>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

**Open Access document
downloaded from King's Research Portal
<https://kclpure.kcl.ac.uk/portal>**

Citation to published version:

[Gartner, R. (2012). The digital object in context: using CERIF with METS. Journal of Library Metadata, 12(1), 39-51, doi: 10.1080/19386389.2012.661689]

The published version is available at:

DOI: [10.1080/19386389.2012.661689]

This version: [Post Print/Author Final Version]

URL identifying the publication in the King's Portal:

[[https://kclpure.kcl.ac.uk/portal/en/publications/the-digital-object-in-context-using-cerif-with-mets\(0be11e1d-2854-4992-b8e4-14491bba5fa7\).html](https://kclpure.kcl.ac.uk/portal/en/publications/the-digital-object-in-context-using-cerif-with-mets(0be11e1d-2854-4992-b8e4-14491bba5fa7).html)]

The copyright in the published version resides with the publisher.

When referring to this paper, please check the page numbers in the published version and cite these.

General rights

Copyright and moral rights for the publications made accessible in King's Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications in King's Research Portal that users recognise and abide by the legal requirements associated with these rights.'

- Users may download and print one copy of any publication from King's Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the King's Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

This is an electronic version of an article published in *Journal of Library Metadata*, 12(1), pp. 39–51. Journal of Library Metadata is available online at:

<http://www.tandfonline.com/doi/full/10.1080/19386389.2012.661689>

The digital object in context: using CERIF with METS

Running title: The digital object in context

Abstract

The article examines the potential for using the Common European Research Information Format in conjunction with the Metadata Encoding and Transmission Standard to provide contextual information for a digital research output. Both are key standards within their respective communities (the former in research information management, the latter in digital library metadata), but each employs a different approach to information architecture which renders their combination problematic. The article examines three possible ways to using CERIF and METS in conjunction, and suggests possible changes to approach of the METS standard to resolve some of the problems that arise.

Introduction

A report in the Technology and Standards Watch series published by the UK's Joint Information Systems Committee (Gartner 2008) argues for an integrated

metadata strategy for digital library objects based on the METS (Metadata Encoding and Transmission Standard) XML schema (Library of Congress 2011a) . The key advantage of using METS which this report cites is its ability to incorporate external XML schemas (extension schemas) for descriptive or administrative metadata while maintaining a single overall architecture (Gartner 2008, p.15); this, it is argued, provides a degree of integration whilst maintaining the flexibility of using the metadata schemes most relevant to a given application.

Common extension schemas used with METS include MODS (Metadata Object Description Schema) for descriptive (Library of Congress 2010) and PREMIS (PREservation Metadata: Implementation Strategies) for preservation metadata (Library of Congress 2011b). These have proved themselves highly effective in meeting the needs of the users and administrators of digital libraries, but in the case of digital objects which are themselves research outputs (for instance, journal articles held in an institutional repository) further contextual information on the research environments in which they were created may be particularly valuable in critically evaluating their contribution to the academic record.

A standard for encoding information of this type has been existence for over ten years in the form of the CERIF (Common European Research Information Format) data model (European Organisation for International Research Information 2010).

This scheme is now well established, particularly in Europe, as the core standard for interoperable research information management; it is, for instance, the recommended format for exchanging this data in the United Kingdom (Rogers, Huxley & Ferguson 2010, p.23), and will be a submission format for the UK's next national research assessment exercise for higher education (Bolton 2010, p.21).

This article argues that there is a strong case for using CERIF within METS as an adjunct to more standard metadata schemes in order to provide important contextual information for a digital object. However, the complexity of the CERIF standard, which is based on a data model derived from relational database tables, raises questions as to the usability of METS's currently limited facilities for addressing the contents of the metadata it embeds or references. Although an approach to resolving this within the current METS framework of metadata 'buckets' is discussed here, it is arguable that METS needs to be extended to allow more sophisticated approaches to metadata in general.

The research object in context

The need to provide contextual information on a research output is inherent in the concept of the Current Research Information System (CRIS) which holds information on the projects, organisations, people and funding that contribute to the production of the output itself (Jeffery et al. 2002, p.78). Such information is important to a wide range of potential stakeholders in the

research environment. Most obviously, funding authorities require this information for allocating resources: the UK's 2008 Research Assessment Exercise (RAE), for instance, on the basis of which much of its research funding to higher education institutions was distributed, required a detailed set of information on such factors as funding sources, studentships, numbers of research students, research groupings and the overall research environment (Higher Education Funding Council for England 2007) in order to allow a valid assessment of the value of the outputs submitted for evaluation.

For the individual researcher, contextual information of this type may also be valuable in assessing the relevance and potential value of their peers' outputs. A researcher may, for instance, want to evaluate a paper in the sciences by finding out what software was used in the experiment it documents or by examining the original raw data on which the paper is based (Jeffery, Lopatenko & Asserson 2002, p.79). A counterpart in the humanities or social sciences may perhaps attempt to deduce the limitations or potential biases in a work by examining the context of the institutional setting in which it was produced (and its possible effects on research agendas) or by examining the funding sources which supported its creation. In current environments where the overall direction of so much research is dictated by the priorities of funding bodies, information of this type can be invaluable as an adjunct to more standard methods for assessing the relevance of a work to the individual's own research.

A metadata environment for the digital research output

Where the output of research is available in digital form, particularly as part of an institutional repository, similar metadata requirements apply as in any digital library environment. In particular, a packaging standard is necessary to structure the often complex diversity of descriptive, administrative and structural metadata necessary to ingest, maintain and deliver the object. The most obvious choice for this is METS which is already well established as standard of this type for digital library metadata.

METS can function as a Submission Information Package (SIP), Archival Information Package (AIP) or Dissemination Information Package (DIP) under the OAIS (Open Archival Information System) (Consultative Committee for Space Data Systems 2009), and so can be used throughout the submission, archiving and delivery chain for digital objects. It has also been used successfully in institutional repositories, and is supported by such key repository systems as ePrints (University of Southampton 2010), Dspace (Massachusetts Institute of Technology 2010) and Fedora (Fedora Project 2006). For all of these reasons, using METS as the underlying architecture for an integrated metadata strategy for digital research objects is a sensible option.

The choice of CERIF for providing the contextual information for a research object is equally compelling. This standard provides the only comprehensive,

interoperable mechanism for such metadata, and has now gained sufficient traction as an approved standard within higher education sectors to be regarded as the only viable option for this.

Using METS and CERIF in conjunction therefore appears to be a sensible option for a digital repository or CRIS: the former brings all of the advantages of a coherent packaging standard which is now well established in the digital library community, the latter the ability to contextualise the digital research output using an interoperable and highly flexible scheme which has established itself as the *lingua franca* for research information management.

Using CERIF for research information management

CERIF was initially published as a data model realised in a set of relational SQL tables. Since 2006, it has also been made available as a set of 192 XML schemas which duplicate the architecture of the SQL original. The CERIF model in both SQL and XML defines a small set of entities and then provides an extensive set of linking mechanisms to realise their relationships in a working environment.

The core of CERIF is a set of very basic 'base' information entities, which document projects, people and organisational units: for projects, for instance, the information recorded here limited to fields for an internal ID, URI,

acronym, start and end dates, title, abstract and keywords. This core set is supplemented by a further small set of 'result' entities which record information on research outputs (publications, patents and products). These again include only basic components: for publications, for instance, only the core metadata required to identify a work (such as identification numbers, date and pagination) is included here.

A further set of 16 entities, termed "second-level" in the CERIF specification, includes an array of subsidiary concepts which may be used to qualify project, person, organisational unit or result entities: these include such components as metrics, events, qualifications, facilities, equipment or expertise. Another set of 49 entities handles the multi-lingual features of a CERIF application, allowing multiple language versions of any textual information to be encoded: for example, multiple language versions of the title of this article would be rendered as follows:-

<cfResPublTitle>

<cfResPublId>1abc</cfResPublId>

<cfTitle cfLangCode="en-UK">The digital object in context: using CERIF with METS</cfTitle>

<cfTitle cfLangCode="de">Das digitale Objekt im Kontext: CERIF mit METS benutzen</cfTitle>

<cfTitle cfLangCode="it">L'oggetto digitale nel contesto: utilizzando CERIF con METS</cfTitle>

</cfResPublTitle>

The core of the CERIF application is a set of 95 linking entities which mirror the relational database linkages of the original data model and allow the base, result, second-level and multi-lingual entities to be joined together. These entities usually contain nothing more than the IDs of the two components being linked, the semantic terms which establish the nature of the link itself and start and end dates for its validity: for instance, linking a researcher to their institutional affiliation may be done as shown below, where the **cfClassId** element contains the semantic term and the **cfClassSchemeId** indicates the vocabulary from which it is taken:-

<cfPers_OrgUnit>

<cfPersId>9876543210123</cfPersId>

<cfOrgUnitId>9999/17/A</cfOrgUnitId>

<cfClassId>class-is-affiliated-with</cfClassId>

<cfClassSchemeId>class-scheme-pers-orgunit-roles</cfClassSchemeId>

<cfStartDate>2002-01-02T00:00:00-00:00</cfStartDate>

<cfEndDate>2007-10-31T00:00:00-00:00</cfEndDate>

</cfPers_OrgUnit>

All **cfClassIds** and **cfClassSchemeIds**, and their associations, are in their turn defined in “Class” schemas.

A CERIF application therefore requires the use of a substantial number of XML instances, each conforming to its own schema; these are then linked together by a coherent system of IDs and a series of semantic vocabularies.

METS and CERIF: possible approaches to integration

Several possible approaches could feasibly be adopted to allow CERIF to be integrated with the METS framework. The simplest is to use CERIF as an extension schema in a manner analogous to MODS or PREMIS and embed or reference its component XML files from within METS’s metadata sections. A further approach may be to embed a simpler and more constrained XML file conforming to an ‘intermediary’ schema from which the CERIF files could be derived. A third approach, which perhaps extends METS beyond its intended functionality, would be to attempt to serialize the relationships expressed within the CERIF files in

the METS structural map.

CERIF as an extension schema

Because CERIF is available in XML form, it is feasible to either embed all of the CERIF files associated with a digital research output within METS `<mdWrap>` elements or hold them externally and reference them from within `<mdRef>` elements within the `<dmdSec>` or `<admSec>` sections. This is by far the simplest approach and fits cleanly within the METS model for extension schemas.

An initial problem with this approach is a degree of obscurity as to where in METS's architecture the CERIF files belong. Although some, such as those which provide bibliographic information akin to such standards as MODS or MARC, fit neatly into METS's `<dmdSec>` element for descriptive metadata, the majority of files, particularly the linking components which form the majority of a CERIF application, do not fall so readily into the METS framework. These files may perhaps be rationally considered part of an object's administrative metadata, but few fit neatly into the four sub-components (technical, source, rights or digital provenance) into which METS's section for this type of metadata, `<amdSec>`, is divided.

The problems of fitting some types of metadata into the METS taxonomy have been noted by a number of implementers (METS Editorial Board 2010, p.11), but at

present METS's requirement that metadata (in particular administrative metadata) is subdivided into separate "buckets" prevents any clean fit with CERIF. Although unsatisfactory in many ways, for the purpose of this article CERIF files will be put into METS `<dmdSec>` elements, an approach analogous to that taken by Habing and Cole (2008) in encoding aggregation RDF metadata within METS.

More difficulties arise from the complexity of the CERIF standard itself. As shown above, a CERIF application may involve a total of 192 separate files linked together by a complex, and highly flexible, system of semantic classes. Navigating this set of linkages in order to extract relevant metadata and present it in a meaningful way requires complicated processing and extensive documentation of the data architecture employed in a given CERIF application.

These problems become more acute when the METS framework is intended to facilitate interoperability between systems. The use of METS for digital object interchange has been acknowledged as one of its less successful features and some writers, such as Maslov, claim with some reason that it can only function as a packaging rather than exchange protocol as it lacks the specificity needed to generate an unambiguous interpretation of its encoded data and metadata (Maslov et al. 2010).

The standard mechanism for facilitating exchange between METS applications is

the METS Profile, which documents key features, such as the external schemas and structural requirements, of an application in a standardised way. It are not, however, intended to be machine actionable: all parts are merely human-readable descriptions which must be actioned by those responsible for the design of the receiving system. Consequently, it prone to inconsistencies and idiosyncracies, and lacks the specificity which Maslov et al. point to as an impediment to interoperability.

These problems become particularly acute when an application as complex as CERIF is incorporated into METS. Interpreting the possible 192 files that make up an application requires a much greater degree of precision of documentation than a simpler, unitary schema such as MODS necessitates, and even careful recording of this type within a METS Profile will require an intimidating degree of analysis and development by the administrators of a system ingesting a CERIF-enabled METS file before it can sensibly be actioned.

While this simple use of METS as an extension schema is, therefore, entirely feasible when it is intended to function as a packaging protocol alone, it may prove more problematic when it is intended to transfer objects between systems.

Use of an 'intermediary' schema

An alternative approach to the handling the complexity of CERIF was adopted by the UK's Readiness for REF (R4R) (Centre for e-Research 2011) project in the form of 'intermediary' XML schemas. These are heavily constrained schemas which cut down a complex and flexible model such as CERIF to a simplified form designed specifically for a given application. XSLT transformations are then used to generate the required CERIF files with consistent semantics and linkages (Gartner 2011). The R4R project created a bespoke schema of this kind, CERIF4REF, designed specifically to produce CERIF-compliant submissions to the UK's periodic research assessment exercise.

Adopting this approach would entail embedding or referencing an XML instance conforming to such an 'intermediary' schema within a METS <dmdSec> rather than the CERIF files themselves. A <behavior> element within the METS behavior section would then contain a <mechanism> element which would reference the associated XSLT needed to generate the raw CERIF files. The STRUCTID attribute of the <behavior> element would reference a <div> element in the structural map which in turn references an entry in the <fileSec> that points to the <dmdSec>

```

<dmdSec ID="respub1-0001-dmd-0002">
  <mdWrap MDTYPE="OTHER" OTHERMDTYPE="CERIF4REF">
    <xmlData>
      <c4r:cerif4Ref>
        <c4r:c4rREFData hesalID="9999">
          <c4r:c4rResearchGroups>
            <c4r:c4rResearchGroup ID="rg-a-12-999">
              <r4rResearchGroupName>Arts Research Group</r4rResearchGroupName>
              <c4r:c4rResearchGroupCode>A</c4r:c4rResearchGroupCode>
              <c4r:c4rUOA>12</c4r:c4rUOA>
            </c4r:c4rResearchGroup>
            ...
          </c4r:c4rREFData>
        </cerif4Ref>
      </xmlData>
    </mdWrap>
  </dmdSec>

<fileSec>
  <fileGrp ID="respub1-0001-dmd-filegrp002" USE="CERIF4REF">
    <file ID="respub1-0001-dmd-file001-002">
      <FLocat LOCTYPE="URL" xlink:href="#respub1-0001-dmd-0002"/>
    </file>
  </fileGrp>
</fileSec>

<structMap TYPE="physical" ID="respub1-0001-strmap002">
  <div TYPE="cerif4ref">
    <fptr FILEID="respub1-0001-dmd-filegrp002"/>
  </div>
</structMap>

<mets:behaviorSec LABEL="Generate CERIF from CERIF4REF">
  <mets:behavior ID="cerif4ref2cerif" STRUCTID="respub1-0001-strmap002">
    <mets:mechanism LOCTYPE="URL" href="xslt/cerif4ref2cerif.xsl"/>
  </mets:behavior>
</mets:behaviorSec>

```

The diagram illustrates the relationships between four XML sections: **dmdSec**, **fileSec**, **structMap**, and **behaviorSec**. Arrows indicate the following references:

- A line from the **fileGrp** element in **fileSec** points to the **dmdSec** element.
- A line from the **fptr** element in **structMap** points to the **fileGrp** element in **fileSec**.
- A line from the **STRUCTID** attribute of the **behavior** element in **behaviorSec** points to the **div** element in **structMap**.

containing the intermediary XML instance; this requires a recursive referencing of the `<dmdSec>` from the `href` attribute of the `<FLocat>` element, an approach similar to that previously advocated by Habing & Cole (2008) as a method for structuring OAI-ORE aggregation data within the METS framework.

Figure 1 illustrates this chain of references.

This approach should, if the intermediary schema is well designed, substantially reduce the complexity of the CERIF application and the consequent difficulties in designing suitable mechanisms for rendering its constituent components for delivery. The exchange of METS files containing contextual metadata encoded in this way should also be considerably easier, and require less documentation and interpretation by the recipient system.

Such an approach does, however, involve the use of a bespoke, non-standard, schema of limited use beyond the project for which it is designed. Although such schemas would almost certainly be simpler than CERIF itself, a proliferation of unstandardised schemas of this kind would nullify the benefits of employing a standard for encoding and exchanging research information, and probably generate consequent problems for future administration and delivery. It is therefore probably not a fully viable solution to employ these mechanisms beyond the

narrow, project-specific confines for which they were designed.

Serializing CERIF into the METS structural map

A third strategy for integrating CERIF into METS attempts to create a logical structural map which encodes the structure of the complex web of relationships embedded within the CERIF metadata itself. This is undoubtedly a non-standard use of the METS schema: the structural map is designed to encode relationships between the component data files that make up a digital object, not the metadata associated with it. Nonetheless, a case may be made for such complex and fragmented metadata as CERIF to be handled within METS in this way: the implications for the standard will be discussed after the methodology itself is explained.

Under this approach, a series of structural maps are constructed in which the internal hierarchy of <div> elements serializes the connections expressed in CERIF's linking files. A 'striped' syntax (as advocated by Habing & Cole (2008) for expressing aggregations) is here used to serialize the required relationships. A structural map <div> for a typical CERIF linkage (once more between an author and a research output) takes the following form:-

The linkage itself is expressed in three nesting `<div>` elements (marked by a in the diagram): the outermost delineates the first component (here the author identifier) of the link, using the method of an xpointer in the BEGIN attribute of an `<area>` element which references the element with the embedded CERIF file containing this information. The next level of the hierarchy indicates the semantic term used to characterise the link, again using the xpointer syntax to point to the element in CERIF's class definition file in which this term is declared. The lowest level of the hierarchy contains the object of the linkage, in this case details of the research publication.

Although the nesting itself is sufficient to designate the function of each `<div>` within this hierarchy, the TYPE attribute of `<div>` may be used as in this example to render the function more explicit by employing a controlled set of terms (here “link-subject”, “cfClassId” and “link-object”). Similarly, although not necessary, the LABEL attribute may be used as here to record the value of each referenced CERIF element, in order to render the structural map itself clearer to the reader.

Two further `<div>` elements, marked by b in the diagram, contain information on the temporal limits of this linkage which are an obligatory part of the CERIF model. Here the xpointers reference the link table itself, the only part of the CERIF data set in which this information is contained. It will be noticed that one further element in the linking table, `<cfClassSchemeId>`, the identifier for the classification scheme from which the semantic element `<cfClassId>` is taken, is not encoded in the `<div>` hierarchy: this, however, is unambiguously derivable from the CERIF class file as the following sibling of the `<cfClassId>` element referenced by the second `<div>` in hierarchy *a*.

A similar approach may also be used to handle such features as CERIF’s tables for encoding multiple-language versions of textual data (such as titles or abstracts of bibliographic works) as shown in figure 3 below. Encoding relationships of this type using “striped” `<div>` elements is relatively

simple:-

In this case only two levels of hierarchy are necessary to express the relationship between a component and its multi-lingual manifestations. Generating multiple language views of the CERIF object is then a simple matter of using XSLT transformations with a parameter for desired language codes to select those components with matching **cfLangCode** attributes.

Serializing CERIF relationships in this way does appear to combine the best parts of both of the two approaches discussed earlier. It employs the CERIF standard itself, rather than a bespoke intermediary, and also renders the

resolution of the linkages inherent in CERIF much easier to handle in a working environment, and easier to document (for instance, in a METS profile) than raw CERIF itself may prove. It is, therefore, likely to be easier to exchange METS files with CERIF encoded metadata using this technique than would be possible using the first two methods.

Using the structural map in this way may, however, legitimately be seen an inappropriate use of the METS schema as it is currently constructed. The structural map is not specifically designed as a mechanism for describing relationships within metadata, and it is necessary to employ something of a fudge (in the form of recursive referencing) to make it work here. It may also be argued with some validity that the structural map is not designed to allow the semantic modelling of the type envisaged here, and to do so severely stretches the METS model beyond its intended use.

The latter argument may reasonably be countered by arguing that the METS structural map has often been used for encoding logical as well as physical structures in which, it could be argued, some degree of semantic definition is inherent (even in such basic notions as 'chapter' or 'section' in a book, for example). It could be argued that the approach taken here defines logical structures in an analogous way, albeit one that extends the concept of a logical structural map beyond that originally envisaged.

The arguments against the use of the structural map to serialize metadata relationships are certainly more serious. This approach undoubtedly extends the METS' s functionality into areas for which it was not originally designed, but a case could be made, in applications such as this where the metadata is highly complex and fragmented, for allowing some of the powerful structural features METS provides for data to be extended to metadata. As alluded to earlier, METS presently takes a rather monolithic 'bucket'-like approach to metadata (METS Editorial Board 2010, p.11) each bucket of which can only be referenced as a whole (through the DMDID attribute of <div> for instance). In cases such as this, where more sophisticated referencing would be useful, the standard could usefully be extended to allow it to address individual metadata components as envisaged here. A change of this kind would obviate the need to take a non-standard approach of this kind.

Conclusions

Providing contextual information on the research environment surrounding a publication or output is becoming increasingly useful, and so employing a methodology which allows a key digital object metadata packaging scheme to be used with the only established research information standard currently available is a sensible option. The complexity of the CERIF standard, however, does present a number of problems for integrating it with METS.

Although it is possible to employ CERIF as a standard extension schema, or to employ an 'intermediary' schema instead, there is a case for serializing within METS itself some of the CERIF relationships which require decoding to make sense of the metadata expressed with it. At present METS does not readily allow this treatment of metadata, and so a non-standard methodology is required to allow it to be employed. There is a possible case for amending the METS standard to allow this type of metadata handling, although this would require extensive further investigation. Integrating CERIF into METS, however, is undoubtedly an important way to an integrated approach to digital library metadata to incorporate complex contextual information of this type without abandoning its core architectural principles.

References

Bolton, S. (2010). The Business Case for the Adoption of a UK Standard for Research Information Interchange. Joint Information Systems Committee. Retrieved from

<http://www.jisc.ac.uk/publications/reports/2010/businesscasefinalreport.aspx>

Centre for e-Research (2011) 'R4R: Readiness for REF' , *R4R: Readiness for REF*, [online] Available from: <http://r4r.cerch.kcl.ac.uk/> (Accessed 7 February 2011).

Consultative Committee for Space Data Systems. (2009). Reference Model for an

Open Archival Information System (OAIS): Draft Recommended Standard. Retrieved from <http://public.ccsds.org/sites/cwe/rids/Lists/CCSDS%206500P11/Attachments/650x0p11.pdf>

Digital Library Federation. (2010). <METS>: Metadata Encoding And Transmission Standard: Primer And Reference Manual. Digital Library Federation. Retrieved from <http://www.loc.gov/standards/mets/METSPRimerRevised.pdf>

European Organisation for International Research Information. (2010). CERIF Introduction. Retrieved February 8, 2011, from <http://www.eurocris.org/Index.php?page=CERIFintroduction&t=1>

Fedora Project. (2006). Encoding Fedora Objects in METS (Fedora Extension). Retrieved April 19, 2011, from <http://fedora-commons.org/download/2.1.1/userdocs/digitalobjects/rulesForMETS.html>

Gartner, R. (2008). *Metadata for digital libraries: state of the art and future directions*. JISC Technology and Standards Watch. JISC. Retrieved from http://www.jisc.ac.uk/media/documents/techwatch/tsw_0801pdf.pdf

Gartner, R. (2011) 'Intermediary schemas for complex XML publications: an example from research information management', *Journal of Digital Information*,

12(3), [online] Available from: <https://journals.tdl.org/jodi/article/view/2069>
(Accessed 7 July 2001).

Guenther, R., & Myrick, L. (2006). Archiving Web Sites for Preservation and Access: MODS, METS and MINERVA. *Journal of Archival Organization*, 4(1-2), 141-166.

Habing, T., & Cole, T. (2008). ORE Resource Map Implementation in METS: draft proposal. Retrieved April 1, 2011, from
http://ratri.grainger.uiuc.edu/oremets/1structMap_1REM.htm

Higher Education Funding Council for England. (2007). Data collection user guide: how to input, manage and submit data for the RAE. HEFCE. Retrieved from
<http://www.rae.ac.uk/pubs/2007/01/userguide/RAE%20User%20Guide.pdf>

Jeffery, K. G., Lopatenko, A., & Asserson, A. (2002). Comparative study of metadata for scientific information: the place of CERIF in CRISs and scientific repositories. *Gaining Insight from Research Information* (pp. 77-87). Presented at the Sixth International Conference on Current Research Information Systems, University of Kassel: Kassel University Press.

Library of Congress. (2010). Metadata Object Description Schema: MODS. Retrieved January 28, 2010, from <http://www.loc.gov/standards/mods/>

Library of Congress. (2011a). Metadata Encoding and Transmission Standard (METS) Official Web

Site. Retrieved March 4, 2011, from <http://www.loc.gov/standards/mets/>

Library of Congress. (2011b). PREMIS: Preservation Metadata Maintenance Activity (Library of Congress). Retrieved January 28, 2010, from <http://www.loc.gov/standards/premis/>

Massachusetts Institute of Technology. (2010). DSpaceMETSSIPProfile - DSpace - DuraSpace Wiki. Retrieved April 19, 2011, from <https://wiki.duraspace.org/display/DSPACE/DSpaceMETSSIPProfile>

Maslov, A. et al. (2010) 'Adding OAI-ORE Support to Repository Platforms' , *Journal of Digital*, 11(1), [online] Available from: <http://journals.tdl.org/jodi/article/view/749/640> (Accessed 31 May 2011).

METS Editorial Board. (2010). Reimagining METS: an exploration: draft, draft, draft for discussion at Fall 2010 DLF. Retrieved from <https://docs.google.com/fileview?id=0BzMhAsKzul-tN2Y0N2RjM2EtNzFmMC00NGQ3LTkyYjctNGZlZjEwODUzNmFm&hl=en&authkey=CIWlms8F>

Rogers, N., Huxley, L., & Ferguson, N. (2010). Exchanging Research Information in the UK. Retrieved from http://ie-repository.jisc.ac.uk/448/1/exri_final_v2.pdf

University of Southampton. (2010). Preservation Support – EPrints. Retrieved April 19, 2011, from http://wiki.eprints.org/w/Preservation_Support