

The International Journal of Digital Curation

Issue 2, Volume 4 | 2009

Modelling Organizational Preservation Goals to Guide Digital Preservation

Angela Dappert, Adam Farquhar,
The British Library,
Boston Spa, Wetherby, West Yorkshire, UK

Summary

This paper is an extended and updated version of the work reported at iPres 2008. Digital preservation activities can only succeed if they go beyond the technical properties of digital objects. They must consider the strategy, policy, goals, and constraints of the institution that undertakes them and take into account the cultural and institutional framework in which data, documents and records are preserved. Furthermore, because organizations differ in many ways, a one-size-fits-all approach cannot be appropriate. Fortunately, organizations involved in digital preservation have created documents describing their policies, strategies, work-flows, plans, and goals to provide guidance. They also have skilled staff who are aware of sometimes unwritten considerations. Within Planets (Farquhar & Hockx-Yu, [2007](#)), a four-year project co-funded by the European Union to address core digital preservation challenges, we have analyzed preservation guiding documents and interviewed staff from libraries, archives, and data centres that are actively engaged in digital preservation. This paper introduces a conceptual model for expressing the core concepts and requirements that appear in preservation guiding documents. It defines a specific vocabulary that institutions can reuse for expressing their own policies and strategies. In addition to providing a conceptual framework, the model and vocabulary support automated preservation planning tools through an XML representation¹.

¹ This article is based on the paper given by the authors at iPRES 2008; received September 2009, published October 2009.

The *International Journal of Digital Curation* is an international journal committed to scholarly excellence and dedicated to the advancement of digital curation across a wide range of sectors. ISSN: 1746-8256 The IJDC is published by UKOLN at the University of Bath and is a publication of the Digital Curation Centre.



Introduction

This article introduces a conceptual model and vocabulary for preservation guiding documents. Preservation guiding documents include documents, in a broad sense, which specify requirements that make the institution's values or constraints explicit and influence the preservation planning process. They may be policy, strategy, or business documents, applicable legislation, guidelines, rules, or even a choice of temporary runtime parameters. They may be oral representations as well as written representations in databases, source code, websites, etc.

The model and vocabulary can be shared and exchanged by software applications. They offer a starting point for creating individualized models for an institution. We show how they can be used to describe requirements for individual institutions, possibly, but not necessarily, in a machine-interpretable form. Furthermore, we show how these requirements can then be used in the context of comprehensive preservation planning.

To perform the analysis, the team used a combination of top-down and bottom-up methods. We examined the literature (e.g., American Library Association [ALA], 2007; Center for Research Libraries [CRL]/Online Computer Library Center [OCLC], 2007; Cornell University Library, 2003-2009; Electronic Research Preservation and Access Network [ERPANET], 2003; Joint Information Systems Committee [JISC], 2006; Solinet, n.d.) to create a top-down model from first principles. To complement this, we analyzed actual preservation guiding documents from archives, national libraries, and data centres for their content (e.g., National Archives of Australia, 2002; Florida, 2007; Digital Archives of Georgia, 2005; Hampshire Record Office, 2005; UK Data Archive [UKDA], 2008), and interviewed decision makers (Dappert, Ballaux, Mayr & van Bussel, 2008) to determine factors that influence their preservation choices. We extracted relevant concepts and vocabulary from the material to populate our model and compiled a list of example requirements. A more detailed [description of the approach](#) can be found in Dappert et al. (2008); a more detailed description of the conceptual model can be found in Dappert (2009). Aspects of the model draw from the Pronom conceptual model (R. Sharpe, 2009) and the OAIS model (Consultative Committee for Space Data Systems [CCSDS], 2002).

Context

The context of our conceptual model is the process of preservation planning for a digital collection. The goals of this process are to:

- identify which parts of the collection present the greatest risks.
- identify candidate preservation actions that could be taken to mitigate the risks.
- evaluate the candidate preservation actions to determine their potential costs and benefits. The cost includes the cost of executing the action, the cost of infrastructure required to sustain the results of the action, and the cost of significant characteristics lost in the action (e.g., loss of authenticity) etc. The benefits come from mitigating the risks and increase in proportion to the value of the object and the severity of the risk. The costs and benefits are not necessarily monetary.
- provide justified recommendations for actions to execute on parts of the collection.

All of these activities should be based on institutional requirements that extend beyond considering file formats and characteristics of individual digital objects to take into account the goals and limitations of the institution, features of its user community, and the environment in which its users access digital content.

The Core Conceptual Model

The core conceptual model consists of the components in Figure 1. In summary, any Preservation Object has one or more Environments. Every Environment in which the Preservation Object is embedded consists of one or more Sub-Environments, such as hardware and software environments, the legal system, and other internal and external factors. Preservation Objects and Environments are described by Characteristics, which are Property / Value pairs.

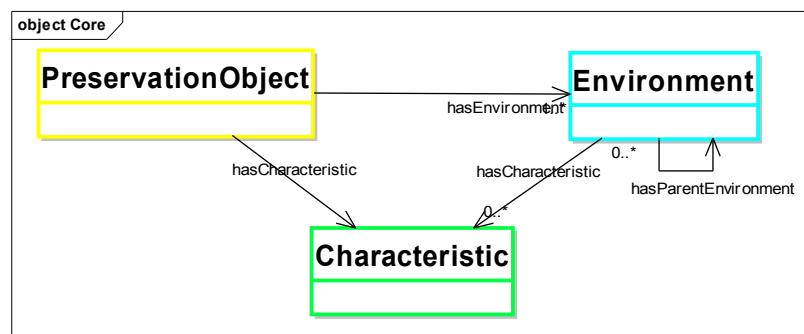


Figure 1. Core Conceptual Model.

We realized early that requirements express constraints on many levels of granularity.

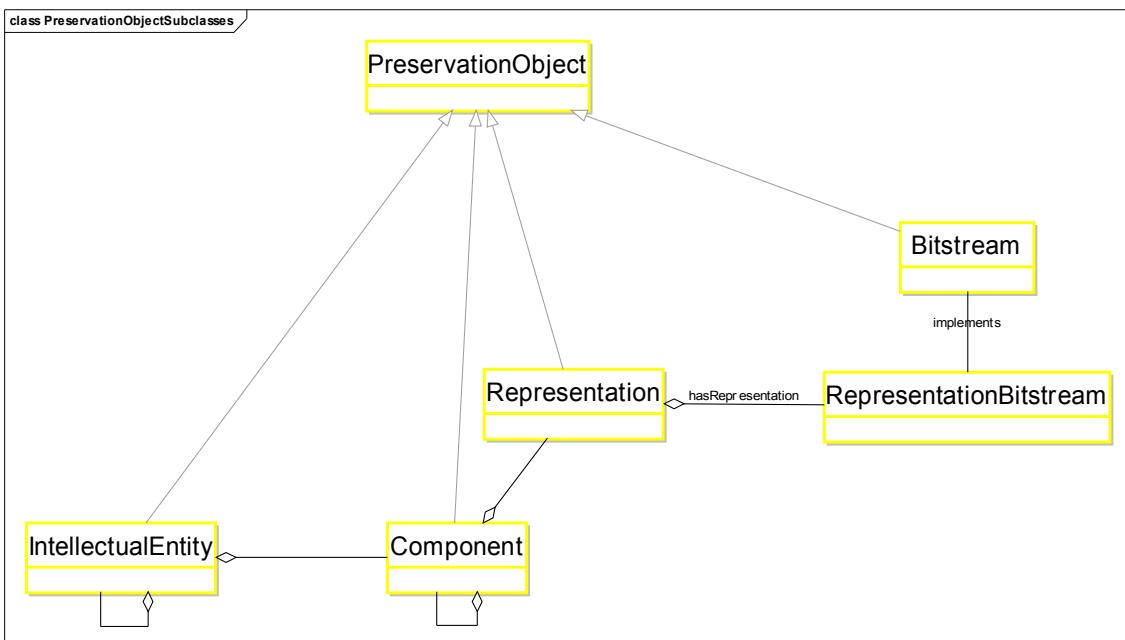


Figure 2. Preservation Object Subclasses.

We, therefore, define **Preservation Objects** as follows: A Preservation Object is any object that is directly or indirectly at risk and needs to be digitally preserved. Subclasses of **Preservation Object** are illustrated in Figure 2: Intellectual Entity, Component, Representation, Representation Bitstream, Bitstream.

Each Preservation Object subclass is related to another with the “containedIn” relationship (except that a Bitstream is contained in its Representation via its Representation Bitstream).

A **Bitstream** (and its subclasses, including Bytestreams and Files) is the primary physical Preservation Object. If it is at risk of decay or obsolescence, it becomes the object of preservation. We create and execute preservation plans to preserve it. A Bitstream is, however, embedded in a larger context.

A **Representation** is the set of all Representation Bitstreams that are needed to create one rendition of a logical digital object. A Bitstream is realised by its **Representation Bitstream**. Representation and Representation Bitstream are logical descriptions of physical Bitstreams.

Intellectual Entity and **Component** are logical objects. An Intellectual Entity is a distinct intellectual or artistic creation. The Intellectual Entity can be refined in ways to meet the needs of stakeholders. For example, in the library setting, common subclasses include Collection, Work, and Expression. Components are fine-grained parts of an Intellectual Entity, such as Table, Image, Title, Substring, that need to be described individually.

In the simplest case, a Bitstream, Representation Bitstream, and Representation have a one-to-one correspondence. For example, a book might be represented as a single PDF file in the PDF format. In other cases, however, several Bitstreams may be contained in one Representation Bitstream and several Representation Bitstreams may be contained in one Representation. For example, a book might be represented with one PDF file per chapter, each of which contains an embedded image for each of several pages.

Preservation activities that take place in the context of a content-holding institution, such as a library, involve considerations that go beyond an individual Representation. Consider the case of a library that has a substantial Collection:

- The overall Collection may be composed of smaller Collections. Some of these may be static for the institution, such as the Science Collection, or determined dynamically, such as the Collection of all articles that contain TIFF3.0 files. Collections may contain digital and non-digital objects.
- A Journal may belong to one or more collections. It is the logical object describing all Issues with the same title (setting aside some complexities involving name changes, etc.).
- An Issue is part of a Journal. It is the logical object containing all of the Articles published in a single Issue.
- An Article is part of an Issue. It is the abstract concept representing a distinct intellectual creation – the article.
- An Expression is the specific intellectual or artistic form that an Article takes when it is realised. There may be multiple Expressions of an Article and each Expression may have multiple Representations.
- A Representation is a set of Bitstreams that are required to render the Expression. There might be several Representations of an Expression of an Article (e.g., an HTML, a PDF, an XML, and a publisher-specific format).
- A Bitstream, such as a file, is part of a Representation (e.g., an MPEG-4 video File is part of an HTML Representation of an Article).

This model supports an essential property of preservation activity. Institutions

need to take a global view of their Collections and resources in order to coordinate preservation activity and take the appropriate actions. It is not enough to consider each digital object in isolation. This is the reason that the model goes well beyond the individual digital object.

Every Preservation Object has one or more **Environments**. Environment is defined as: A factor which constrains a Preservation Object and that is necessary to interpret it. Environments may fulfil different roles. For example, a Bitstream or a Representation may have creation, ingest, preservation, and access Environments; a Collection of Intellectual Entities may have an internal, a physical delivery, and an online delivery Environment.

Environments for Preservation Objects at a higher level must accommodate the requirements of Preservation Objects at a lower level. As long as a Bytestream is part of its Representation, it will live in the Representation's Environment. When it is taken out of the Representation's Environment, for example to be used in a migration, then the Bytestream's individual Environment requirements will influence the Environment of its new Representation.

It is worth noting that it may not be possible to derive the best Environment from a Bytestream's file format. If, for example, a Word file contains only text without formatting, headers and tables, and so forth, then a TXT output might be considered perfectly adequate, even though this would in general not be considered an ideal migration format for a Word file.

Every Environment consists of a number of Sub-Environments. The subclasses of Environment (see Figure 3) include software, hardware, and community, including legal or budgetary restrictions. See the full report (Dappert, 2009) for additional Environment subclasses that have been extracted from preservation guiding documents. Policy Factors, in particular, are discussed in depth.

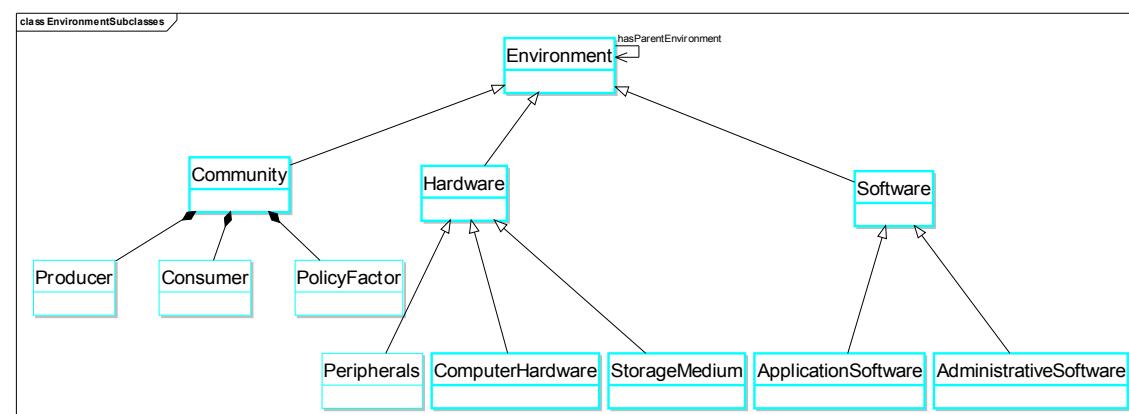


Figure 3. Environment Subclasses.

There is a close relationship between an Environment and an extended notion of Representation Information as it is defined in OAIS (CCSDS, 2002). Other examples of extended notions of Representation Information are discussed in Brown (2008).

Characteristics describe the state of Preservation Objects and Environments as

Property / Value pairs. Values may be stored directly as object values, referenced indirectly through registries or inventories, or extracted dynamically through characterisation processes. Some vocabulary for Properties can be found in the full report (Dappert, 2009).

The Full Conceptual Model

The full conceptual model extends the core model with concepts from the preservation planning domain: Preservation Risk, Preservation Action and Requirement, as illustrated in Figure 4.

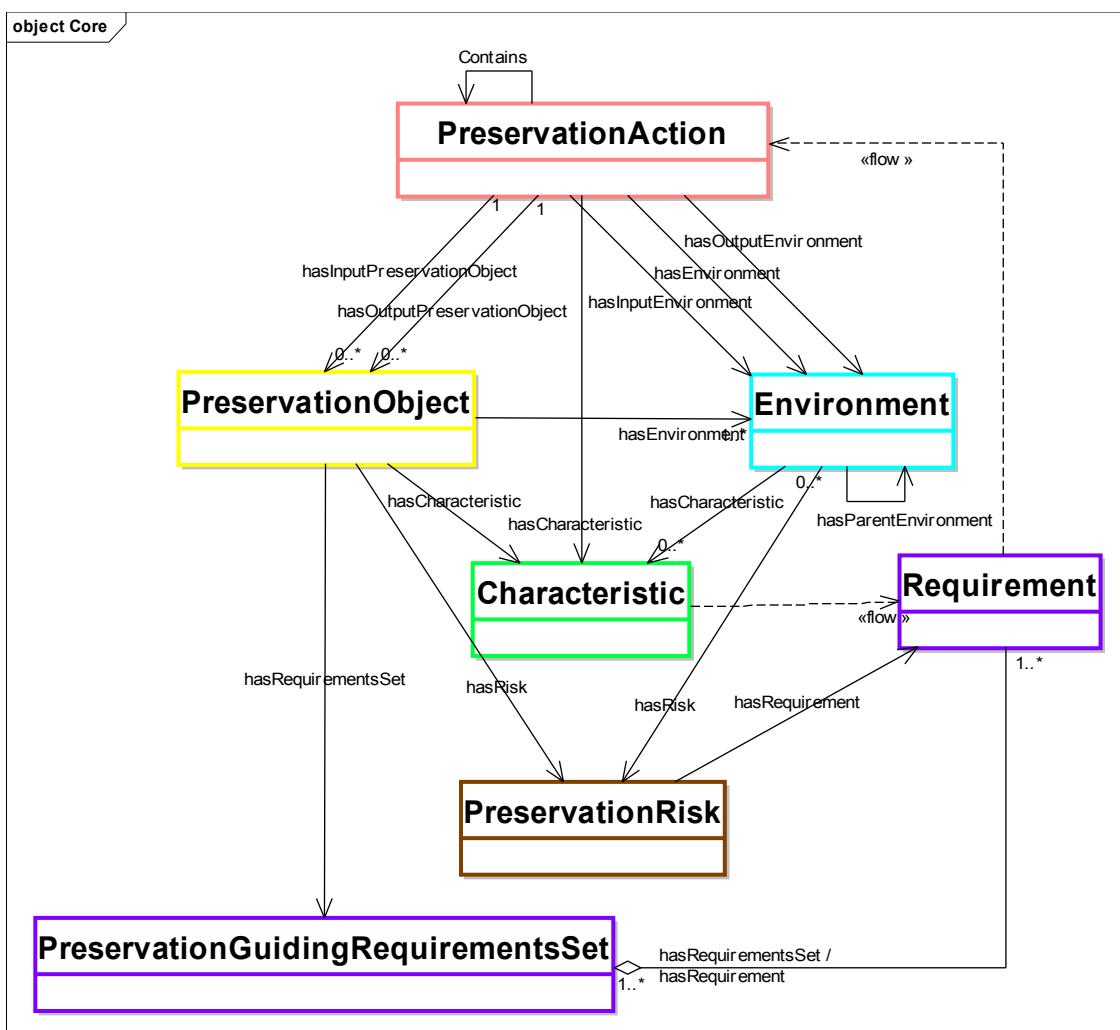


Figure 4. Full Conceptual Model.

Degradation of Preservation Objects is caused by two things:

- Preservation Risks
- Executing imperfect, lossy Preservation Actions

Acceptable levels of degradation are defined in an institution's Requirements, which specify permissible or desirable Characteristics of Preservation Objects and Environments. They make the institution's values explicit, influence the preservation process, and are captured in Preservation Guiding Documents.

Changes to a Preservation Object or Environment, such as obsolescence of hardware or software components, decay of data carriers, or changes to the legal

framework may introduce **Preservation Risks**. An individual institution's Preservation Risks are specified in **Risk Specifying Requirements**. Whenever Characteristics of a Preservation Object or its Environments violate the constraints which are specified in the Requirement, then the Preservation Object is considered at risk. Once a Risk Specifying Requirement is violated, a preservation monitoring process should notice this and trigger the preservation planning process. It, in turn, determines the best Preservation Action to mitigate this risk. **Preservation Object Selecting Requirements** are a subclass of Risk Specifying Requirements which specify which subset of Preservation Objects is at risk.

A composite **Preservation Action** may consist of elementary Preservation Actions and may include conditional branches and other control-flow constructs. When a Preservation Action is applied to a Preservation Object and its Environment, it produces a new Preservation Object and/or a new Environment in which the Preservation Risk has been mitigated. Every Preservation Action, therefore, has not only an Input Preservation Object and (at least one) Input Environment, but also an Output Preservation Object and Output Environment. For example, if a Microsoft Word File is migrated to a PDF File, this results in a new Preservation Object, which has different Characteristics, but also a new Environment in which it can be used – in this case the platform needs at least to contain a PDF viewer. This approach works for migration, emulation, hardware and other solutions.

For any given Preservation Object and its Environment, there may be multiple possible Preservation Actions to mitigate a Preservation Risk. Which of these Preservation Actions is the most suitable for the Preservation Object can be derived from the information in the **Requirements**. In order to determine whether an abstract Requirement is applicable and satisfied, one needs to evaluate the concrete Values of the Characteristics of Preservation Objects and their Environments or the concrete Values of a candidate Preservation Action at a given time.

Some Requirements can be expressed in a machine-interpretable way. They refer solely to concepts and vocabulary contained in the model. They may provide a conditional context, pre- and post-conditions, and sometimes complex expressions. In addition, it is useful to specify the relative importance and acceptable tolerances for requirements. Importance factors specify the importance of a requirement for an institution. A tolerance threshold specifies the degree to which deviation from the Requirement can be accepted.

Requirement Subclasses

During our literature and document analysis, we extracted Requirements that we categorized into the subclasses depicted in Figure 5. Risk Specifying Requirements were discussed earlier.

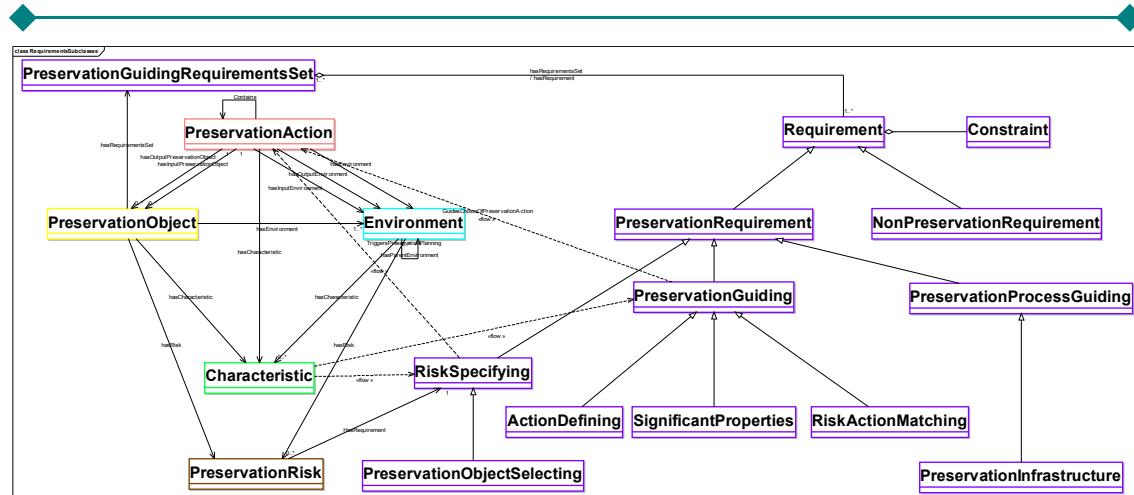


Figure 5. Requirement Subclasses.

Preservation Guiding Requirements specify which kinds of Preservation Actions are desirable for the Preservation Object and its Environments. For example: The size of the Preservation Action's output Preservation Object should not exceed a maximal size as set by the institution. They are dependent on

- which input Characteristics need to be met to consider the Preservation Action.
- which output Characteristics are permissible or desirable (either in absolute terms or in relationship to Characteristics of the input Preservation Object, which might be a derivative or the original submitted to the institution).
- which Characteristics of the Preservation Action itself are desirable.

Action Defining Requirements (subclass of Preservation Guiding Requirement) define which kinds of Preservation Actions are desirable independent of the Characteristics of the Preservation Object and its Environments, but dependent only on the Characteristics of the Preservation Action itself. For example PDF may, for a given institution, not be an acceptable preservation output format of a Preservation Action (independent of any input Characteristics of Preservation Objects and Environments).

Significant Characteristics (subclass of Preservation Guiding Requirement). Our definition of significant characteristics is close to the one expressed by Andrew Wilson, National Archives of Australia: “the Characteristics of digital objects that must be preserved over time in order to ensure the continued accessibility, usability, and meaning of the objects, and their capacity to be accepted as evidence of what they purport to record.” (Wilson, 2007) We, however, treat them as Requirements rather than Characteristics. Significant Characteristics are often limited to Characteristics of Bitstreams or Representations for which it is possible to evaluate Values automatically. We consider Significant Characteristics for any Preservation Object or Environment subclass. See (Dappert & Farquhar, 2009) for a more complete discussion.

Preservation Process Guiding Requirements (subclass of Preservation Requirement) describe the preservation process itself independent of the Characteristics of the Preservation Object or the Preservation Actions. For example: A preservation planning process should be executed for every digital object at least every 5 years, independent of the Preservation Risks that are established for this digital object. These requirements do not guide the preservation planning process.

Preservation Infrastructure Requirements (subclass of Preservation Process Guiding Requirement) are particularly prominent in preservation guiding documents. They specify required infrastructure Characteristics with respect to security, networking, connectivity, storage, etc.. For example: Mirror versions of on-site systems must be provided.

Non-Preservation Requirements (subclass of Requirement) specify processes relevant to preservation, but not part of preservation itself. For example: The Preservation Action must produce metadata that is needed by the electronic resource management system.

Risk / Action Matching Requirements (subclass of Preservation Guiding Requirement) specify that a candidate Preservation Action has to be an appropriate match to a given Preservation Risk. They are rarely stated explicitly in preservation guiding documents.

Preservation Risk subclasses include (see Figure 6):

- NewVersion: A new version of the Preservation Object or Environment is available. This creates a risk of future obsolescence, or a risk of having to support too many versions.
- LackingSupport: The Preservation Object or Environment is no longer sufficiently supported. This creates a risk that support will cease altogether, rendering the Preservation Object or Environment inaccessible.
- DeteriorationOrLoss: The Preservation Object or Environment is deteriorating or has been lost. Reconstruction or replacement become necessary.
- Proprietary: The Preservation Object or Environment is proprietary. There is a risk that it cannot be replaced since the specifications for it are unknown.
- UnmanagedGrowth: The institution's Preservation Objects or Environments are becoming too diverse to manage. A normalization Preservation Action is needed to simplify or unify them.

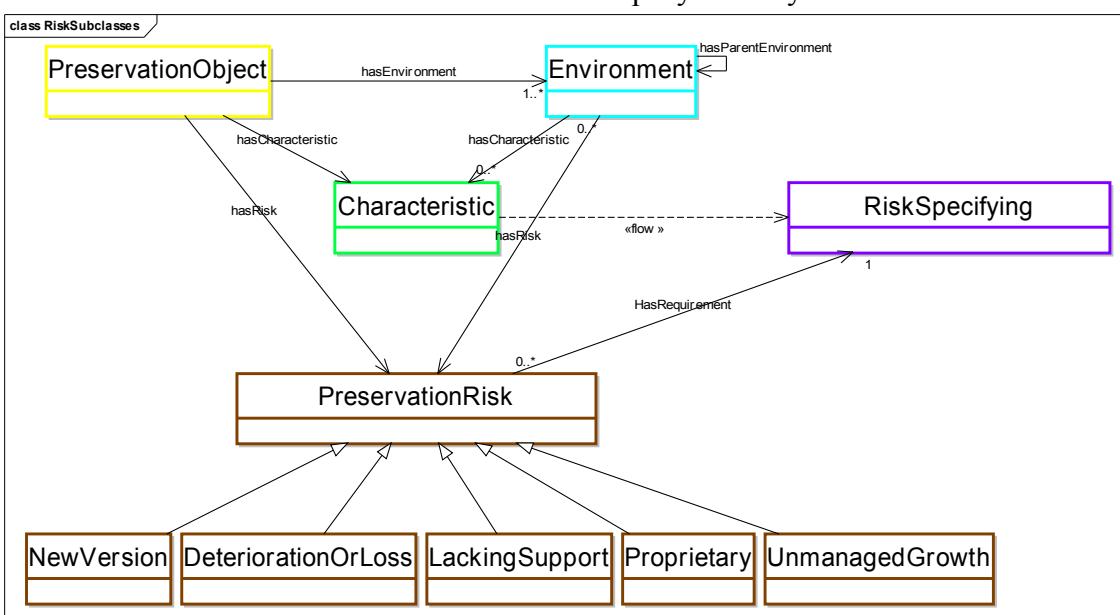


Figure 6. Risk Subclasses.

For every Preservation Risk, Preservation Object, and Environment there is an appropriate Preservation Action to mitigate the risk. For example, the risk of data carrier failure can be mitigated by a carrier refresh. The risk of file format obsolescence can be mitigated by migrating objects to an alternative format. The three key subclasses of Preservation Action are Replacement, Repair and Reconstruction (see Figure 7).

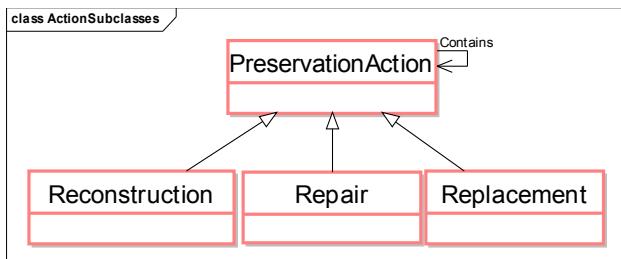


Figure 7. Action Subclasses.

The diagram (Figure 8) and table (Table 1) illustrate refined Preservation Action subclasses depending on the subclasses of the Preservation Risk and the affected Preservation Object or Environment.

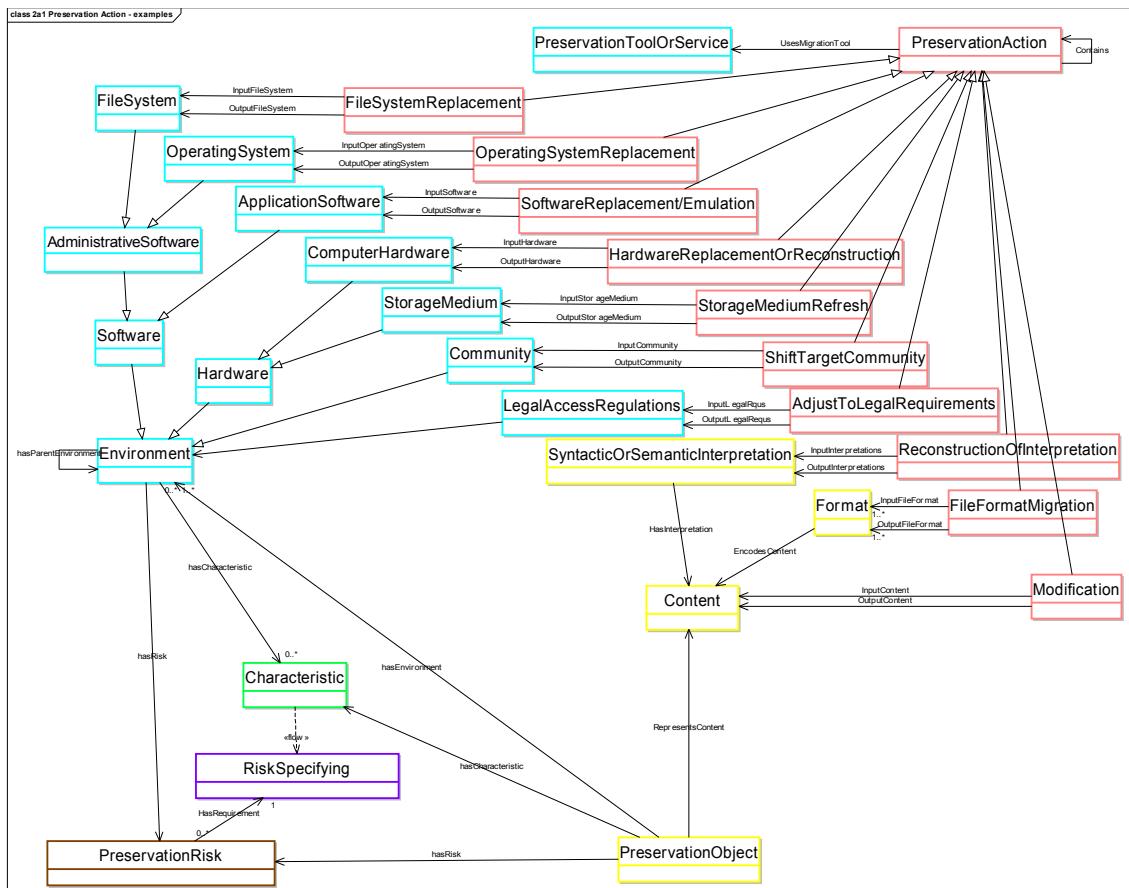


Figure 8. Preservation Action Subclasses Depend on Preservation Objects or Environments and Risk

Most of them are self-explanatory. Some deserve some comment:

- Modification of Content might represent an action such as the reconstruction of a deteriorated file, or a file that is modified in order to satisfy new legal requirements.

- One possible Preservation Action is not to do anything (“wait and see”).
- Migration does not always imply that a different file format is chosen. One might, for example replace an XML file with another XML file. In that case the input and output file formats happen to be the same. The output Preservation Object might nonetheless have different Characteristics to the input Preservation Object because of the different information captured within the XML tags.
- The needs of the target community might be a deciding factor for the choice of PreservationActions, and, conversely, the choice of PreservationActions will shape and change the community, just as it changes the other Environment subclasses.
- Shifting the target community might be a somewhat unintuitive PreservationAction, which is parallel to all other forms of Environment replacement. An example might be turning a research data collection into a history-of-science repository, as the material contained in the collection ceases to live up to contemporary standards of scientific use.
- Community consists of producers and consumers. Both types are either technical (e.g., repository or IT staff, publishing staff) or content-oriented (authors or readers) and will consider the digital object obsolete under different circumstances and according to their needs.

Example Risks	Preservation Object subclass	Environment subclass	Preservation Risk subclass	Preservation Action subclass
Data carriers deteriorate and cannot be read		Storage Medium	Deterioration	Replacement
The digital object becomes corrupted on the carrier and the original byte stream cannot be retrieved.	Bitstream		Deterioration	Reconstruction
Essential hardware components are no longer supported or available		Hardware	Lacking support	Replacement
Software components are proprietary and the dependence is unacceptable to the institution.		Software	Proprietary	Replacement
The community requires new patterns of access, such as access on a mobile phone, rather than a workstation		Hardware and Software	Obsolete	Replacement
File formats become obsolete.	File		Obsolete	Replacement
The legislative framework changes and the data or access to it has to be adapted to the new regulations		Legislation	New Version	Replacement

Table 1. Examples of Refined Preservation Action Subclasses.

Use to Model Institutional Requirements

The diagram in Figure 9 gives an overview of how the model described in this report can be used to create an institutional preservation guiding document. It introduces the General Model that consists of the concepts and vocabulary that are described in this paper, and the Instantiated Model that an institution might create to reflect its individual state and requirements.

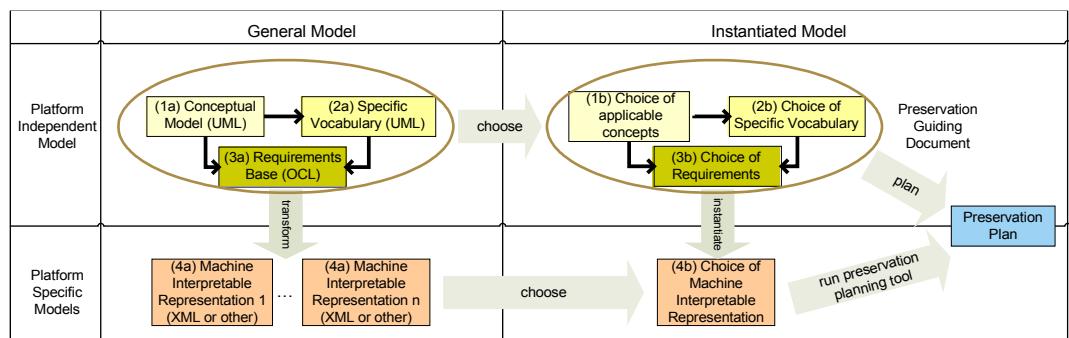


Figure 9. Modelling institutional requirements.

The numbering in the text refers to components in the diagram. Numbering including the letter “a” describes components in the general model. Numbering including the letter “b” describes components in an instantiated model.

(1a) The conceptual model, as discussed in this paper, defines the basic concepts that are needed in the domain of organizational preservation guiding documents and the relationships between them. They comprise Preservation Objects, Environments, Characteristics, Preservation Actions, Risks and Requirements.

(2a) The specific vocabulary defines

- subclasses of the basic concepts,
- properties of the basic concepts and their subclasses,
- allowable values for these properties.

(3a) The requirements base describes sets of organizational requirements which may be contained in preservation guiding documents. They are expressed solely in terms of the concepts and attributes of the conceptual model and its specific vocabulary. They may be parameterized so that they can be instantiated for a specific institution’s conditions.

(4a) The elements in the conceptual model, the specific vocabulary, and the requirements base can be translated into several implementation-specific machine-interpretable representations, for example, based on an XML schema.

(1b) The institution chooses which of these concepts are supported in its setting and are needed by its preservation planning service. Since the conceptual model is very concise, in most cases all of the concepts would be used.

(2b) The institution chooses which specific vocabulary applies to it. The institution also assigns values to the Characteristics of its Preservation Objects and Environments if these values are not to be measured automatically, or otherwise specifies the method of obtaining measurements or derivations. It will, for example, need registries of tools, formats, and legislative requirements, and need inventories of its collections, software licenses and staff members.

(3b) The institution chooses which Requirements in the Requirements base apply and instantiates them, so that they are now un-parameterized. It specifies importance factors, operators, and tolerances.

The outputs of steps (1b), (2b) and (3b) form the core part of a preservation guiding document.

(4b) From the choices of steps (1b), (2b), (3b), and the choice of machine-interpretable language, an instantiated machine-interpretable description of the institutional Requirements is derived. This serves as a basis for automated preservation planning. Many requirements in preservation guiding documents, especially on higher institutional levels, may not be machine-interpretable, but it can still be useful to represent the machine-interpretable subset for automatic evaluation.

The planning tool now matches the Requirements in the machine-interpretable version of the preservation guiding document (4b) against the state of the institution. It can then identify which Preservation Actions can best satisfy the Requirements under the given state.

Use to Perform Comprehensive Preservation Services

This model is well-suited for describing any Preservation Object subclass and a wide range of preservation services (e.g., monitoring, planning, characterisation).

For example, characterisation tools are defined to work on the Representation and Bitstream level. But there are also tools that characterise on a higher level, such as collection profiling tools which analyse Characteristics of a Collection at a given time and produce profiles describing the Collection. They could in principle share the conceptual model and associated processes.

As a further example, preservation planning needs to compare the Characteristics of a Preservation Object and its Environments before and after the execution of a candidate Preservation Action in order to evaluate the action against an institution's Requirements. The result is an evaluation score for how suitable each candidate Preservation Action is with respect to the Institution's Requirements. The utility analysis of Plato (Becker, Kulovits, Rauber & Hofman, [2008](#)) is an example of this.

Preservation Requirements express constraints on all levels of Preservation Objects in the Preservation Object hierarchy (e.g., budgetary constraints on the Collection level; preserving interactivity at the Representation level) and might even mix Characteristics from several levels (e.g., specifying constraints on Collections which contain Bitstreams with a certain Characteristic).

Since each possible Preservation Action may affect multiple levels in the Preservation Object hierarchy, the evaluation of a Preservation Action must be determined on all levels. That is, for every candidate Action, we can evaluate how well it satisfies the Requirements associated with a specific Bitstream, as well as how well it satisfies the Requirements for the whole of its Representation, Component, or even Intellectual Entity.

If for example, a concrete Preservation Action exceeds the Institution's budget, then it need not be considered for a given Bitstream. Equally, if it violates a Collection principle, even though it would be very suitable for preserving a specific Representation, it need not be considered. This sort of higher-level constraint is very useful to rule out unsuitable candidate Preservation Actions at a lower level.

Conversely, it is necessary not just to evaluate a concrete Preservation Action's utility in isolation on a lower level, but rather place it in a higher level context. When combining the evaluations from lower levels, with constraints on the higher level, then the evaluation of a Preservation Action might shift in the more global perspective. Planning algorithms need to take this into account.

For example: Consider the case in which an organization has decided to migrate a PNG to a GIF file. When we look at the enclosing Web Page Intellectual Entity, we see that the references to the image are broken and that the best Action would now add the Preservation Action "rename the links". When we look at the next higher Website Intellectual Entity, we see that they use java script for their links. The renamed links would not work. The best option now is to use a redirect list for the web server to the image on the server side.

Conclusion

This article introduced a conceptual model and vocabulary for preservation guiding documents. We showed how the model and vocabulary can be used to model requirements for individual institutions, possibly in a machine-interpretable form, and how these requirements can then be used to perform *comprehensive* preservation services that:

- accommodate a full range of preservation services such as monitoring, characterisation, comparison of characteristics, and evaluation of candidate preservation actions.
- allow processes to be associated with a full range of entities from institutions, and collections, down to bitstreams and atomic logical components of digital objects.
- consider technical as well as organizational properties.
- accommodate all types of preservation actions, from software actions (e.g., migration, emulation, file repair), hardware-related actions (e.g., data carrier replacement or hardware replacement / reconstruction / repair), to organizational actions (e.g., adapt processes to new legislation, adapt to new requirements of the designated community).

The conceptual model presents a simple but expressive representation of the preservation planning domain. The model and vocabulary can be shared and exchanged by software applications. They offer a convenient starting point for creating individualized models for an institution; this holds true even if the institution does not require a machine-interpretable specification. The model views preservation planning as a process that identifies and mitigates risks to current and future access to digital objects.

This paper is an extended version of the work reported at iPres 2008. It is informed by analysis undertaken in the Planets project. It will be partially implemented during the project, and also serves as a basis for further development and implementation.

Acknowledgements

The authors would like to thank colleagues on the Planets Project, in particular Bart Ballaux, Michaela Mayr, Rob Sharpe and Sara van Bussel for their contributions to this work.

Work presented in this paper is partially supported by the European Community under the Information Society Technologies (IST) Programme of the 6th FP for RTD - Project IST-033789. The authors are solely responsible for the content of this paper.

References

- American Library Association. (2007). Definitions of digital preservation. *ALA Annual Conference, Washington, D.C.*, June 24, 2007. Retrieved from <http://www.ala.org/ala/mgrps/divs/accts/resources/preserv/defdigpres0408.pdf>
- Becker, C., Kulovits, H., Rauber, A., & Hofman, H. (2008). Plato: A service oriented decision support system for preservation planning. *Joint Conference on Digital Libraries, June 16-20, 2008* - Pittsburgh, Pennsylvania.
- Brown, A. (2008). *White paper: Representation information registries*. PLANETS report PC/3-D7, January 29, 2008. Planets. Retrieved September 22, 2009, from http://www.planets-project.eu/docs/reports/Planets_PC3-D7_RepInformationRegistries.pdf
- Consultative Committee for Space Data Systems. (2002). *Reference Model for an open archival information system (OAIS)*, from *Blue Book* (the full ISO standard), January, 2002. Retrieved September 22, 2009, from <http://public.ccsds.org/publications/archive/650x0b1.pdf>
- Cornell University Library. (2003-2009). *Digital preservation management: Implementing short-term strategies for long-term problems*. Tutorial. Retrieved September 22, 2009, from <http://www.library.cornell.edu/iris/tutorial/dpm/terminology/strategies.html>
- Center for Research Libraries/Online Computer Library Center. (2007). *Trustworthy repositories audit & certification (TRAC): Criteria and checklist*. Version 1.0, February 2007. Retrieved September 22, 2009, from <http://www.crl.edu/PDF/trac.pdf>
- Dappert, A., Ballaux, B., Mayr, M., & van Bussel, S. (2008). *Report on policy and strategy models for libraries, archives and data centres*. PLANETS report PP2-D2. June 24, 2008. Planets. Retrieved September 22, 2009, from http://www.planets-project.eu/docs/reports/Planets_PP2_D2_ReportOnPolicyAndStrategyModelsM24_Ext.pdf
- Dappert, A. (2009). *Report on the conceptual aspects of preservation, based on policy and strategy models for libraries, archives and data centres*. PLANETS report PP2-D3. July 9, 2009. Planets. Retrieved October 12, 2009, from http://www.planets-project.eu/docs/reports/Planets_PP2_D3_ReportOnPolicyAndStrategyModelsM36_Ext.pdf
- Dappert, A., & Farquhar, A. (2009). Significance is in the eye of the stakeholder. *European Conference on Digital Libraries (ECDL)*, September 27 – October 2, 2009. Springer LNCS 5714.

Digital Archives of Georgia. (2005). *Digital Preservation Policy*. Retrieved September 22, 2009, from
http://www.sos.ga.gov/archives/who_are_we/rims/digital_History/policies/policy%20-%20Digital%20Preservation%20Policy.pdf

Electronic Research Preservation and Access Network. (2003). *Digital policy preservation tool*. ERPLANET guidance document, September, 2003. Retrieved September 22, 2009, from
<http://www.erpanet.org/guidance/docs/ERPLANETPolicyTool.pdf>

Farquhar, A., & Hockx-Yu, H. (2007). Planets: Integrated services for digital preservation. *International Journal of Digital Curation* 2,(2), p. 88–99.

Florida Digital Archive (2007). *Florida Digital Archive (FDA) Policy guide*. Version 2.4, August 2007. Retrieved September 23, 2009, from
<http://www.fcla.edu/digitalArchive/pdfs/DigitalArchivePolicyGuide.pdf>

Hampshire Record Office. (2005). *Digital preservation policy*. Retrieved September 24, 2009, from <http://www3.hants.gov.uk/archives/hro-policies/hro-digital-preservation-policy.htm>

Joint Information Systems Committee. (2006). *Continued access to authentic digital assets*. Digital Preservation and Records Management Programme Digital Preservation briefing paper, November, 2006. Retrieved September 22, 2009, from
http://www.jisc.ac.uk/publications/publications/pub_digipreservationbp.aspx

National Archives of Australia. (2002). *An approach to the preservation of digital records*. December, 2002. Canberra, Australia. Retrieved September 22, 2009, from http://www.naa.gov.au/Images/An-approach-Green-Paper_tcm2-888.pdf

Sharpe, R. (2009) PLANETS Data Model Overview, Planets Report IF8-D1, please request from info@planets-project.eu

Solinet (nd). *Contents of a digital preservation policy*. Retrieved September 24, 2009, from <http://www.lyrasis.org/Preservation/Resources%20and%20Publications/Contents%20of%20a%20Digital%20Preservation%20Policy.aspx>

UK Data Archive. (2008). *UK Data Archive preservation policy*. March, 2008. Retrieved September 22, 2009, from <http://www.data-archive.ac.uk/news/publications/UKDAPreservationPolicy0308.pdf>

Wilson, A. (2007). *Significant properties report*, InSPECT Work Package 2.2, Draft/Version 2. Retrieved October 12, 2009, from
http://www.significantproperties.org.uk/documents/wp22_significant_properties.pdf