A Model for Digital Preservation Repository Risk Relationships

Andrew McHugh

Abstract

The paper introduces the Preserved Object and Repository Risk Ontology (PORRO), a model that relates preservation functionality with associated risks and opportunities for their mitigation. Building on work undertaken in a range of EU and UK funded research projects (including the Digital Curation Centre[1], DigitalPreservationEurope[2] and DELOS[3]), this ontology illustrates relationships between fundamental digital library goals and their parameters; associated rights and responsibilities; practical activities and resources involved in their accomplishment; and risks facing digital libraries and their collections. Its purpose is to facilitate a comprehensive understanding of risk causality and to illustrate opportunities for mitigation and avoidance.

The ontology reflects evidence accumulated from a series of institutional audits and evaluations, including a specific subset of digital libraries in the DELOS project which led to the definition of a digital library preservation risk profile. Its applicability is intended to be widespread, and its coverage expected to evolve to reflect developments within the community.

Attendees will gain an understanding of the model and learn how they can utilize this online resource to inform their own risk management activities.

Introduction

Considerable work has been done in recent years to explore the extent to which digital preservation resources and processes can be considered by stakeholders to be *trustworthy*. Results of this work have included the publication of a number of standards and *de facto* standards, the establishment of accreditation schemes[4] and an increasing community sense of baseline requirements for successfully tackling preservation of digital materials. With the emergence of wider preservation requirements (such as UK funder-issued research data management policies) means to understand and express readiness have never been more sought-after. Risk and risk management have been suggested as useful metrics to quantify capacity and capability (McHugh et al 2008) and efforts to date suggest that these are intuitive concepts for those with curatorial responsibility. Despite increasing awareness of best practice and associated expectations, this knowledge has been either generically expressed, or limited to individuals' or specific practitioner groups. The *Preserved Object and Repository Risk Ontology (PORRO)* is a means for presenting a coherent mapping of critical

[1] See http://www.dcc.ac.uk
[2] See http://www.digitalpreservationeurope.eu
[3] See http://www.delos.info
[4] See DANS Data Seal of Approval, http://www.datasealofapproval.org/

preservation risks to the range of resources, activities and drivers that characterize organisations' preservation efforts.

Foundation Work

Trustworthy Repositories Audit and Certification Criteria and Checklist

Around a decade ago the seminal *Trusted Digital Repositories: Attributes and Responsibilities* called for appropriate audit and certification metrics to support validation of preservation repositories' infrastructures and processes (Dale et al, 2002). This demand was realised as *the Trustworthy Repositories Audit and Certification Criteria and Checklist (TRAC:CC)*, authored by a range of international experts and maintained at various times by *RLG/OCLC*, *NARA* and the *Center for Research Libraries* (CRL-OCLC, 2007). The document presented a set of criteria that should be demonstrable by repositories aspiring to a trustworthy status, and was expected to provide an intellectual foundation for repository certification. Criteria are divided into issues associated with organisation, technology and digital preservation process. It has subsequently been used as the basis of a *Consultative Committee for Space Data Systems* document approved for publication by the ISO technical committee (TC 20/SC 13) as ISO 16363 (CCSDS, 2011) and an accompanying standard, ISO 16919 describing requirements for auditing organisations (CCSDS, 2011 [2]). The documents have provided platforms for several pilot audit programmes, including those undertaken by CRL, the *Digital Curation Centre*, *DELOS* and *DigitalPreservationEurope* projects and most recently the *MOIMS-RAC Digital Repository Audit and Certification Working Group*[5]. Notable feedback from many of these pilot studies is that although useful, there are often difficulties in demonstrating conformity with criteria that remain somewhat opaque (necessarily to ensure widespread applicability to a heterogeneous range of repository contexts). Although the standard's most recent accompanies each criterion with examples of how repositories can demonstrate that they meet associated requirements, as well as further discussion, doubt continues to persist. Furthermore, there is little interconnectivity made explicit between individual criteria, but for the appearance of common evidential requirements (such as similar policy requirements for multiple criteria).

Digital Repository Audit Method Based on Risk Assessment

In 2008 the *Digital Curation Centre* and *DigitalPreservationEurope* projects released a repository risk management methodology and online tool called the *Digital Repository Audit Method Based on Risk Assessment (DRAMBORA)* (McHugh et al, 2007). Its development was the culmination of a period of sustained investigation into the issues associated with repository assessment. Its core workflow was informed by six preliminary repository audits whereby a process was developed upon existing draft certification standards. This period lead to the iterative prototyping of paper-based and online tools, which were validated in a further series of supported assessments, which included a selection of digital libraries (Ross et al, 2008).

The process of tool development and deployment has revealed considerable insights into the practical demands associated with the maintenance of a successful digital preservation service. Core preservation objectives (broadly implicit in existing standard documentation), typical rights and responsibilities, facets of policy, required resources, activities undertaken, and finally associated

[5] See http://cwe.ccsds.org/moims/default.aspx

risks are all better understood as an outcome of these efforts. However, the audits were undertaken by a comparatively small team, and the knowledge and experience accrued was for the most part limited to these individuals. Audit participants typically commended the process but explicitly noted that the expert insights available from auditors were integral and irreplaceable requirements to ensure the process's value. The auditor's role was to provide clarity to the question of *how* to satisfy trustworthiness criteria, how to identify and improve existing shortcomings. Fundamentally, to provide mapping between best practice requirements, the realities of existing preservation infrastructure, and what must be done in order to shore up evident gaps.

Process of Ontology Development

The Value of Relating Risk

When the *DCC*, *DPE* and *Delos* reflected on their work in the area of repository evaluation, conclusions were somewhat mixed. There is demonstrable value in a tool that streamlines the process of organisational self-awareness and risk management, but self-assessment remains fundamentally limited by the extent to which individuals are aware of what is wrong and what may be done to improve. Best practice standards are good at describing an optimal preservation landscape, but provide fewer insights in terms of how to get there. Furthermore, the constraints of a top-down, prescriptive model can be uncomfortably limiting to a domain that is hugely diverse. Digital preservation is an issue that spans disciplines, data, cultures, technology and legal jurisdictions. Likewise it is characterised by a temporal dimension that implies further change over time. Exacerbating this, digital preservation is a complex challenge with many contributing factors that span individuals, departments and even organisations. Criteria lists are by their nature somewhat ineffective in illustrating the interconnectedness of infrastructural facets that can increase or limit risk exposure in various ways.

Better means for risk definition and understanding prompt greater awareness of influential factors, their consequences and appropriate responses. Work in the related domain of internet security has revealed ontologies' value in the identification and classification of attacks and threats to networked systems, in terms of their relationships with technology, policy and use (Ahmed et al, 2007; Ekelhart et al, 2007; Fenz & Neubayer, 2009; Raskin et al, 2001; Tsoumas & Gritzalis, 2006). A more holistic view of risk and its cause and effects seems well suited to the complex environments within which information is preserved, and its availability threatened. Recent work has explored conceptual modelling of preservation goals with risk in mind (Dappert & Farquhar, 2009), and the role of risk management in designing preservation solutions (Barateiro et al, 2010). Other efforts have sought to form relationships between heterogeneous metadata. The P2 registry (Tarrant et al, 2009) uses the semantic web to link data from Pronom[6] to support rudimentary risk assessment based on file format characteristics. Ontologies were used in the PANIC project's prototypical preservation alert and response system (Hunter & Choudhury, 2004; Hunter & Choudhury, 2005). More general documentation projects such as PREMIS OWL, and the CIDOC CRM present logic-driven structures that enable documentation to capture underlying semantics concealed beneath domain-dependent documentation structures. Digital library models (Candela et al., 2007; Kovacs & Micsik, 2005) describe the digital library environment in terms of classes, subclasses and implicit relationships. As

[6] See http://www.nationalarchives.gov.uk/PRONOM/

part of the work of the *CASPAR* project[7] a Core Ontology for Dependencies facilitates documentation of information dependencies, both semantic and structural, and *PreScan* (Marketakis et al, 2009) supports automatic extraction of metadata and its encoding in RDF.

Sculpting the World

We conceived an ontology structure to effectively illustrate the range of influential factors within a preservation environment that collectively determine the extent to which preservation objectives can be realistically accomplished. We build this on an assumption (largely demonstrated throughout the deployment of the *DRAMBORA* audit process) that preservation can be understood as a complex interrelationship of objectives, policies, activities, resources and rights. Relating these individual facets can enhance our understanding of risk causation and recovery, illustrating in high level and more specific, practical terms how information preservation can be improved, and flexible enough to reflect a wide range of information environments. Use cases include risk identification tools (whereby users can identify familiar circumstances and trace to hitherto unacknowledged risks) and mapping tools, risk resolution, where risk can be traced to appropriate mitigation measures and preservation system gap analysis, where objectives and facets can be mapped to ontology elements.

Implicit within a suitable ontology is the capacity to trace risk exposure, cost and preservation best practice in terms of relationships with organizational processes, assets and contextual factors. An obvious application is the traversal of a network of related risks, in order to determine the factors exacerbating each, and represent more clearly the wider implications of a particular circumstance. Linear means for recording such information, as currently exists (like a conventional organizational risk register) makes exploration of such relationships tedious and error prone. Other applications include risk resolution, where known risks can be traced to appropriate mitigation measures documented within the ontology and preservation gap analysis, where objectives and facets can be mapped to ontology elements in order to ensure the completeness of existing or proposed systems.

Primary Datasets

The process of ontology development depended on two primary datasets, with an initial structure already largely conceived as described in the previous section. The dataset that informed the main intellectual development of *PORRO* (and was most influential in establishing its structure) was the set of detailed reports, and accumulated evidence from the repository evaluations already undertaken and facilitated within the course of *DRAMBORA's* development and validation (see above). The datasets had to be curated and anonymized to conform to pre-existing agreements and generalized to ensure their wider applicability.

In order to correctly model a particular domain the ontology architect must reflect on that which distinguishes it. Audit reports and associated "raw" evidence presented a broad perspective of a range of preservation contexts. We parsed each report, isolating information elements (representative of both strengths and omissions from each assessed repository) and determined from these a set of 130 generic categories for preservation that could be widely applicable. These were subdivided into issues associated with principal functions of ingest, preservation and access; physical or technical infrastructure; organizational provisions and policy. The next step was to begin to extrapolate from these information elements a relational structure – within *Semantic Mediawiki*

[7] See http://www.casparpreserves.eu

entities were modeled and populated with isolated information elements. This resulted in initial groupings of preservation "actions", "resources", "policies" and "mandates" related accorded to common associated preservation objectives. Further mappings were made from information elements to corresponding parts of the TRAC checklist, intended to offer illustrations of how in practice these often opaque criteria may be satisfied.

A bespoke ontology manager application was developed to enable the ontology to be iterated and defined. The first step was the generalization of elements considered too specific for widespread applicability, resulting in a two tier hierarchy of resources, activities, rights and responsibilities (formerly mandates) and parameters (formerly policy). This was intended to ensure the discoverability of more specific example practice via terminology with more general meaning. Appropriate relationships were defined and modeled, intended to reflect and illustrate systematic, functional relationships within the example preservation contexts, as well as risk causality relationships, highlighting where elements are threatened by particular risks, and where they influence risk likelihood and/or impact. The ontology currently consists of around 7500 triples, representative of relationships between 625 top level relatable elements. More specific (pre-generalization) data is also related and discoverable, accounting for around 1233 additional sub-elements. A simple AJAX web application for traversing the ontology was developed[8].

Relationships were iteratively developed. Preservation goals seemed an appropriate starting point, and these are derived from reference literature such as OAIS (CCSDS, 2002) and TRAC as much as from the individual (and more subjective) case studies. Reflecting their origins, these are related to the rights and responsibilities which motivate or legitimize them. Preservation goals find their first practical expression in parameters, which characterize them, illustrating what is required for their accomplishment (those characteristics that transform broadly defined goals into SMART objectives). Parameters also direct the activities which are undertaken to satisfy them and are evidenced by specific resources (most often documentation such as policy, but sometimes implicitly, like in software algorithms). Activities are supported by, and may also enhance resources. Resources may be required by other resources. Using this latter relationship we can represent semantic or structural dependencies between content information and infrastructure. Figure 2 illustrates these core relationships.

[8] See http://mchughontology.hatii.arts.gla.ac.uk/ontologybrowser/

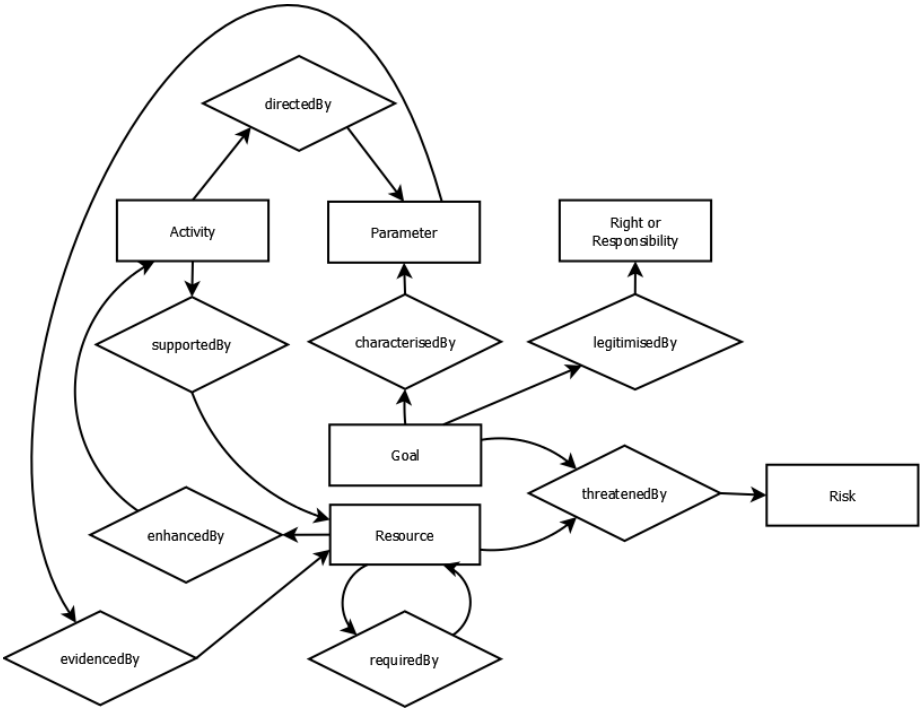

Figure 1 Ontology Relationships

PreservationGoals: largely corresponds to criteria in documents such as TRAC, intended to provide an expression of preservation aims, in order to reflect the diversity of goals evident throughout the digital library and broader preservation landscape.

PreservationParameter: enacted in some aspect of operational policy, lending a specificity and measurability to preservationGoals. Their value defines in a more tangible way the meaning of satisfying individual preservation objectives, and reflects the diversity of the preservation landscape.

PreservationRightsOrResponsibilities: arise from the context within which preservation is undertaken, includes any kinds of debts, obligations, liabilities or enablers. Contracts or legislative mandates are obvious examples. The conferment of mandate may be a risk limiter, or expose greater risks. A common risk that may arise is associated with incompatibilities between a particular liability and a business objective. A digitization project may face a conflict between its objective to make available digitized copies of its whole collection with intellectual property law liabilities which restrict the dissemination of copyright content. A consequential risk is that either the objective will fail or the liabilities escalate.

PreservationActivities: processes intrinsic to the preservation context that is associated with the accomplishment of particular objectives in terms of the rules defined in policy.

PreservationResources: tangible or non-tangible resources within the preservation context that influence the severity or existence of risk. These are generally divisible between those things fundamental to the preservation process (and normally intended to assist in the management of risk e.g. software access systems), and those that are valued in and of themselves, as part of a core business objective (e.g. preserved digital objects, financial profit). Resources are exposed to threats of loss or failure with consequences in terms of their contribution to risk management activities, and wider implications in terms of success of procedures and associated wider objectives. Resources'

contributions to risk causation may be in terms of their insufficiency, associated conflicts, arising liabilities or lack of appropriate deployment. Resources includes expressions of repository roles, including 'typical' preservation roles including data archivists and librarians, information architects, system administrators and developers and those external role-holders such as depositors, consumers and information owners.

PreservationRisks: risks are negative outcomes that can be related to other facets as both causation and mitigation factors. Risks can influence, and have their severity or realization determined by other risks, in various ways. Risk types may be broadly subdivided into risks of failure (threatening goals), loss (threatening resources), or liability. They can be influenced by the existence, absence or specific characteristics of individual activities, resources, rights and responsibilities, and by the cumulative effects of multiple concurrent factors. More importantly, risks can both follow from or be rendered more severe as a consequence of other risks.

Applications for PORRO

The ontology is ever-evolving as more content is added (mainly as sub-elements to the near-comprehensive generalized top level elements) and the ontology continues to be enriched with additional relationships. Figure 3 illustrates one of the ontology's user interfaces, displaying an entry corresponding to the activity "*Automate Metadata Extraction*". Each such entry is hyperlinked, enabling straightforward traversal of ontology relationships. Furthermore, two additional applications have been developed that rely upon the ontology, leveraging its value as a knowledge base for planning and validating digital preservation and data management activities.

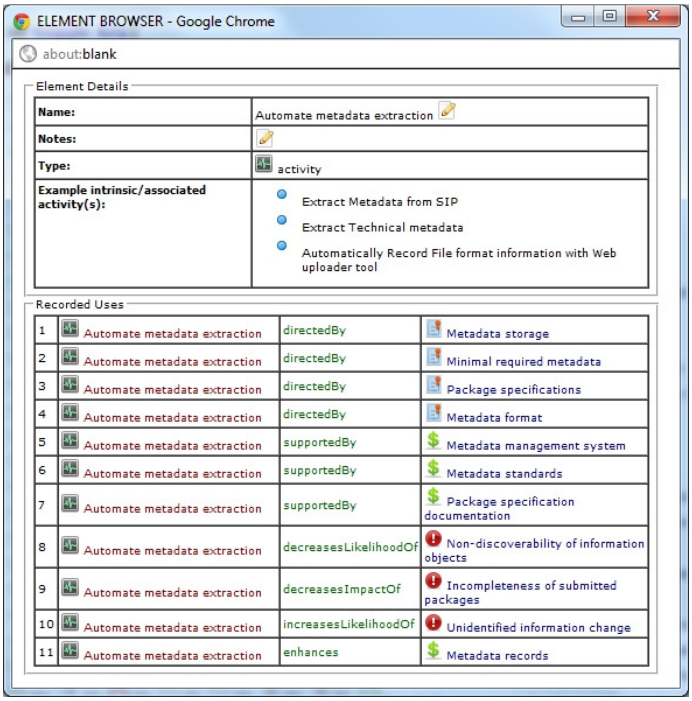

Figure 2: "Automate Metadata Extraction"

The *Long Term Digital Preservation* component developed as part of the 3D-Coform project[9] aims to provide a repository model for the storage, description, distribution and management of three-dimensional object data. Included is a risk association manager, intended to illustrate where and how generic risk factors and risks are manifested within the 3D-Coform information space. Instead of manually encoding risk relationships between 3D-Coform content we take PORRO's more generic examples and map these to 3D-Coform information elements to understand where risk exposure may reside. A variably granular level of mappings are permitted between PORRO elements and either 3D-Coform information types (a taxonomy was created in the project) or specific instance values. One may map a particular generic resource (e.g. *Ingest Platform*), with the broadly encompassing DeviceType *Laser Scanners*, or, if it's more appropriate a specific individual model or example of laser scanner. Traversing the ontology then reveals relationships with other mapped content, or challenges the user to determine whether generic activities, policies or resources which appear to be required have been adequately implemented. In tandem with a preservation package manifest, which is also created within this long term management tool, this enables a clear risk profile to be presented, with closely associated risks and potential additional risk mitigation approaches clearly identifiable, albeit generically expressed. This also enables the continued population of practical expressions within the ontology, and the development of domain-specific mapping sets.

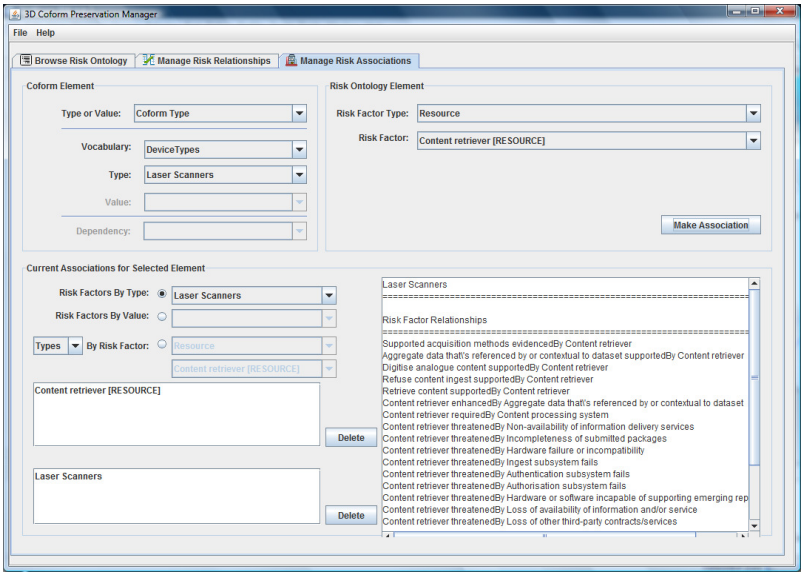

Figure 3: 3D Coform Preservation Manager

PORRO is also deployed within the *Collaborative Assessment of Research Data Infrastructure and Objectives* (CARDIO) tool[9] developed in association with the Digital Curation Centre (DCC) project. This is a tool and associated workflow for performing data management maturity capacity and capability assessment across a data context, which supports and requires a collaborative approach. It enables comparison against a consolidated collection of real world data, via the DCC CARDIO Knowledge Base, founded upon the PORRO ontology. The ontology has been specifically encoded to correspond with the thirty or so infrastructural facets that structure the CARDIO assessment, ranging between issues of organization, resources and technology. This enables the provision of relevant best practice considerations in each specific area, as well as more complex risk relationships.

[9] See http://www.3d-coform.eu

Reaction to CARDIO has been positive –its role in promoting widespread analysis across multiple perspectives within a single institutional setting has been particularly well received. It is now established as a core utility within the Digital Curation Centre's institutional engagement program, whereby around sixteen UK HE institutions will benefit from consultancy services aimed at boosting data management capacity and understanding[10].

Evaluation and Conclusion

Our proposal implies a capacity to represent and illustrate a full range of digital preservation facets and their relationships irrespective of discipline or domain and support and inform validation and planning activities (and ultimately their automation). It is in these terms that evaluation is undertaken.

These objectives can be distilled into core desirable qualities of completeness, applicability and usefulness. The first is evaluated by comparison with evidence of existing practice. To date, discounting spam other erroneous registrations, around four hundred repositories are registered as users of *DRAMBORA's* interactive online tool, representative of institutions including national libraries and archives, academic research repositories, commercial data centers and financial services institutions. Of these, the database and logging systems reveal high usage activity from around one hundred repositories. Mapping user submissions to *PORRO* has revealed the ontology's breadth is sufficient. Its dual tiered approach to recording information facets (with generalized entries linking to more specific example 'implementations') ensures its scalability to encapsulate emerging trends while maintaining its generic qualities and without becoming skewed in any specific disciplinary direction. Likewise, its alignment with acknowledged 'best practice' resources such as TRAC (as described above, the ontology contains mappings to TRAC criteria) provide further reassurance of its completeness, at least in terms of scope.

The ontology relationships are evaluable only by exposure to real world circumstances, and their successful deployment in practical scenarios. PORRO's integration within novel tools and its increasingly prominent role in high profile *Digital Curation Centre* engagement activities provide some assurances of its effectiveness. Likewise, this provides some evidence of its applicability to diverse domains. Within *3D-Coform* mappings have been made between ontology elements and discipline-specific terminology with success.

Full evaluation case studies are not particularly widely available within the preservation and data management context, but brief reports from the *Center for Research Libraries' Certification of Digital Archives* and *Certification and Assessment of Digital Repositories* projects were published via its website[11]. These projects included assessments of *Portico* (on two separate occasions); the *Inter-University Consortium for Political and Social Research*; the *LOCKSS* distributed archiving system; and *HathiTrust* at the University of Michigan. Aligning the broad findings of these to *PORRO* reveals broader recommendations than issued within these brief reports.

PORRO has been retrospectively used in association with these reports to simulate the ways in which the ontology may support the evaluation process. The first conclusion is that the ontology supports

[10] See http://www.dcc.ac.uk/news/working-together-institutional-engagements
[11] See http://www.crl.edu/archiving-preservation/digital-archives/digital-archive-reports

the straightforward mapping of these findings to its implicit information elements, most obviously via the corresponding TRAC criteria, but also with ease via corresponding activities, resources, policies or rights whether by virtue of their existence or omission. Its added value is the straightforward revelation of associated issues and additional risk exposure. Taking the issue of succession planning which consistently appears we can look to one of two corresponding objectives within *PORRO*, Establish relationships with succession partners or Establish appropriate strategies for facilitating succession of organization or content. Traversing the ontology we see that the *Succession arrangement* parameter which helps characterize the first objective in turn directs the activity *Establish succession arrangements*. While this appears perfectly intuitive it is at this point that we begin to attain greater insights, where we see that this activity may be supported with resources such as *Membership of partners' network*. We also reveal additional motivations for pursuing this activity, since it can limit the impact of risks such as *Loss of mandate*, *Budgetary reduction* or *Enforced cessation of repository activities*.

Clearly this example is simple, but illustrative of the value of a related network of elements – one may more intuitively use these links to explore from a starting point of risk exposure to find the elements best suited to their resolution. The ontology is more useful still when used in tandem with TRAC (and by extension the forthcoming ISO standard). The generic TRAC criteria are accompanied by examples of evidence, intended to illustrate what must be demonstrable to achieve conformity. *PORRO* not only reflects this example evidence (which typically amounts to types of documentation within which evidence might be found), it exceeds it, with details of all the associated contextual arrangements which may indicate the criteria's satisfaction. TRAC criterion B1.2 (a random selection) is entitled "Repository clearly specifies the information that needs to be associated with digital material at the time of its deposit (i.e., SIP)". Suggested example evidence for auditors includes transfer requirements and producer-archive agreements. This is mapped to *PORRO* via the objective *Define ingest package specification* and provides means for interpretation of TRAC's frequently misunderstood provisions.

Define ingest package specification	
characterisedBy	Policy on relationship between ingest, archival and dissemination packages
characterisedBy	Minimal required metadata
characterisedBy	Package specifications
characterisedBy	Metadata creation responsibility
characterisedBy	Metadata creation workflow
legitimisedBy	Has prescribed minimal metadata requirements
threatenedBy	Extent of what is within the archival object is unclear
threatenedBy	Shortcomings in semantic or technical understandability of information

threatenedBy	Archival information cannot be traced to a received package
threatenedBy	Loss of authenticity of information
threatenedBy	Incompleteness of submitted packages
threatenedBy	Structural non-validity or malformedness of received packages
threatenedBy	Destruction of primary documentation
threatenedBy	Loss of information provenance

Figure 3: "Define Ingest Package Specification"

In terms of its ability to accommodate diverse information facets and support the expression of myriad associations, *PORRO* is successful. *PORRO's* content is navigable, relatable and intuitive, capable of illustrating close and distant relationships between various system and information components. Even in isolation, the ontology enables digital libraries to reference the encapsulated knowledge in order to support their own risk assessment and preservation planning exercises. Since the ontology is intended to present a holistic vision of managed risks one can determine risk exposure by reference to infrastructural components that are lacking in an example institution, or focus on risks threatening vital provisions in priority areas. In the context of risk management respondents' confidence in their perceived organizational maturity would be challenged by exposure to possible risks (with real world precedent) that may pose threats. For example, if respondents consider elements of their legal infrastructure to be very mature they can traverse a small number of relationships to confront possible risk scenarios concerning IPR infringement, Freedom of Information liabilities or contractual breach. If satisfied that these risks are adequately countered they will have greater faith in their assertion. Conversely, the process may prompt an awareness of shortcomings that were not previously well understood. The ontology is expected to scale to reflect the very latest perspectives in preservation decision making, and with additional population present further insights.

Bibliography

Ahmed, M., Anjomshoaa, A., Nguyen, T.M., & Tjoa, A.M., (2007). Towards an ontology-based risk assessment in collaborative environment using the semanticlife. In ARES '07: Proceedings of the The Second International Conference on Availability, Reliability and Security, pp 400- 407, Washington, DC, USA, 2007. IEEE Computer Society.

Barateiro, J., Antunes, G., Freitas, F., Borbinha, J. (2010). Designing Digital Preservation Solutions: A Risk Management-Based Approach. International Journal of Digital Curation (IJDC), ISSN: 1746-8256, Issue 1, Volume 5, pp 4 – 17.

Candela, L., Castelli, D., Ferro, N., Koutrika, G., Meghini, C., Pagano, P., Ross, S., Soergel, D., Agosti, M., Dobreva, M., Katifori, V., & Schuldt, H. (2007). The DELOS Digital Library Reference Model - Foundations for Digital Libraries. Pisa, Italy.

Consultative Committee for Space Data Systems (CCSDS) (2002). Reference Model for an Open Archival Information System (OAIS). ISO 14721:2003.

Consultative Committee for Space Data Systems (CCSDS) (2011). Audit and Certification of Trustworthy Digital Repositories, Recommended Practice. CCSDS 652.0-M-1. Magenta Book, September 2011 (Draft ISO 16363).

Consultative Committee for Space Data Systems (CCSDS) (2011). Requirements for Bodies providing Audit and Certification. CCSDS 652.1-M-1. Magenta Book, November 2011 (Draft ISO 16919).

Dappert, A., Farquhar, A. (2009). Modelling Organizational Preservation Goals to Guide Digital Preservation. International Journal of Digital Curation (IJDC), ISSN: 1746-8256, Issue 2, Volume 4, pp 119 - 134

Ekelhart, A., Fenz, S., Klemen, M., & Weippl, E. (2007). Security ontologies: Improving quantitative risk analysis. In System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on, pages 156a-156a.

Fenz, S. & Neubauer, T. (2009). How to determine threat probabilities using ontologies and bayesian networks. In CSIIRW '09: Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research, pages 1-3, New York, NY, USA, ACM.

Hunter, J. & Choudhury, S. (2004). A semi-automated digital preservation system based on semantic web services. In JCDL '04: Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries, pages 269-278, New York, NY, USA. ACM.

Hunter, J. & Choudhury, S. (2005). Semi-automated preservation and archival of scientific data using semantic grid services. In CCGRID '05: Proceedings of the Fifth IEEE International Symposium on Cluster Computing and the Grid (CCGrid'05) - Volume 1, pages 160-167, Washington, DC, USA. IEEE Computer Society.

Kovacs, L. & Micsik, A. (2005). An ontology-based model of digital libraries. In Digital Libraries: Implementing Strategies and Sharing Experiences, pages 38-43.

Marketakis, Y., Tzanakis, M., & Tzitzikas, Y. (2009). Prescan: towards automating the preservation of digital objects. In MEDES '09: Proceedings of the International Conference on Management of Emergent Digital EcoSystems, pages 404-411, New York, NY, USA. ACM.

McHugh, A., Ruusalepp, R., Ross, S., & Hofman, H. (2007). The Digital Repository Audit Method Based on Risk Assessment. DigitalPreservationEurope and Digital Curation Centre.

Raskin, V., Hempelmann, C.F., Triezenberg, K.E., & Nirenburg, S. (2001). Ontology in information security: a useful theoretical foundation and methodological tool. In NSPW '01: Proceedings of the 2001 workshop on New security paradigms, pages 53-59, New York, NY, USA, ACM.

Ross, S., McHugh, A., Innocenti, P., & Ruusalepp, R. (2008). Investigation of the potential application of the Drambora toolkit: an assessment of the repository aspects in digital libraries. Glasgow, UK. DELOS Network of Excellence.

CRL-OCLC. (2007). Trustworthy Repositories Audit and Certification: Criteria and Checklist. Chicago, IL and Dublin, OH: Center for Research Libraries and OCLC, 2007. Available at: http://www.crl.edu/sites/default/files/attachments/pages/trac 0.pdf.

Dale, R., Bellinger, M. (eds). (2002). Trusted Digital Repositories : Attributes and Responsibilities. System, 65(May), p.70. Available at: http://www.oclc.org/research/activities/past/rlg/trustedrep/repositories.pdf.