



OCLC Systems & Services: International digital library perspectives

Scalable decision support for digital preservation: an assessment
Christoph Becker Luis Faria Kresimir Duretec

Article information:

To cite this document:

Christoph Becker Luis Faria Kresimir Duretec , (2015), "Scalable decision support for digital preservation: an assessment", OCLC Systems & Services: International digital library perspectives, Vol. 31 Iss 1 pp. 11 - 34

Permanent link to this document:

<http://dx.doi.org/10.1108/OCLC-06-2014-0026>

Downloaded on: 01 February 2015, At: 16:32 (PT)

References: this document contains references to 29 other documents.

To copy this document: permissions@emeraldinsight.com

The fulltext of this document has been downloaded 26 times since 2015*

Users who downloaded this article also downloaded:

Christoph Becker, Luis Faria, Kresimir Duretec, (2014), "Scalable decision support for digital preservation", OCLC Systems & Services: International digital library perspectives, Vol. 30 Iss 4 pp. 249-284 <http://dx.doi.org/10.1108/OCLC-06-2014-0025>

Laurie Borchard, Michael Biondo, Stephen Kutay, David Morck, Andrew Philip Weiss, (2015), "Making journals accessible front & back: examining open journal systems at CSU Northridge", OCLC Systems & Services: International digital library perspectives, Vol. 31 Iss 1 pp. 35-50 <http://dx.doi.org/10.1108/OCLC-02-2014-0013>

H Frank Cervone, (2015), "Three trends in higher education and their potential impact on information agencies", OCLC Systems & Services: International digital library perspectives, Vol. 31 Iss 1 pp. 7-10 <http://dx.doi.org/10.1108/OCLC-10-2014-0034>

Access to this document was granted through an Emerald subscription provided by 417532 []

For Authors

If you would like to write for this, or any other Emerald publication, then please use our Emerald for Authors service information about how to choose which publication to write for and submission guidelines are available for all. Please visit www.emeraldinsight.com/authors for more information.

About Emerald www.emeraldinsight.com

Emerald is a global publisher linking research and practice to the benefit of society. The company manages a portfolio of more than 290 journals and over 2,350 books and book series volumes, as well as providing an extensive range of online products and additional customer resources and services.

Emerald is both COUNTER 4 and TRANSFER compliant. The organization is a partner of the Committee on Publication Ethics (COPE) and also works with Portico and the LOCKSS initiative for digital archive preservation.

*Related content and download information correct at time of download.

Scalable decision support for digital preservation: an assessment

Christoph Becker

Faculty of Information, University of Toronto, Toronto, Canada

Luis Faria

Keep Solutions, Braga, Portugal, and

Kresimir Duretec

*Information and Software Engineering Group,
Vienna University of Technology, Vienna, Austria*

11

Received 16 July 2013
Accepted 21 August 2013

Abstract

Purpose – This article aims to evaluate a new architecture for scalable decision-making and control in preservation environments for its ability to address five key goals: scalable content profiling; monitoring of compliance, risks and opportunities; efficient creation of trustworthy plans; context awareness; and loosely coupled preservation ecosystems. Scalable decision support and business intelligence capabilities are required to effectively secure content over time.

Design/methodology/approach – We conduct a systematic evaluation of the contributions of the SCAPE Planning and Watch suite to provide effective and scalable decision support capabilities. We discuss the quantitative and qualitative evaluation of advancing the state of art and report on a case study with a national library.

Findings – The system provides substantial capabilities for semi-automated, scalable decision-making and control of preservation functions in repositories. Well-defined interfaces allow a flexible integration with diverse institutional environments. The free and open nature of the tool suite further encourages global take-up in the repository communities.

Research limitations/implications – The article discusses a number of bottlenecks and factors limiting the real-world scalability of preservation environments. This includes data-intensive processing of large volumes of information, automated quality assurance for preservation actions, and the element of human decision-making. We outline open issues and future work.

Practical implications – The open nature of the software suite enables stewardship organizations to integrate the components with their own preservation environments and to contribute to the ongoing improvement of the systems.

Originality/value – The paper reports on innovative research and development to provide preservation capabilities. The results of the assessment demonstrate how the system advances the control of digital preservation operations from *ad hoc* decision-making to proactive, continuous preservation management, through a context-aware planning and monitoring cycle integrated with operational systems.

Keywords Business intelligence, Preservation planning, Scalability, Digital libraries, Repositories, Decision support

Paper type Research paper

Part of this work was supported by the European Union in the 7th Framework Program, IST, through the SCAPE project, Contract 270137, and by the Vienna Science and Technology Fund (WWTF) through the project BenchmarkDP (ICT12-046).



OCLC Systems & Services:
International digital library
perspectives

Vol. 31 No. 1, 2015

pp. 11-34

© Emerald Group Publishing Limited
1065-075X

DOI 10.1108/OCLC-06-2014-0026

1. Introduction

This article continues the discussion in [Becker et al. \(2014\)](#) and presents a systematic assessment and evaluation of the SCAPE decision support environment comprising PLATO, SCOUT and C3PO. We discuss the improvements and identified limitations of the presented system. We, furthermore, discuss the quantitative and qualitative evaluation of advancing the state of art and report on a case study with a national library. Finally, we summarize the contributions and provide an outlook on future work.

2. Evaluation and assessment

While some of the questions that are raised by the design goals discussed in [Becker et al. \(2014\)](#) can be readily evaluated using standard metrics, others require a detailed qualitative assessment. This section will discuss how to systematically assess improvements on the dimensions of trust and scalability. We will report on a typical case study conducted with the State and University Library Denmark, discuss key metrics that can be used for evaluation, apply them to assess recent advances, and discuss a set of limitations. We further discuss how these findings can be applied on a wider scale.

2.1 Evaluation dimensions and challenges

Five major design goals have been proposed by [Becker et al. \(2014\)](#):

- *G1*: Scalable content profiling is required to create and maintain an awareness of the holdings of an organization, including the technical variety and the risk factors that cause difficulties in continued access and successful preservation.
- *G2*: Monitoring compliance, risks and opportunities is a key enabler to ensure that the continued preservation activities are effective and efficient.
- *G3*: Efficient creation of trustworthy plans is required so that preservation can function as a normal part of an organization's processes in a cost-efficient way.
- *G4*: Context awareness of the systems ensures that they can adapt to the specific situation rather than provide generic recommendations or require extensive manual configuration.
- *G5*: Loosely coupled preservation ecosystems, finally, enable organizations to follow a stepwise adoption path and support continuous evolution of the preservation systems as new solutions and improved systems emerge.

Given this set of design goals, it is clear that a systematic evaluation has to be based on both qualitative and quantitative criteria and account for the various socio-technical dimensions of the design problem.

Scalable content profiling requires, first and foremost, efficiency in the data processing system. This can be measured in terms of the amount of data processed in a certain timeframe using a defined set of resources. This applies to the content profiling tool C3PO. The efficiency of decision-making, on the other hand, can be measured in controlled experiments. This, however, has to be done in a real-world environment to be meaningful, which creates additional challenges for a large-scale assessment and has to be interpreted with caution.

The effectiveness of a preservation system is composed of several heterogeneous and asynchronous processes, collaborating over time and controlled by decision-makers in a

real organization, is much harder to measure, as very often what needs to be measured in terms of the effects is time-delayed, and to a large extent defies objective measures in the present time. Similarly, trust is extremely hard to measure, and the preservation community has for a decade discussed different ways of assessing the trustworthiness of a repository (Ross and McHugh, 2006; OCLC and CRL, 2007). The resulting criteria catalogue ISO 16363 (ISO, 2010) provides a useful checklist for assessing the assumed trust of an organization and, hence, can form a guideline for evaluation, but does not apply to the actual operations and the preservation lifecycle on the operational level. The Plato planning approach that forms the basis for the architecture presented here has been designed with these criteria in mind and evaluated for adherence with and support of these criteria (Becker *et al.*, 2009). However, it can be argued that more holistic perspectives are required to assess and improve an organization's trustworthiness, perspectives that emphasize enterprise governance of IT and the maturity of organizational processes (Becker *et al.*, 2011).

The following discussion is designed loosely along the Goal–Question–Metric paradigm (Basili *et al.*, 1994). Each goal is associated with a set of questions corresponding to the objectives outlined by Becker *et al.* (2014). The answers to these should support an assessment as to how far the goal has been achieved. To this end, each question is further linked to a set of metrics that provide objective indicators to support an answer to the question. We discuss each of the design goals in turn and discuss the specific questions that need to be answered to provide an assessment of how the state of art is improved with the proposed system design and implementation. This forms the basis of a systematic discussion, taking into account the quantitative indicators and the qualitative discussion of the state of art.

2.2 Evaluation of design goals

2.2.1 *G1 provides a mechanism for scalable in-depth content profiling.* Figure 1 poses the key questions we need to answer to evaluate the scalability and quality of in-depth profiling. Content profiles need to be meaningful, i.e. cover the interesting features that are known to be relevant for preservation, and trustworthy. Clearly, a profile covering only the size of files will be less meaningful than a profile including mime-types, formats and validity. Additionally, a plethora of features influence the success in continued access, ranging from the presence of digital rights management settings to the numbers of embedded tables in electronic documents and other dependencies.

The features needed for in-depth characterization process are broad coverage in terms of supported file formats and the features extracted, usage of a common vocabulary for identification of the formats, feature names and its values and a reasonable low resource consumption and performance so they can be used in large-scale frameworks. By relying on the fits file information toolset, the C3PO profiler maximizes the coverage of features, arguably providing the highest possible feature coverage that can currently be achieved (Petrov and Becker, 2012).

The correctness of aggregation itself can be verified in a straightforward way, as the operations are basic statistical calculations. The correctness of general MapReduce-based operations in themselves can be assumed. On the other hand, the correctness of characterization components to provide accurate feature descriptors based on arbitrary input is far from proven. In fact, current datasets are entirely insufficient for proving correctness of the complex interpretation processes that take place. This, however, is a

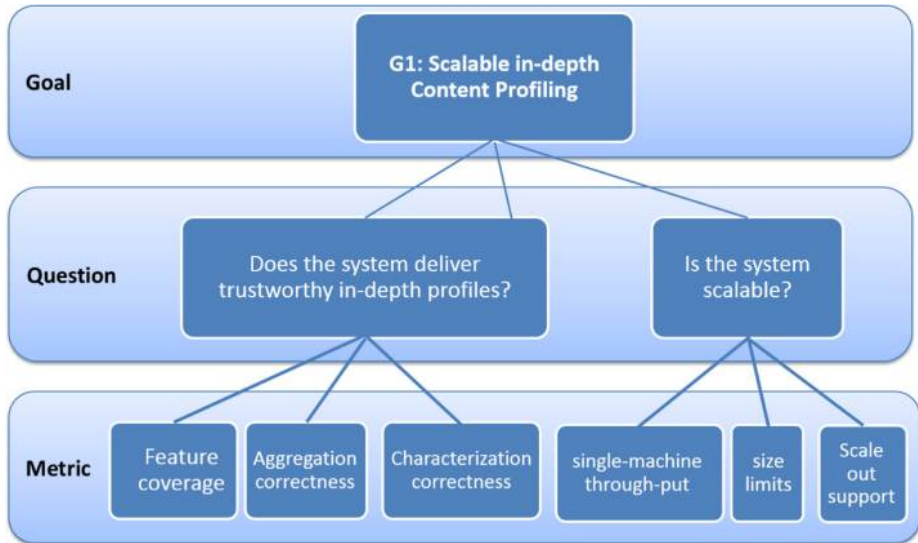


Figure 1.
Scalable profiling
goals

problem on the level of operations and cannot be attributed to the aggregation step of content profiling. Separate efforts are underway to verify the correctness of characterization tools using model-driven engineering to generate annotated test data (Becker and Duretec, 2013). To enable future evolution, yet another aspect of scalability, a meaningful content profiler must be flexible enough to work on arbitrary property sets. C3PO supports this by relying on a generic data model, so that any additional property sets can be profiled.

This also supports the integration of further characterization tools. While FITS fulfils coverage and vocabulary requirements, it consumes considerable resources and takes a substantial amount of time to execute[1]. Hence, C3PO also supports Apache Tika, which in experiments showed far better resource consumption and performance with a throughput of up to 18 GB per minute[2] when used with large-scale platforms such as Apache Hadoop[3]. While in this case, Apache Tika was only used for file format identification, it supports feature extraction and has good coverage of file formats[4], but does not yet use a well-defined vocabulary for the identification of extracted features.

The objectively measurable throughput and resource usage in profiling, then, is the crucial final question. To measure the time and resource behavior of C3PO, a set of controlled experiments has been conducted. The first measured the throughput of C3PO on a single standard machine, while the second used a server with strong hardware and explored the boundaries of scalability by attempting to profile up to 400 million resources (12 terabyte [TB]) in a single profile, enabling a further extrapolation of these results to the entire set of 300 TB in this collection. The third examined the limits of the Web visualization platform to cope with these amounts of data.

The first experiment (Petrov and Becker, 2012) tested the performance on limited resources and showed that on a standard computer with 4 GB of RAM and 2.3 GHz CPU, the ingesting and generation of a profile of 42,000 FITS takes about 1.5 minutes.

Large-scale tests were performed by Niels Bjarke Reimer from the Danish State and University Library[5]. A 12 TB sample was taken from a dataset with 300 TB of the Danish Web archive. FITS was run on the sample content, resulting in 441 million FITS files. This characterization process took about a year to complete[6].

For content profiling, two processing parts need to be considered:

- (1) the gathering of files into the internal data structure; and
- (2) the analysis of that dataset using MapReduce queries.

The experiments were executed on a single machine with the specifications described in Table I.

The first step, which ingests the FITS files into a MongoDB server, was tested with the 441 million FITS items. The graph depicted in Figure 2 shows the import time for samples of around 3,600 files. The y-axis unit shows time in milliseconds and the x-axis unit is a sample number, which can also be considered as a timeline.

The complete import process took less than 80 hours, with an average execution time of 0.65 ms per FITS file. This import time is quite constant with only a few outliers, which implies that the platform and the software are acceptable for importing large amounts of data.

Processor	2 × Intel Xeon × 5660 2.8 GHZ (12 core)
RAM	72 GB
Storage	Isilon storage box with 20 TB storage and 400Gb SSD, connected by a 1 Gbit/s Ethernet network
Operating System	MongoDB Linux × 86 64-bit v2.4
MongoDB	Version 2.4
Application service	Apache Tomcat version 7

Table I.
C3PO scalability test machine specifications




Figure 2.
Performance of C3PO import process using FITS metadata

Source: Reimer *et al.* (2013)

The second step, the analysis of the data using map reduce queries, was tested with a dataset of about 12 million FITS files and took 15 hours and 18 minutes, which is about 4.63 ms per FITS file. Using sharding and MapReduce technologies, the processing time of the second step should also be linear.

In conclusions, both steps are linear and together take on average 5.28 ms per FITS file. This means that processing the current 300 TB dataset would take about 16,170 hours, or about 677 days, on a single machine. As both processes are massively parallelizable and the MongoDB platform already supports sharding and MapReduce, the processing time can be highly reduced by distributing the load on several servers. Substantial resources may be needed to bring the processing time down to a practical time, but this profile does not have to be re-generated frequently.

The C3PO tool also provides a Web interface that supports real-time analytics on the gathered data. While this is not one of the requirements strictly required for automated monitoring, it provides interesting insights into content profiles that is considered highly valuable by the decision-makers. However, in this scenario, the limits of the Web application are revealed. Several test runs were made with different dataset sizes to ascertain the limits of the application. For each test run, two manual procedures on the Web interface were made:

- (1) opening the overview page that calculates, in real time, distributions of several extracted features; and
- (2) drilling down into the characteristics of a subset of the collection, as all the files of a defined format.

Table II shows the results of the tests, which show acceptable results up to around 2.5 million files, with a waiting time of around 6.5 minutes. Above this limit, the Web system does not respond within 3 hours, which is considered unacceptable. Hence, it is not feasible to perform real-time analytics with the current solution on the set of 440 million FITS files of the 12 TB dataset. It should be noted that the analysis of the entire dataset provides a view of over 400 million rows in a table with over 5,000 columns, with a resulting database size of over a TB.

2.2.2 *G2 enables monitoring of operational compliance, risks and opportunities.* This section analyzes how the mechanisms presented by Becker *et al.* (2014) can be used to

Test run	No. of FITS files	Elements size GB	No. of properties	Overview processing time	Drill down processing time
1	13,962	0.03	80	Fast	Fast
2	108,348	0.26	96	18 sec	11 sec
3	363,991	1.00	106	30 sec	34 sec
4	1,020,514	2.46	113	2 min 25 sec	1 min 42 sec
5	1,639,842	3.95	119	3 min 52 sec	2 min 50 sec
6	2,683,596	6.44	119	6 min 28 sec	4 min 25 sec
7	11,905,935	28.63	211	Not finished within 3 hours	N/A
8	441,923,560	1183.50	5,122	N/A	N/A

Table II. Testing the limits of real-time analytics in the C3PO Web application

Source: Reimer *et al.* (2013)

accomplish proactive monitoring of operational compliance, risks and opportunities in a preservation environment.

As outlined in Figure 3, the key questions relate to the identification of aspects that need to be monitored and to the coverage of measures available to provide indicators related to these aspects. Based on an analysis of a reference model for drivers and constraints (Antunes *et al.*, 2011), which classifies each of the influencers a preservation organization should be aware of, the discussion in Becker *et al.* (2012) showed that relevant questions and measures can be derived for each of the influencers of interest. This enables the development of appropriate adaptors for measuring specific indicators pertaining to this driver. Table III shows key examples, while a full discussion and detailed table is provided in Becker *et al.* (2012).

In practice, the achieved coverage of measures is by no means complete, but increasing. Currently, supported sources include format registries, semantic policies,



Figure 3. Monitoring

Driver	Question	Indicator	Sources
Content	Is the content volume growing unexpectedly?	Rate of growth changes dramatically in Ingest	Content profile, repository report API
Operations	Are our content profiles policy-compliant?	Mismatch between content profiles and policy statements	Content profiles, control policy statements
Format	How many organizations have content in this format?	Number of shared content profiles containing a format	Content profiles shared by organizations
Format	What is the predicted lifespan of format X?	Life span estimates based on historic profiles	Model-based simulation

Table III. Selected preservation drivers and related information sources

Source: Becker *et al.* (2012)

content profiles and an automated rendering and comparison tool (Law *et al.*, 2012). A prioritization approach is taken to target first and foremost those aspects that are perceived most critical. The open nature of the adaptor design, the data model and the licensing model has the effect that additional sources can be integrated by anybody in the preservation community, and the coverage is rising steadily.

Monitoring of operational plans is illustrated schematically in Figure 4 on a simplified example. Consider a preservation plan that evaluates four potential actions (“alternatives”) against a set of four decision criteria. These criteria evaluate the important aspects of the data to be preserved, the environment and the actions to be applied. Based on these criteria, preservation actions in question are evaluated and a ranking is calculated. The planner then chooses the best suited action and adopts it. In this example, a check mark denotes a best-in-class performance, a tilde denotes acceptable performance and a cross reflects an unacceptable performance; for example, a process that did not terminate or an image conversion that shows a distorted image. We can see that two alternatives have been rejected, and alternative 1 has the highest score and will be selected[7].

Since the decision criteria identified during planning lead to the adoption of a certain action, they must be monitored during operational executions, as well to enable the organization to track whether the action keeps performing according to the expectations. This is shown on the bottom left of the figure. However, this is not the only aspect prone to evolution:

- New alternatives will emerge over time that may perform better than the chosen alternative. In cases where no alternative was acceptable, this will sometimes be the only thing monitored, as the organization would wait for a better solution to become available before embarking on premature preservation actions. For example, this was the case in Kulovits *et al.* (2009).
- Updated or new quality assurance tools can emerge that provide more reliable or more efficient measures for quality assurance or even the first automated way to measure relevant quality. For example, these could be of the kind described in Jurik and Nielsen (2012) or Bauer and Becker (2011).

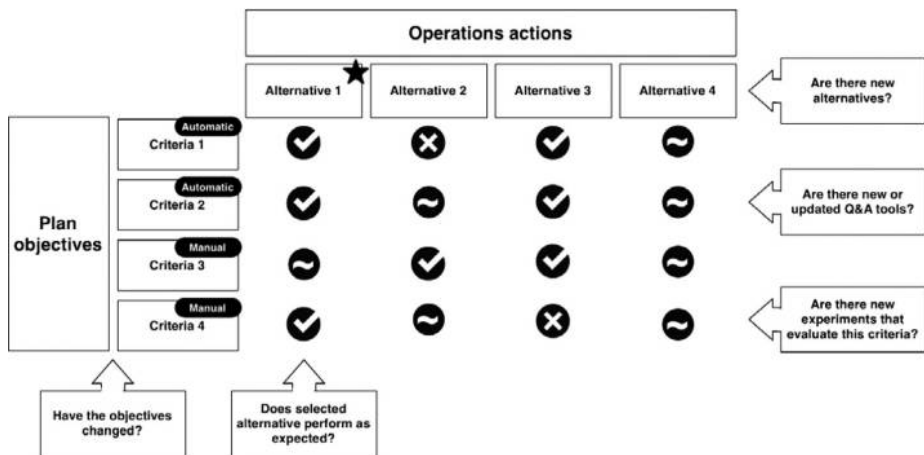


Figure 4. Evaluation of monitoring compliance

- Related to this, experiments including certain criteria may be conducted by other individuals or organizations that can reveal risks and opportunities related to this plan. For example, the chosen quality assurance tool might be shown to malfunction on similar objects, which poses a major risk (Bauer and Becker, 2011).
- Finally, the organization's objectives themselves may shift over time as goals change. This would be reflected by a change in the control policies.

The tool suite described in this article is designed to provide full support for this monitoring scenario. The upcoming release of Plato generates specifications describing the expected quality of service (QoS), similar to a service-level agreement, for the set of decision criteria considered, linked to the corresponding organizational policies and deposits corresponding monitoring conditions in Scout on deployment of a preservation plan. Such QoS specifications are created for those criteria in a tree which are influenced by the dynamic behavior of the service, i.e. the components. That means that they are not created for aspects relating to the format, such as the ISO standardization of PDF versions, but they include criteria such as whether the created files are well-formed. QoS is then measured within executable workflows and monitored for fulfillment. Aspects pertaining to the format and other non-dynamic aspects are monitored as risks and opportunities using Scout.

While Scout is able to collect a wide variety of measures, these are naturally limited by the availability of operations that support such measures. The controlled vocabulary encourages developers to declare which measures their tools deliver to support discovery, but the coverage of measures will naturally vary across different scenarios. It is important to note, however, that any required measures can be integrated by any organization due to the open nature of the ecosystem.

Finally, transparency of the monitoring process is achieved through the usage of the permanent shared vocabulary and the explicit declaration of tolerance levels in the QoS, corresponding to the specified acceptance thresholds that are derived from the organization's control policies.

2.2.3 G3 improve planning efficiency. Previous work has shown that the key challenge in planning is to make the decision-making process more efficient (Kulovits *et al.*, 2009; Becker and Rauber, 2011b). In Becker *et al.* (2014), we reflected on the dimension of trust that should not be sacrificed along this quest. Correspondingly, the key questions shown in Figure 5 relate to the aspect of effort: How long does it take to create one preservation plan now, and how much further improvement is possible?

Previous discussions have shown the trustworthiness of plans produced by Plato (Becker *et al.*, 2009), which is based on the evidence-based measures of decision criteria directly linked to organizational goals and based on factual evidence, documented with full change tracking assigned to acting users. These strengths continue to form the backbone of trustworthy planning. While it is clear that fully automated, i.e. *autonomous*, preservation planning is contradicting the goal of trustworthiness in this domain, the goal, nevertheless, must be to achieve a substantial increase in efficiency (Becker and Rauber, 2011b).

We will focus our discussion on measurements of effort on a controlled case study conducted with the Danish State and University Library, described in detail by Kulovits *et al.* (2013a, 2013b). In this study, a set of responsible decision-makers and experts from the library set out to create a preservation plan, with the assistance of a planning expert

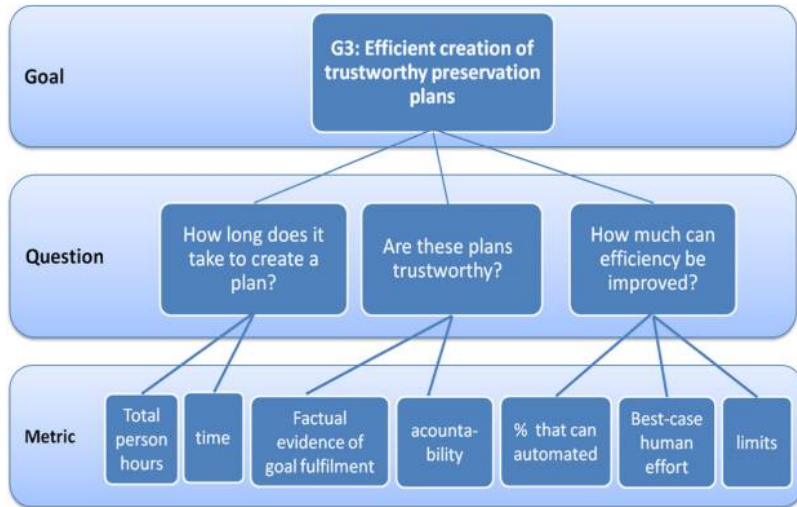


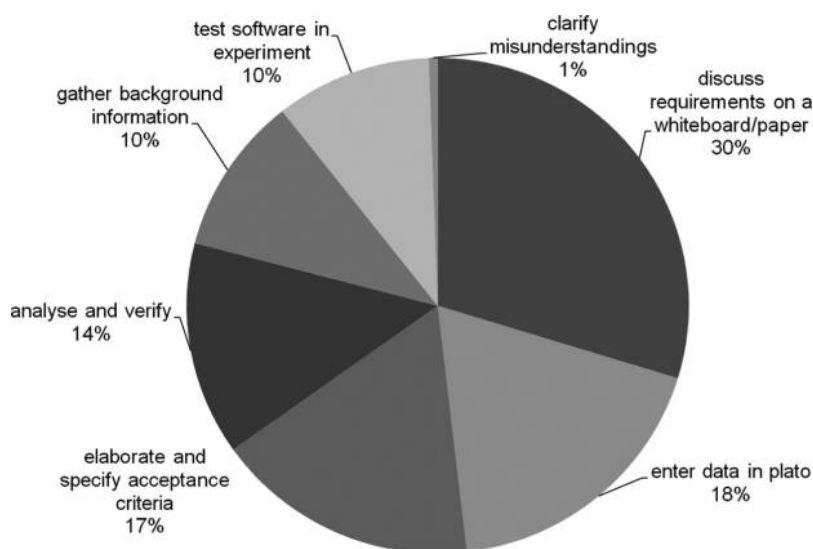
Figure 5.
Efficient creation of
trustworthy plans

and a moderator who kept time of all activities throughout the planning process. The goal of planning was to create a preservation plan for a large set of audio recordings; the drivers and events motivating the plan included the goal to homogenize the formats of the library's holdings and provide well-supported and efficient access to authentic content. The team at the library has comprehensive expertise in all relevant areas, which range from technical knowledge on audio formats and quality assurance mechanisms for comparing audio files to a documented understanding of the designated communities and preservation purpose of the content set at hand.

The preservation plan was created using the then-current version 3 of the planning tool Plato, which is the precursor of the solution presented here. The goal was to identify the major areas of decision-making effort and measure the potential improvement that can be realistically achieved.

The total time required to create a preservation plan amounted to 35.5 person hours, completed over a period of two days. This shows, on one hand, that efficient teams in well-established settings can already plan quite efficiently. Nevertheless, the effort must be further reduced to make planning truly a part of "business-as-usual" preservation in practice. To contextualize the effort required in this case, it is important to understand that this effort strongly depends on a well-defined understanding of the decision-making context, including the understanding of the goals and constraints, the expertise of decision makers and the technical proficiency of the staff carrying out the experimental steps of preservation planning. Finally, a strong driver for cost is the homogeneity of content: For large object sets that are very diverse, several preservation plans will have to be created, each respecting to a certain degree the specific aspects of a subset of the content and the means available to ensure access to this subset.

Figure 6 shows the distribution of effort across each of the types of activities that were part of this planning process. It should be noted that several of these activities were in fact on the upper end of the efficiency range for several reasons:



Source: Kulovits *et al.* (2013, 2013b)

Figure 6.
Distribution of effort
across activities in
preservation
planning

- Experiment execution often takes more time. The experiments conducted were highly efficient due to the minimum number of alternatives evaluated, the high technical proficiency of staff, the homogeneity of content and the quality assurance mechanisms used. In many cases, the experimentation processes consumes a multiple of the time. The integration of Taverna workflows and myExperiment can reduce this massively, as potential components can be discovered and automatically invoked within planning. This automation is similar to an existing integration of automated measures in Plato (Becker and Rauber, 2011a), but makes these mechanisms available on an open, standardized and easily extensible basis.
- Background information often is unavailable. This applies in particular to the user communities and the statements of *preservation intent* that many organizations are only now beginning to document systematically (Webb *et al.*, 2013). The organization in question, however, has a stable and well-supported definition of collections and user communities, from which the preservation goals could be derived rather efficiently. Formal policy specification makes this background explicit and known by the systems, so that the effort can be further reduced.
- Analysis and verification is complex. Even with the support of a planning expert, 14 per cent of the time was spent in sense-making, analyzing the completed set of evidence and assessment in the decision-making tool to arrive at a conclusion that was well understood by the stakeholders. This points to the need for improving the decision support tool in visualizing results in a more easily understandable and user-friendly way. Improved summaries in Plato are planned to this end.
- Entering data into the system is tedious, in particular for users not familiar with the tools. This was alleviated by the involvement of a planning expert familiar

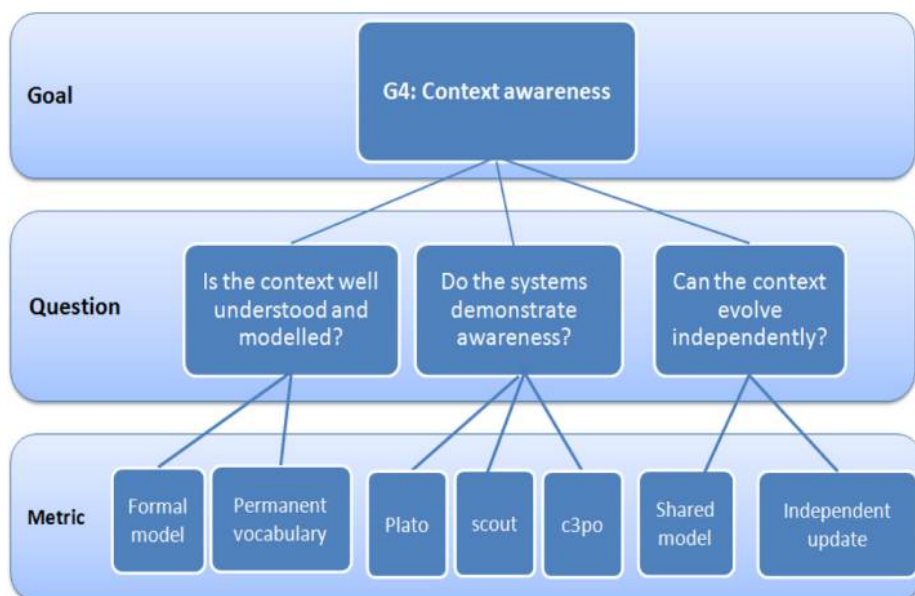
with the tool. Similar to other aspects, this benefits greatly from the integration of the tool with workflows and from the explicit endowment of Plato with an understanding for the policy models of organizations. A subsequent controlled experiment showed that Plato 4 reduced this effort by over 50 per cent (Kulovits *et al.*, 2013a, 2013b).

In an ideal case, the effort required to cover the above aspects (software testing, background information, analysis and verification and data input) can be removed almost entirely. Still, 50 per cent of the time, in this case, would be spent discussing requirements. However, the majority of these concerns objectives about formats, significant properties and technical encoding or representation (Becker and Rauber, 2011a). For all of these aspects, standard definitions are now available as part of the controlled vocabulary that would enable decision-makers to reuse definitions and formalize these aspects on a policy level, removing this activity from the operational planning process. This applies to the designated community and preservation intent statement, as well as to format and risk factors and technical constraints. The control policy statements thus can improve his effort by enabling reuse of these goals and constraints across plans. As the discussion on Goal 4 will show, the context awareness of Plato can eliminate the need for in-depth discussions of requirements *as part of planning* almost entirely.

For an organization that establishes planning as a proper function in its roles and responsibilities and possesses a solid skills and expertise base, we estimate that preservation planning should on average take about one to two person days per plan, provided that policies and content profiles are known and documented. However, a large variance across organizations is to be expected. This estimate will strongly depend on a variety of specific factors and certainly needs to be further validated in longer-term empirical studies. These should in particular also cover the question of homogeneity of content sets covered in a plan: how many plans are required to safeguard a particular heterogeneous set of objects? A detailed discussion on the activities in this process and the relevant skills and expertise is presented in Kulovits *et al.* (2012).

2.2.4 G4 make systems aware of their context. By providing a model that enables decision-makers to formulate policies so they can be understood by automated processes, the systems can understand their context and stay informed about its state. To assess the context awareness of the systems in question, we investigate three distinct aspects: on one hand, the context needs to be well understood and modeled to ensure a solid approach has been taken. On the other hand, each of the systems needs to demonstrate that it can use the part of the context that is relevant for its function appropriately. Finally, it is crucial to ensure that this does not come at the cost of coupling the context too closely with the systems. To this end, we discuss how this context can evolve independently from each of the systems. This is illustrated in Figure 7.

The modular approach of the semantic models has been discussed in Becker *et al.* (2014). A detailed documentation of the model is provided in Kulovits *et al.* (2013a, 2013b). The models are based on W3C-approved standards and follow established Linked Data principles. At the heart of the model is the Resource Description Framework[8] (RDF), a standard model for enabling the representation of data and metadata in subject-object-predicate triples. The Web Ontology Language[9] (OWL) provides the mechanisms for the description of vocabularies, defining

**Figure 7.**
Context awareness

classes and properties. These are used to annotate, describe and define resources. Having well-defined semantics, OWL facilitates reasoning, ontology management and querying of data. The model is represented as an RDF graph and queried using SPARQL. The vocabulary domains have permanent identifiers according to the following ontologies:

- <http://purl.org/DP/preservation-case> contains the basic elements that link a preservation case together.
- <http://purl.org/DP/quality> describes the quality ontology, linking attributes and measures in a domain-specific quality model.
- <http://purl.org/DP/quality/measures> contains the vocabulary individuals that are used for annotating, describing and discovering measures and the mechanisms for measuring.
- <http://purl.org/DP/control-policy>, finally, defines the classes of objectives relevant for making a preservation case operational.

Each of the systems presented is aware of those parts of the model that are relevant for its domain. Correspondingly, each system shows its awareness of this model in different manners.

Plato uses the control policy model in several ways. On one hand, the preservation case provides the basic cornerstones of planning. Instead of providing the documentation of the planning context in a textual form, as it used to be standard (Becker and Rauber, 2011b), a planner who has specified the policy model selects a preservation case to start planning, and the contextual information from this case is extracted from the policy model. Additionally, the objectives and measures specified in the control policy enable the decision support tool to derive the complete goal hierarchy

automatically from the model, leaving it to the decision-maker only to revise, verify and confirm the decision criteria to be used for the experimental evaluation. In the case study discussed above, this requirements specification alone accounted for 30 per cent of the effort. While the policies of course require a similar discussion, much of the objective specification has to be discussed only once and can then be carried forward across preservation cases, which represents a substantial efficiency gain as soon as more than one plan is created. Similarly, the acceptable values, and thus the utility functions associated with each measure, can be computed in a straightforward way based on the objectives specified in the control policies, which presented a potential gain of another 17 per cent in our case study discussed above.

C3PO uses the vocabulary relevant to characterization in the content profile, referencing elements from the quality measures catalogue (such as <http://purl.org/DP/quality/measures#55>). Given that it only provides objective analytics of factual statements about the domain elements, it has no understanding of the policy model and does not require any.

Scout, finally, leverages the policy model for monitoring the alignment of operations and plans to the policies, and also monitors the policy itself: If it is updated, that means that affected plans should be re-evaluated. Specific standard queries are provided as templates that monitor policy compliance. These can be activated by the user. For example, Figure 8 shows Scout starting a monitoring activity on the policy conformance of a specific content set (identified by a collection key). In this case, it shows in a preview

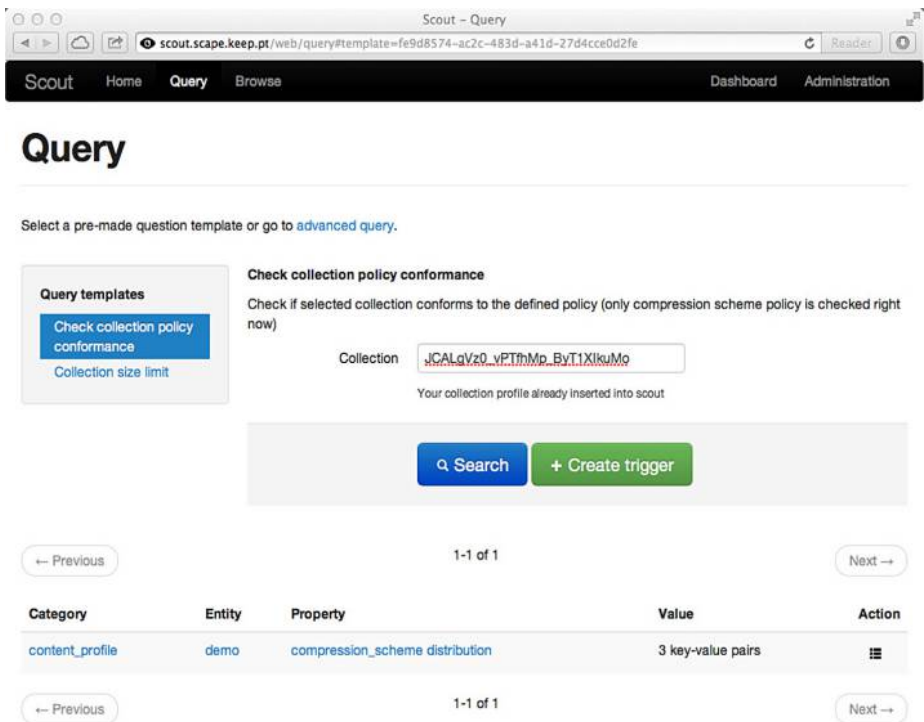


Figure 8. Checking collection policy conformance in Scout

that the property *compression scheme* is violated by three entries, and provides the option to create a continuous monitoring process by specifying a trigger with a condition and an event.

It can be seen that the model of the context is shared between the tools, with the decision-maker updating the ontology independently of the tools. A crucial requirement is that the context model can evolve independently of the systems. This is especially important considering that the current model is very much focused on operational support and can benefit greatly from being expanded to cover aspects of decision-making that are further removed from operations. Similarly, it can be expected that meaningful linkages will surface that connect the existing ontologies to emerging ontologies from neighboring areas ranging from software quality and ontologies for describing software dependencies and platforms to preservation metadata and related policies. The potential for such evolution is guaranteed by the choice of representation and languages, since the Linked Data principles that the model adheres to are designed with these very goals in mind.

2.2.5 G5 design for loosely coupled preservation ecosystems. The design goal of loosely coupled systems is relevant for several reasons. On one hand, it is crucial to enable the stepwise adoption approach preferred by many organizations (Sinclair *et al.*, 2009). On the other hand, it ensures that evolution can take place independently, enabling each organization to replace parts of its system without negatively affecting continued operations, and enables each component of the ecosystem to be sustained independently (to a degree) of the others.

Figure 9 relates these goals to more specific questions. While it is clear that the components are open source, licensed under Open Source Initiative-approved conditions[10], and highly modular, it is useful to consider closely both the functional specifications and the data structures. The API specifications for the SCAPE Planning and Watch suite are in the process of being published openly on

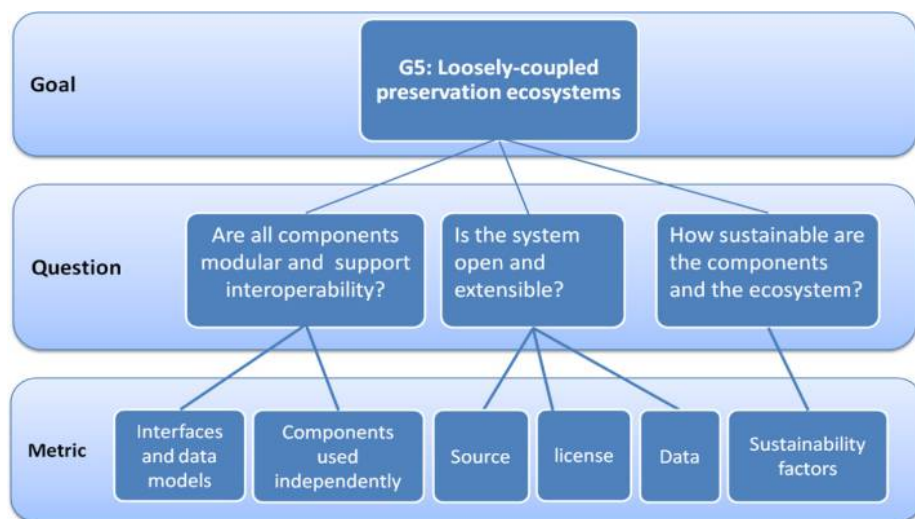


Figure 9.
Loosely coupled
preservation
ecosystems

OCLC
31,1

github. Data exchanged between components is standardized and supported by schemas, as shown in [Table IV](#).

The controlled vocabulary as the glue that connects much of the ecosystem is maintained on github[11]. Curating this vocabulary over the long term will be sustained by a community effort. Recent discussions in the communities of metadata and preservation have brought forward long-term requirements for such evolution that will be considered carefully ([Gallagher, 2013a, 2013b](#)).

26

The components are functionally independent in that every component can and is actually used independently. Nevertheless, it is clear that the compound value proposition is larger than the sum of its parts, serving to encourage take-up of the suite as a whole. Similarly, the usage of this tool suite benefits greatly from integrating also with the workflow development, execution and sharing platforms Taverna and myExperiment, whose latest releases provide specific support for semantic annotation, driven by the requirements outlined in this article.

Since such an ecosystem should be built with sustainable evolution in mind, we consider a recent discussion that identified 11 factors affecting the sustainability of a modular preservation system ([Gallagher, 2013a, 2013b](#)). [Table V](#) shows how our system performs on each of these criteria.

While the ecosystem is well positioned for future sustainability, there is still room for improvements. This includes the development and publication of Technical Compatibility Kits that can automatically test the functional compliance of a component to an API specification, as it has been done for the Data Connector API[13], but also the long-term evolution of vocabularies and any future extensions of the tool suite.

2.3 Practical adoption

Considering the preservation lifecycle outlined in [Becker et al. \(2014\)](#), what does the availability of the described system mean for an organization that has content and a preservation mandate, has set up a reasonable organizational structure and defined corresponding responsibilities, but has not yet ventured to create and maintain specific, actionable preservation plans? The exact measures to be taken will certainly depend on the specific institutional context, but essentially, such an organization can follow a series of steps:

- (1) Getting started entails several aspects.
 - Start content profiling. Run format identification and characterization components such as fits on the set of content to extract metadata, deploy the

Criterion	Plato	Scout	C3PO
All functional interfaces openly published	In progress	In progress	In progress
All data structures documented using standards and schemas	XML schemas published for each version	Linked data model Policy model	XML schema published
Component is used independently	Yes	Yes	Yes
Component follows the controlled vocabulary	Objectives, measures, control policies, preservation cases	Objectives, measures, control policies	Measures

Table IV.
Interoperability of components

Sustainability factors	How does the SCAPE Planning and Watch suite perform?
Ability to view and modify source code	All components are openly licensed, and all source code elements are freely available on a github repository
Widely used	C3PO and Scout are relatively new, but enjoying quick take-up in the community, while Plato has been growing to over 1,000 user accounts since first publication 2008. However, usage so far has been limited to prototypical evaluation rather than production-level deployment, mostly due to the level of effort involved
Well tested, few bugs or security flaws	All tools support automated tests and have an active ticketing system, and the major releases are considered very stable. No security incident has been reported so far
Actively developed, supported	All tools are part of an active development community, continuously supported, and the development platform is hosted by the Open Planets Foundation[12]
Standards aware	All components follow standards on multiple levels wherever possible. This ranges from standard technologies such as Java Server Faces to XML Schema declarations and Linked Data principles
Well documented	All components have extensive code documentation, manuals, built-in help and tutorials, as well as scientific publications explaining the theoretical foundations and practical implications of the software
Unrestricted licensing	All software components are licensed under OSI-approved open licenses such as LGPL and Apache Software License 2.0. All documentation is licensed under the Creative Commons license
Ability to import and export data and code	Preservation plans, executable plans and content profiles can be freely imported and exported, and shared between users. The Scout knowledge base is a Linked Data triple store and hence equally portable
Compatible with multiple platforms	Being based on standard server technologies, all components are compatible with multiple platforms. Plato even integrates with multiple platforms at once in the case of preservation action discovery (Kraxner <i>et al.</i> , 2013)
Backward compatible	This is very relevant in the context of Plato, which is an online service since 2008. Here, there is full backward compatibility with a fully traceable forward conversion upon import of legacy preservation plans. All plans created on the online service have been automatically migrated for all releases Similarly, the knowledge base of Scout is designed to keep growing incrementally, without disposing of accumulated historical data
Minimal customization	There is almost no customization required, since all contextual adaptation of the systems' behavior can be achieved through the configuration of API endpoints and the corresponding definition of control policies

Table V.
Sustainability
evaluation of the
SCAPE planning and
watch suite

Source: Gallagher (2013a, 2013b)

content profiling tool C3PO, gather the metadata and conduct an analysis of the content profile.

- Sign up with SCAPE Planning and Watch, either on the online service[14] or on an organization-specific deployment based on a public code release[15].
- Connect the organization's repository with SCAPE Planning and Watch, either through configuring a standard adaptor or implementing a specific adaptor.

- (2) Specify control policies based on a thorough analysis of the organization's collection, the user communities, and the preservation cases that are considered relevant.
- (3) Activate the monitoring of policies and content profiles in Scout to detect policy violations.
- (4) Create preservation plans to increase the alignment of the organization's content and operations to the goals, as declared in the policies. This planning will be done by evaluating action components using characterization and quality assurance components in Taverna workflows, all integrated in planning. The finished Plans contain workflow specification including QoS criteria that can be automatically monitored.
- (5) Deploy the operational plans to the repository through the plan management API, connected to a workflow engine such as Taverna.
- (6) Establish responsibility for continuous monitoring. This is supported by Scout, which will monitor the compliance of operations to plans and detect risks and opportunities connected to these plans and policies.

2.4 Limitations

From the discussion above, a number of limitations can be observed. These can be divided into limitations of the current capabilities of available tools, which can be expected to grow, and more fundamental limitations of current approaches, which require new perspectives to be overcome; limitations of the problem space that set natural limits to further improvement; and limitations on the quantitative evaluation that can feasibly and meaningfully be conducted. This section discusses those limitations that are seen as central to the further advancement of the state of art.

2.4.1 Coverage and correctness of available measurement techniques. The availability of tools and mechanisms to deliver objective and well-defined measures that are shown to be correct and reliable is a key challenge holding back operational preservation today (Becker and Rauber, 2011b; Becker and Duretec, 2013). Scout supports a growing set of adaptors to feed in measures into the knowledge base, and, by nature of the design, alleviates some of the shortcomings and gaps in existing tools through the free combination of multiple information sources, but still is limited by the availability of these information sources. Similarly, experiment automation in Plato and, equally important, the feasibility of large-scale preservation operations in general, is entirely dependent on the existence of well-tested, efficient and effective mechanisms for quality assurance. Recent work is showing promising advances (Jurik and Nielsen, 2012; Bauer and Becker, 2011; Law *et al.*, 2012), but there is still a wide gap to be addressed for preservation operations to be broadly supported. It seems crucial that this gap is made explicit and shared with a wide community so that efforts to close it can be based on a solid assessment of the shortcomings of existing tools rather than isolated *ad hoc* identification of application scenarios within single institutions, as is often practiced today. Scalable preservation operations are only possible with fully automated, reliable and trustworthy quality assurance; and such quality assurance is expensive to develop and difficult to verify. Only through coordinated community efforts based on solid experimentation can the evidence be constructed to make a convincing case on authenticity (Bauer and Becker, 2011).

The utter lack of solid, reliable and open benchmark dataset with full ground truth is a fundamental inhibitor to validating the correctness of such measures. To address this gap, we are investigating innovative approaches to turn around the publication of test datasets from ex-post annotation, inherently plagued by unreliable ground truth and copyright problems, to an open, model-driven generative approach (Becker and Duretec, 2013).

2.4.2 Scalable distributed and cost-efficient processing: how to profile a petabyte? As shown above, the content profiling tool C3PO provides support for scaling out on distributed platforms. However, it requires considerable resources if the content to be profiled is approaching the petabyte range, and visual analytics are not currently supported on such amounts of data. Yet, it is important to point out that the core goals of content profiling are achieved regardless of the collection size: Visual analytics is an additional capability on top of the processing activity.

To enable cost-efficient creation of large content profiles without visual analytics requirements, we are exploring purely sequential profilers with a small footprint as a low-cost alternative, and we are investigating a set of techniques for feature-space pruning and dimensionality reduction prior to the more expensive processing steps.

Similar considerations apply to preservation operations such as actions, characterization and quality assurance. As noted, the execution of *fits* on the 440 M resources on the Danish Web archive took a year to complete, which clearly indicates the need for improvements. Similarly, automated quality assurance mechanisms are computationally demanding (Bauer and Becker, 2011). These processes need to be supported by parallel execution environments and more efficient algorithms to be truly applicable on large-scale volumes.

2.4.3 The element of human decision-making. As observed above, trustworthy preservation should always be driven by careful decision-making and factual evidence. While this element of human decision-making can be reasonably minimized, replacing it entirely will only be possible once a solid, substantial knowledge base of real-world cases populates the ecosystem described above. Eventually, the human element can in the ideal case be reduced to a policy specification activity and a monitoring oversight function. This is clearly out of scope for this article, but will provide the logical next step in research on preservation planning and monitoring.

2.4.4 Trust and maturity. The assessment of complex socio-technical systems, such as the one presented, is challenging. Arguably, it will not be complete without an enterprise governance view incorporating a set of dimensions on the level of organizational process performance and maturity. A first view on this perspective has been presented in Becker *et al.* (2011), where a process and maturity model for preservation planning was outlined that was aligned with the IT Governance framework COBIT (IT Governance Institute, 2007). Current efforts are building on this work to develop a full-fledged process and capability maturity model that shall support organizations in systematic improvement of their preservation capabilities[16].

2.5 Summary

This section discussed each of the key design goals of the architecture and system presented in Becker *et al.* (2014) and conducted a quantitative and qualitative evaluation of the key objectives for each of the goals. We showed that the system significantly improves on the existing state of art in digital preservation by combining a

context-aware business intelligence support tool with a scalable mechanism for content profiling, both integrated with a successor of the standard preservation planning tool Plato that is showing substantial efficiency gains over previous solutions. While there are limitations on the scale of content that can be profiled, analyzed and preserved in limited amounts of time, the improvements show that preservation planning and monitoring can be realistically advanced to a continuous preservation management function integrated with operational systems. This will provide a substantial step forward for the many organizations that are looking for ways to enable their repositories for truly supporting the long-term access promise that digital preservation has set out to deliver (Hedstrom, 1998). We pointed out a number of limitations that currently hold back further progress, and outlined current efforts to tackle them.

3. Conclusion and outlook

Ensuring the longevity of digital assets across time and changing social and technical environments requires continuous actions. The volumes of today's digital assets make effective business intelligence and decision support mechanisms crucial in this enterprise. While purely technical scalability of data processing can be handled using state-of-the-art technologies, curators require specific decision support to enable large-scale management of digital assets over time. This demands a set of systems and services that facilitate scalable in-depth content analysis, intelligent information gathering and efficient decision support, designed as loosely coupled systems that are able to interact and connect to the wider preservation context.

This article presented a systematic assessment and evaluation of the SCAPE Planning and Watch suite presented by Becker *et al.* (2014). The results of the assessment demonstrate the possibility to deploy full preservation lifecycle support into preservation systems of real-world scale by adopting a loosely coupled, open and extensible suite of preservation tools that each support particular aspects of the core preservation planning and monitoring capabilities:

- Scalable content profiling is supported by the highly flexible and efficient content profiler C3PO, which has been tested on a dataset of 441 million files.
- Monitoring of compliance, risks and opportunities is supported by the monitoring system Scout, which provides an extensible open platform for drawing together information from a variety of sources to support the much-needed business intelligence insights that are key to continued preservation success.
- Preservation planning efficiency is being continuously improved as the ecosystem grows, and recent advances show that planning can become a well-understood and managed activity of repositories.
- Context awareness of each of the systems is supported by a shared permanent vocabulary set to grow over time through extensions with related ontologies, connecting the domains of solution components and the preservation community with the organizational policies and the decision support and control systems presented here.
- Loose coupling of the components in this ecosystem guarantees that organizations can follow an incremental approach for improving their preservation systems and capabilities.

We discussed the evaluation of key aspects of each tool, as well as the ecosystem as a whole and outlined the key benefits and advances over the existing state of art. Based on a number of limitations, we define a number of key goals for future research. These include real-time profiling of very large datasets in the petabyte range; benchmarking of automated tools against solid, reliable ground truth in open, fully transparent experiments with shared datasets; and a systematic framework for assessing the performance of organizations in terms of process metrics and organizational maturity.

Notes

1. www.openplanetsfoundation.org/blogs/2013-01-09-year-fits
2. www.openplanetsfoundation.org/blogs/2012-11-06-running-apache-tika-over-arc-files-using-apache-hadoop
3. <http://hadoop.apache.org>
4. <https://tika.apache.org/1.4/formats.html>
5. <http://en.statsbiblioteket.dk>
6. www.openplanetsfoundation.org/blogs/2013-01-09-year-fits
7. In Plato, the scoring functions range between 0 and 5, with 0 being unacceptable, and are aggregated across the goal hierarchy. This is discussed in detail in (Becker *et al.*, 2013).
8. www.w3.org/RDF/
9. www.w3.org/TR/owl2-overview/
10. <http://opensource.org/licenses>
11. <https://github.com/openplanets/policies>
12. <http://openplanetsfoundation.org/>
13. <https://github.com/fasseg/scape-tck>
14. www.ifs.tuwien.ac.at/dp/plato/
15. <https://github.com/openplanets/plato>
16. www.benchmark-dp.org

References

- Antunes, G., Borbinha, J., Barateiro, J., Becker, C., Proenca, D. and Vieira, R. (2011), "Shaman reference architecture", SHAMAN Project Report, version 3.0, SHAMAN, available at: www.shaman-ip.eu/sites/default/files/SHAMAN-Reference%20Architecture.pdf
- Basili, V.R., Caldiera, G. and Rombach, H.D. (1994), "The goal question metric approach", *Encyclopedia of Software Engineering*, John Wiley, Toronto, NY, Vol. 2, pp. 528-532.
- Bauer, S. and Becker, C. (2011), "Automated preservation: the case of digital raw photographs", *Digital Libraries: For Cultural Heritage, Knowledge Dissemination, and Future Creation Proceedings of 13th International Conference on Asia-Pacific Digital Libraries (ICADL 2011)*, Beijing, Springer-Verlag.
- Becker, C., Antunes, G., Barateiro, J., Vieira, R. and Borbinha, J. (2011), "Control objectives for DP: digital preservation as an integrated part of IT governance", *Proceedings of the*

ASIST Annual Meeting, American Society for Information Science and Technology, New Orleans, LA.

- Becker, C. and Duretec, K. (2013), "Free benchmark corpora for preservation experiments: using model-driven engineering to generate data sets", *Proceedings of the 13th ACM/IEEE-CS Joint Conference on Digital libraries (JCDL), Indianapolis, IN*, pp. 349-358.
- Becker, C., Duretec, K. and Faria, L. (2014), "Scalable decision support for digital preservation", *OCLC Systems & Services*, Vol. 30 No. 4.
- Becker, C., Duretec, K., Petrov, P., Faria, L., Ferreira, M. and Ramalho, J.C. (2012), "Preservation watch: what to monitor and how", *Proceedings of the 9th International Conference on Preservation of Digital Objects (iPRES), Toronto*.
- Becker, C., Kulovits, H., Guttenbrunner, M., Strodl, S., Rauber, A. and Hofman, H. (2009), "Systematic planning for digital preservation: evaluating potential strategies and building preservation plans", *International Journal on Digital Libraries*, Vol. 10 No. 4, pp. 133-157.
- Becker, C. and Rauber, A. (2011a), "Decision criteria in digital preservation: what to measure and how", *Journal of the American Society for Information Science and Technology*, Vol. 62 No. 6, pp. 1009-1028.
- Becker, C. and Rauber, A. (2011b), "Preservation decisions: terms and conditions apply. Challenges, misperceptions and lessons learned in preservation planning", *Proceedings of the 11th annual international ACM/IEEE Joint Conference on Digital libraries (JCDL), Ottawa*, pp. 67-76.
- Gallagher, M. (2013a), "Improving software sustainability: lessons learned from profiles in science", *Proceeding of Archiving 2013, Washington, DC*, pp. 74-79.
- Gallagher, M. (2013b), "Why can't you just build it and leave it alone?", available at: <http://blogs.loc.gov/digitalpreservation/2013/06/why-cant-you-just-build-it-and-leave-it-alone/>
- Hedstrom, M. (1998), "Digital preservation: a time bomb for digital libraries", *Journal of Computers and the Humanities*, Vol. 31 No. 3, pp. 189-202.
- ISO (2010), "Space data and information transfer systems - Audit and certification of trustworthy digital repositories (ISO/DIS 16363)", International Standards Organisation, available at: www.iso.org/iso/catalogue_detail.htm?csnumber=56510
- IT Governance Institute (2007), *COBIT 4.1 Framework*, IT Governance Institute, available at: www.isaca.org/Knowledge-Center/cobit/Documents/COBIT4.pdf
- Jurik, B. and Nielsen, J. (2012), "Audio quality assurance: an application of cross correlation", *Proceedings of the 9th International Conference on Preservation of Digital Objects (iPRES), Toronto*.
- Kraxner, M., Plangg, M., Duretec, K., Becker, C. and Faria, L. (2013), "The SCAPE Planning and Watch suite", in *Proceedings of the 10th International Conference on Preservation of Digital Objects (iPRES) 2013, Lisbon, Portugal*.
- Kulovits, H., Becker, C. and Andersen, B. (2013a), "Scalable preservation decisions: a controlled case study", *Proceeding of Archiving 2013, Washington, DC*, pp. 167-172.
- Kulovits, H., Becker, C. and Rauber, A. (2012), "Roles and responsibilities in digital preservation decision-making: towards effective governance", *The Memory of the World in the Digital Age: Digitization and Preservation, Vancouver*.
- Kulovits, H., Kraxner, M., Plangg, M., Becker, C. and Bechofer, S. (2013b), "Open preservation data: controlled vocabularies and ontologies for preservation ecosystems", *Proceedings of the 10th International Conference on Preservation of Digital Objects (iPRES), Lisbon*.

- Kulovits, H., Rauber, A., Kugler, A., Brantl, M., Beiner, T. and Schoger, A. (2009), "From TIFF to JPEG2000? Preservation planning at the Bavarian state library using a collection of digitized 16th century printings", *D-Lib Magazine*, Vol. 15 Nos 11/12.
- Law, M.T., Thome, N., Gançarski, S. and Cord, M. (2012), "Structural and visual comparisons for web page archiving", *Proceedings of the 2012 ACM symposium on Document Engineering (DocEng'12)*, New York, NY, pp. 117-120.
- OCLC and CRL (2007), *Trustworthy Repositories Audit & Certification: Criteria and Checklist*, OCLC and CRL, Dublin, OH, Chicago, IL.
- Petrov, P. and Becker, C. (2012), "Large-scale content profiling for preservation analysis", *Proceedings of the 9th International Conference on Preservation of Digital Objects (iPRES)*, Toronto.
- Reimer, N.B. and Moldrup, P.M. (2013), "A case study: content profiling using C3PO", *Internal Technical Report*, SCAPE project.
- Ross, S. and McHugh, A. (2006), "The role of evidence in establishing trust in repositories", *D-Lib Magazine*, Vol. 12 Nos 7/8.
- Sinclair, P., Billenness, C., Duckworth, J., Farquhar, A., Humphreys, J. and Jardine, L. (2009), "Are you ready? Assessing whether organisation are prepared for digital preservation", *Proceedings of the 6th International Conference on Preservation of Digital Objects (iPRES)*, San Francisco, CA, pp. 174-181.
- Webb, C., Pearson, D. and Koerbin, P. (2013), "Oh, you wanted us to preserve that?! Statements of preservation intent for the National Library of Australia's digital collections", *D-Lib Magazine*, Vol. 19 Nos 1/2.

Further reading

- Becker, C., Kraxner, M., Plangg, M. and Rauber, A. (2013), "Improving decision support for software component selection through systematic cross-referencing and analysis of multiple decision criteria", *Proceedings of 46th Hawaii International Conference on System Sciences (HICSS)*, Maui, HI, pp 1193-1202.

About the authors

Christoph Becker is Assistant Professor at the Faculty of Information and Director of the Digital Curation Institute at the University of Toronto, and Senior Scientist at the Information and Software Engineering Group of the Vienna University of Technology in Austria. His research is focused on understanding sustainability in a digital world and enabling better decision-making, especially in the area of digital libraries, digital curation, digital stewardship, digital preservation and digital archiving. It pursues an Information Systems approach to Digital Curation and Preservation, emphasizing sustainable software design, scalable analytics, systematic experimentation and organizational improvement. He completed his doctoral degree in computer science at the Information and Software Engineering Group of the Vienna University of Technology in Austria in 2010. In 2010-2011, he was a visiting researcher with INESC-ID in Lisbon, Portugal. He was involved in DELOS, DPE, PLANETS and SHAMAN, is Principal Investigator of the project BenchmarkDP, and until 2013, he was leading the sub-project Scalable Planning and Watch of the European Commission-funded project SCAPE. Christoph Becker is the corresponding author and can be contacted at: christoph.becker@utoronto.ca

Luis Faria is a Systems and Informatics Engineer at the University of Minho, Portugal. He was part of the original development team of RODA (Repository of Authentic Digital Objects) and has engaged in R&D tasks dedicated to systems design, service-oriented architectures, format migration services and database preservation. Currently, he is the Innovation Director at KEEP SOLUTIONS. He is involved in European research projects focused on digital preservation such

OCLC
31,1

as SCAPE, 4C and EARK, and is pursuing his PhD in digital preservation at the University of Minho.

34

Kresimir Duretec is a Researcher at the Department of Software Technology and Interactive Systems at Vienna University of Technology. He is currently pursuing his PhD at the same department. Previously, he graduated with MSc and BSc in Computer Science from the University of Zagreb in 2011 and 2009, respectively. In 2013-2014, he was leading the sub-project Scalable Planning and Watch at the European Union-funded project SCAPE and is currently a researcher in the BenchmarkDP project. His current major interest is centered around applying model driven engineering principles to the systematic assessment and benchmarking of digital preservation systems.

For instructions on how to order reprints of this article, please visit our website:

www.emeraldgroupublishing.com/licensing/reprints.htm

Or contact us for further details: permissions@emeraldinsight.com