



Rijksdienst voor het Cultureel Erfgoed
*Ministerie van Onderwijs, Cultuur en
Wetenschap*

AN APPRAISAL OF RCE'S VISION FOR A
DUTCH HERITAGE DIGITAL INFRASTRUCTURE

Connected content, connected organisations

Prof. dr. Frank van Harmelen, Vrije Universiteit Amsterdam
Dr. Jacco van Ossenbruggen, Centrum Wiskunde & Informatica

CWI

Centrum Wiskunde & Informatica

VU  **UNIVERSITY
AMSTERDAM**

Table of contents

1.	Introduction	<i>pag. 4</i>
1.1	Background	<i>pag. 4</i>
1.2	Task description	<i>pag. 4</i>
1.3	Approach	<i>pag. 4</i>
1.4	Brief biographies of the authors	<i>pag. 4</i>
2.	Technical Perspective : RCE's Digital Infrastructure	<i>pag. 6</i>
2.1	Architecture	<i>pag. 6</i>
2.2	Linking versus Aggregating	<i>pag. 8</i>
2.3	Statistics versus Semantics	<i>pag. 8</i>
2.4	Infrastructure versus applications	<i>pag. 9</i>
3.	External Perspective: RCE's Role in the Dutch Heritage Landscape	<i>pag. 10</i>
3.1	RCE's role as a hub	<i>pag. 10</i>
3.2	RCE's Relation with DEN	<i>pag. 10</i>
3.3	RCE in relation to the organisations it represents	<i>pag. 10</i>
3.4	Business models	<i>pag. 10</i>
4.	Internal Perspective: RCE as a Heritage Organisation	<i>pag. 12</i>
4.1	Management vision and work floor reality	<i>pag. 12</i>
4.2	Production versus lab environments	<i>pag. 12</i>
4.3	Level of technical expertise	<i>pag. 12</i>
5.	Recommendation summary	<i>pag. 13</i>
6.	Acknowledgements	<i>pag. 14</i>
7.	Appendix: Consulted literature	<i>pag. 15</i>

1 Introduction

1.1 Background

The Netherlands has a rich cultural heritage, and the national collection is considered to be of great economic, educational, cultural, scientific and historic value. This value resides in the tangible and non-tangible artefacts in the collection, but also in the knowledge about the collection, including the relations among objects, within and across collections, and the stories in which these objects play a role.

Substantial efforts have been made in recent years to make this knowledge available in digital form, and to open it up to the public. Traditionally, these digitisation projects have been focused on individual objects. More recently, a broad consensus has emerged that besides this focus on individual objects, the inherent networked structure of heritage information needs to be taken into account to exploit the full potential of its richness.

Since relations often run across collection and organisational boundaries, representing such relations in digital form requires syntactic and semantic interoperability. Internationally, a large number of standards now provide the means to express such relations in an interoperable way, building on top of the existing Web and Internet infrastructure.

1.2 Task description

RCE has the ambition to play an important facilitating role in unlocking the information in such a networked cultural heritage landscape, and has made substantial steps towards the technical and organisational realisation of this vision.

We were asked by RCE management to assess the plans and design of RCEs digital infrastructure for digital heritage. This assessment should concern both vision, approach and implementation. It should cover both short-term and more longer term perspectives, and relate RCE's approach to international developments and their choice for partners and pilot projects. It should discuss pros and cons, from the perspective of the national government, other heritage institutes and science (including e-humanities and ICT).

The essential elements of this infrastructure are, according to RCE:

- a three-tier architecture
- support for single-storage/multiple use, and
- use of open standards.

We were asked to pay specific attention to RCE's approach to the middle tier of the three-tier architecture (the so-called semantic layer) and to its relation with other commonly used GIS architectures and alternative approaches to search.

In addition, we were asked to look at tooling support (Erfgoed suite, DIMCON) for heritage institutions, and its relevance and implications of RCE's infrastructure for the Dutch heritage field.

1.3 Approach

We prepared our assessment by reading the background documents provided to us by RCE (see appendix). We subsequently interviewed a number of key players in the Dutch heritage community:

- Dirk Houtgraaf (RCE, October 4, 2013)
- Kees Hendriks (RCE, October 11, 2013)
- Tjeerd de Boer, Gert-Jan Willighagen (Dutch Ministry OC&W, January 7, 2014)
- Hans Nederbragt (Trezorix, January 15, 2014)
- Marco de Niet (DEN, January 16, 2014)
- Johan Oomen (Sound & Vision, January 17, 2014)
- Elsbeth Kwant, Elco van Staveren (National Library, February 7, 2014)

All interviewees were keen to participate and we thank them for their open and frank attitude during our conversations.

A first version of this report was checked for factual errors with the interviewees. An updated version was then submitted to RCE for feedback, leading to our final version.

This report reflects the opinions of the authors, and not necessarily those of RCE or any of our interviewees.

1.4 Brief biographies of the authors

Frank van Harmelen (1960) is professor of Knowledge Representation and Reasoning at the VU University Amsterdam, where he is scientific director of The Network Institute (<http://www.networkinstitute.org>). Van Harmelen has been active in the development of the Semantic Web since its inception. He was co-PI on the first Semantic Web project in Europe (OnToKnowledge, 1999), which laid the foundations for the Web Ontology Language OWL. Van Harmelen wrote the first academic text book in this field, the "Semantic Web Primer" which has been translated into Chinese, Japanese, Korean, Spanish and Greek. On the technological side, Van Harmelen was one of the architects of Sesame, one of the first RDF storage and retrieval engines, which is in wide use in academia and industry. The work on Sesame earned him and his co-authors the "10 year impact award" of the International Semantic Web Conference. He was scientific director of the Large Knowledge Collider (LarKC), a platform for distributed computation over very large semantic graphs. He is the author of some 200 scientific publications, many of which are highly cited.

Jacco van Ossenbruggen is a senior researcher with the Information Access group at the Centrum Wiskunde & Informatica (CWI) in Amsterdam, and affiliated as an associate professor with the Web & Media research group at VU University in Amsterdam. His research interests include semantic web interfaces for cultural heritage and other linked open data. He was involved in the award-winning MultimediaN E-culture project and the early drafts of the Europeana Data Model. He currently serves on the external advisory board of the linked open data initiative of the Getty Research Institute and in various projects investigating crowdsourcing of expert tasks in cultural heritage applications.

2 Technical Perspective : RCE's Digital Infrastructure

2.1 Architecture

RCE's vision is to provide a digital infrastructure that fully supports and exploits the networked nature of heritage information, and to make this networked information available using interoperable, open and commonly agreed upon formats.

In this vision, cultural heritage institutions are not (or: not only) owners of heritage content, but they are nodes connected in a network of semantically rich links between content sources. Such a vision has direct consequences for the infrastructure in which these institutions and their content are connected.

RCE's infrastructure for such a network of heritage content and organisations consists of three different levels:

1. In the first layer, individual collection owners and other heritage organisations make their curated digital information available using open standards. Content is annotated with vocabularies (thesauri, ontologies) that are available on the second layer.
2. In the second layer, vocabularies are made available for content-description, and cross-collection links can be created,

curated and retrieved (including links between collections from different organisations).

3. On the third layer, different parties can build end-user applications, accessing objects from layer one and the links between from layer two. Such applications can be built by heritage organisations themselves, but also by third party organisations, be they public organisations or third party private organisations which make a business by (re-)using public cultural heritage material

This entire architecture is built by using the Web as a network infrastructure, and by using open standards for content and communication.

This architecture encourages cross-collection linking in a "vertical" manner, via vocabulary terms that can be used in search queries and for interlinking. Such links are more reusable and cost effective than the many direct "horizontal" pairwise direct links between collections that would otherwise be needed (and that are currently common).

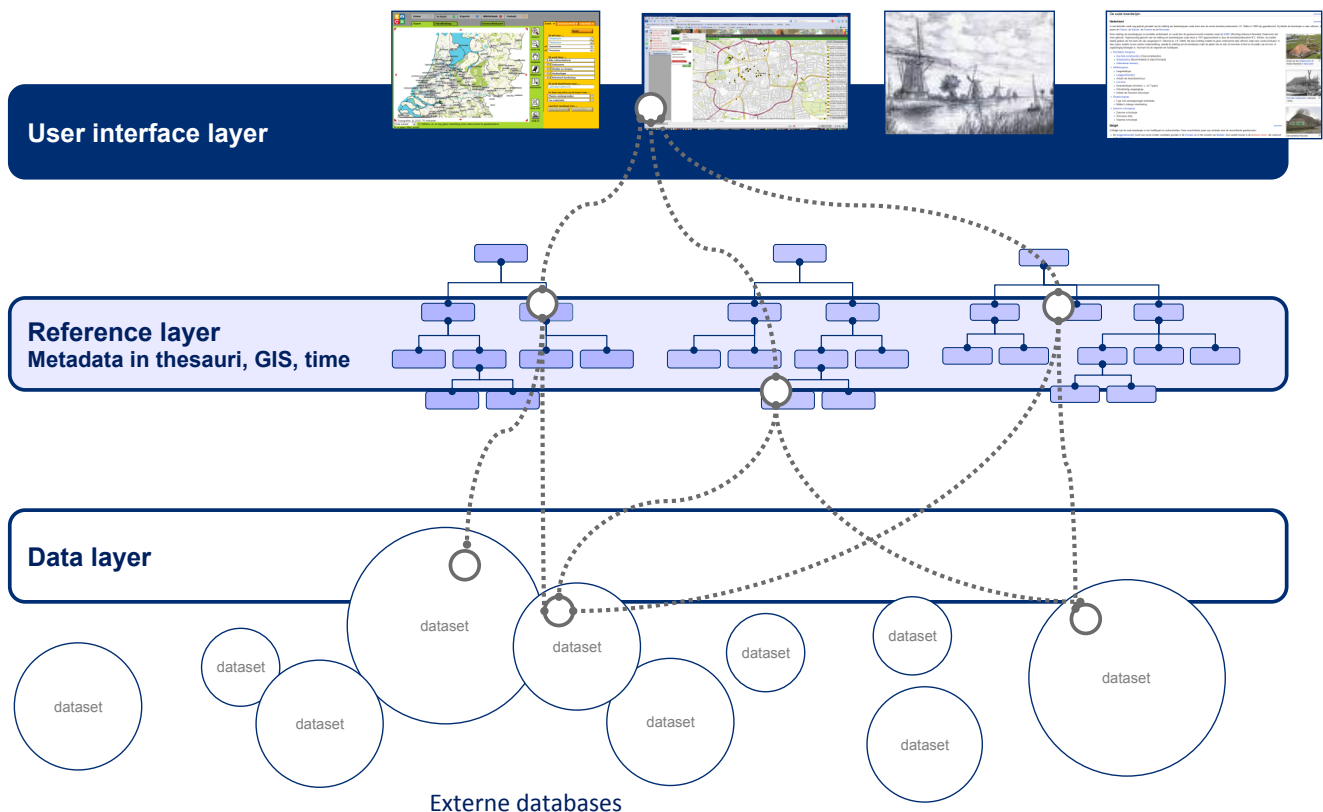


Figure: RCE's basic three tier architecture (adapted from slides provided by RCE)

Assessment and recommendations

Based on our best technical judgement and on our knowledge of the state of the art worldwide, we strongly support the basic design principles of the RCE architecture. It is a vision that is squarely positioned in a vision about an open, networked world, a vision similar to the one that underpinned the success of the World Wide Web.

The three-tier architecture of content (on layer 1), semantic-interoperability (layer 2) and applications (layer 3) is broadly accepted and has been used successfully in a wide variety of settings. The use of the web and of open standards ensures interoperability and avoids vendor lock-in. The following is an illustrative (not exhaustive) list of the successful adoption of these architectural principles worldwide in a variety of sectors, and illustrates that these standards are now sufficiently mature and up to their job:

- e-commerce: GoodRelations¹, schema.org
- public administration: data.gov.uk, data.gov, data.overheid.nl
- news and media: partly data-driven websites from e.g. BBC, NYT
- scientific data publishing: most examples are from the biomedical domain (e.g. on layer 2: bioportal², SNOMED CT³, GO⁴, with on layer one: clinical trial data, research papers, drug descriptions, adverse event reports, genomics data etc)

The architecture is based on standards that are stable and will remain in use many years in the future. At the same time it is sufficiently flexible to adapt to new innovations and internet protocols, as well as future standards and best practices that might emerge in the cultural heritage community. Technical tooling to support this architecture and the underlying open data models are increasingly becoming available on the national and international market, although not yet as well established as more traditional solutions based on XML or relational data models.

In some aspects, the architecture proposed by RCE is a step ahead of the current metadata practice of other key players in the field. These other hubs and (inter)national content aggregators have focussed their efforts on large-scale aggregation of object-centric metadata. These organisations also recognise the need for interlinking, and have also started to investigate how to support the semantic metadata interoperability across collections that has been the core of RCE approach from the start. We feel other players in the field can learn from RCEs approach to improve their own activities on this topic.

¹ <http://www.heppnetz.de/projects/goodrelations/>

² <http://bioportal.bioontology.org/>

³ <http://www.ihtsdo.org/snomed-ct/>

⁴ <http://www.geneontology.org/>

While we are strongly supportive of the RCE vision, we see a number of challenges that have to be met in realising this vision, both organisational and technical.

1. The RCEs architecture has potential to be demonstrated in a broader set of use cases. This would help to increase potential support and buy-in from other players in the network.
2. A strong aspect of the proposed architecture is its vendor neutrality. RCE can capitalise on this by involving a larger number of vendors.
3. RCE's link-centric vision depends on a widely adopted practice of persistent identifiers (PIDs) for heritage objects and metadata terms. We recommend that RCE pro-actively adopts such practices itself and that it provides know-how to other organisations who need a cost-effective PID service.
4. Solving versioning issues would increase the trust of other organisations in the linked data approach. OpenSkos and OAI-PMH partly solve the publishing side of this problem, but unfortunately the consumer side of this problem is still largely unsolved at the time of writing, even in the academic community
5. The openness of the semantic middle layer is crucial in the architecture. We recommend to make decisions on each of the following issues:
 - a. who, from which organisation, has access to the cross-collection parts of the network? Who is responsible and who feels responsible for what part of the network?
 - b. how to prevent that changes that are beneficial in one scenario are detrimental for another?
 - c. how to resolve disagreements between contributors on alternative definitions of a concept or of relations between concepts?
 - d. how should new thesauri be incorporated, in particular when the data model of a new thesaurus does not directly map onto the data model of the infrastructure (as is the case with the poly-hierarchical nature of the commonly used Getty AAT and RKD's IconClass)? A similar question arises for system-level integration with the OpenSkos platform.
6. The "store only once at the content owner" principle has many advantages (as opposed to "duplicate at the aggregator"). Realisation of this principle requires sufficient expertise from the (often small) content owners. RCE's "Erfgoedsuite" plays a crucial role as a content management and storage-as-a-service solution. We recommend a critical analysis, leading to improvements and an increasing uptake of this platform. We also recommend that future releases of this platform are developed in close collaboration with a larger number of potential users to ensure compatibility with their expertise levels.
7. The aim for transparency for collection owners and one-time storage in the architecture contrasts with complaints in the

field about untransparent aggregation, mapping and update policies in the DimCon aggregator.

2.2. Linking versus Aggregating

RCE's vision is centered around frictionless **linking**: the notions of interoperability on a syntactic and semantic level, by using open standards and symbolic links that are meaningful for both human and software agents. The physical storage location of the data is typically decentralized, and uniform access to this distributed data is realized by combining the interoperability links with a good computer network infrastructure.

Our impression is that the daily practice of many of the other hubs in Dutch Cultural Heritage has been centered much more around the notion of frictionless **aggregation**, as opposed to linking. Here the efforts have focused on the creation of an infrastructure where the same open standards are used to enable information to flow easily from the organisations creating it to various domain, thematic or geographic aggregation levels, all the way to national and international aggregators. In this approach, creating an easy way to get information physically at a central place is the primary objective. Once this information is there, getting it meaningfully linked is an important, but secondary step.

While in the long term, semantic linking is crucial in both scenarios, not all interviewees shared that conviction, and some crucial issues remained unclear. For example, in both the aggregation and the linking scenario, it remains currently unclear (a) who is best qualified to do the linking work, and (b) who pays for the linking work. Many hubs envision a per-hub aggregation-model with only a very thin and cheap infrastructure to connect the largely disjoint hubs, and where linking (within and between the hubs) is of secondary importance. To quote one of the other hub representatives: "most of the work should be done within the hubs". This contrasts with the RCE vision, where most of the work, both technical and intellectual, is in building and maintaining the national link infrastructure.

Assessment & recommendation

The RCE vision is clearly "linking not aggregation" (quote from RCE's mission statement on once-only storage of data). In our professional judgement, the "linking" vision will in the long run prove to be more flexible and scaleable (both technically and organisationally):

- Content owners will require support to feed enrichments of their content done on the diverse aggregation levels back into their own collection. However, the current aggregation infrastructure and organisation are mainly directed towards upwards data flows. The store-once/linking approach envisioned by RCE

better supports such needs.

- In the RCE model, the primary storage remains with the collection owners. This model resonates better with the sense of ownership and protection that many collections owners have, which could improve their commitment to digital curation. Additionally, this model is more flexible when dealing with other licensing models than the liberal CCo license that is currently common in the aggregator model.
- The linking model is less sensitive to the technical single point of failure risk than most approaches based on central aggregation

Nevertheless, aggregation is currently more mature (again, both technical and organisational), and most short term gains have been made using the aggregation approach. This difference in maturity is also visible when looking at the products from commercial IT suppliers. Nevertheless, the basic aggregation infrastructure is ready, and there is an emerging need for linking data that has been aggregated by the other hubs. A possibility is a combined linking/aggregation scenario, for example a "centralised" index on top of a distributed network of content⁵. We therefore recommend:

8. a strategy that allows, depending on the use case, deployment of the most effective approach (e.g. the link-centric or aggregation-centric).
9. a migration plan to support transitions from aggregation to linking, or a combination of the two.
10. We recommend a proactive approach from RCE to align linking projects of aggregated content with the ideas about linking in RCEs store-once model. For example, controlled vocabularies and thesauri play a crucial role in the linked architecture envisioned by RCE. RCE could initiate a practical study to investigate which thesauri (that are currently in use) could actually improve the functionality of the different layers in the both the linked and aggregation models.

2.3 Statistics versus Semantics

In both the scientific community and the library community, there is a history of different approaches to search problems. These different approaches can be characterised as:

- the use of manual indexing versus automatic full-text indexes
- the use of search terms from controlled vocabularies and thesauri versus the use of arbitrary keywords and phrases as common in Web search engines
- the use of statistical inference versus semantic inference techniques

⁵ Note that this is precisely the architecture of current web search engines

These approaches are typically seen as mutually exclusive. However, there is growing consensus that these approaches should **not** be seen as mutually exclusive, but rather as complementary. Examples include the Knowledge Graphs extended search results as currently developed by Google and Yahoo, and the structured searches on top of Facebook's social graph, etc.

First, there are collections for which one approach is more suitable than the other. For example, there is sufficient evidence in scientific literature that for ad-hoc queries on large full-text collections, automatic approaches based on statistical analysis of the collection outperform manual approaches. At the same time, fully automatic retrieval techniques are not sufficiently mature to support search on images if these have not been adequately indexed manually.

Second, there are tasks that can only be supported by an innovative combination of both approaches. Examples include semantic inference based on automatically extracted concepts or named entities, and the improvement of statistical techniques for information extraction techniques by using semantic background information to deal with homonyms and synonyms, or to support query narrowing and broadening.

Assessment & recommendation:

In our perception, views on this topic in the heritage field still assume that the approaches are mutually exclusive instead of complementary. As a result, RCE is solely focussed on the semantics and thesaurus-based approaches, while other hubs are mainly focussed on the statistical/full-text approaches.

11. A better understanding of the needs of the other partners, and an understanding of the complementarity of the two approaches are needed, as part of closer collaboration between the hubs.
12. RCE will need access to expertise on methods and techniques from Information Retrieval and Machine Learning to decide when these are useful for RCE's own data, and how they can be combined with semantic approaches to improve the functionality of the RCE digital infrastructure.

2.4 Infrastructure versus applications

There is a wide consensus among our interviewees that while there is much attention for digitisation and online access to the collections (i.e. the data layer) and for semantic interoperability and data aggregation (i.e. the semantic layer), there is less attention for developing end-user applications. In addition, a very crisp view on those end-users and their needs is currently lacking with many of the players in the digital heritage field. We expect that

end-users will differ between hubs, and that each set of end-users comes with its own use-cases and requirements. Many of the potential users will be users of more than one of the hubs, and any cross-hub infrastructure should be informed by their needs and the technical requirements of the (future) end-user applications that will be developed on top of this infrastructure. A clear picture of who these users are, and what functionality they will require from each individual hub and from a cross-hub infrastructure is not always clear, both at RCE but also at the other hubs. Other users might use only RCE as a hub. But also for these users, it is not always clear which urgent or important problems are directly addressed by the infrastructure. This is especially important because in the vision of RCE, there is an important active role for cultural heritage institutes and other layer 3 users, because these should use RCE's open APIs to develop their own applications. Although we subscribe to this vision, one should not underestimate the incentives and the amount of guidance that are required to get this process off the ground. These lessons have also been learned the hard way by the Linked Data community at large.

Assessment & recommendation

Getting a better understanding of user requirements is by no means trivial. In many cases, these users themselves might not yet know what their requirements would be. Nevertheless, we see this as an urgent and important ingredient that is required to realize RCE's vision. Attention to the needs of applications is important from a technical perspective (since it drives the technical requirements), from an organisational perspective (since it helps to prioritise activities and allocate resources) and from a strategic perspective (by showing early gains). For example, there's a feeling with some that the current "windmill demo" is beginning to wear out, and that new exciting showcases are needed to show added value. As another example, we applaud the early release of the recently launched "Doorzoeker⁶" web application as a showcase of the possibilities of RCE's infrastructure. However, although impressive, we still see it as a product that is born mainly from data- and technology-push, and rather little from demand pull.

13. We recommend more attention and sensitivity to applications that will use the linked digital infrastructure. Proposals for future pilot projects and showcase demonstrators such as "Doorzoeker" should be strongly motivated by concrete use cases and added value for users.
14. Convincing examples of use cases are much needed, both to sell the infrastructure "downstream" and increase uptake, but also to convince parties "upstream", e.g. regional and national government and the other hubs.

⁶ <http://doorzoeker.cultureelerfgoed.nl/>

3 External Perspective: RCE's Role in the Dutch Heritage Landscape

3.1 RCE's role as a hub

All interviewees share RCE's vision that Dutch heritage collections need to be widely linked and commonly accessible via open standards. In addition, all acknowledge the role of RCE as one of the five major hubs in the Dutch heritage landscape, together with the National Library, National Archive, Institute for Sound and Vision and DANS. Several interviewees stated that all important decisions in the future should have broad support from the "big five" hubs and the institutes they represent, and that RCE should be a part of this. Some voiced concern whether RCE was playing its role as equal player among these big five. This leads us to a concern whether RCE as an organization is sufficiently connected to realise its vision of connected collections. Some of the other hubs have a better representation than RCE in several collaborative projects, both on a national and an international level, and they are more successful in deploying Europeana-related funding, both to improve their national visibility and to improve their collaboration with their sister institutes abroad. Now that RCE has much of its infrastructure in place, this is a good time to increase RCE's international connectivity.

Several of our interviewees mentioned RCE's dual role: both as a hub being controlled by the Dutch government like all other hubs, and as a hub that is also directly a governmental organisation itself. Some interviewees were of the opinion that RCE should focus on the data, and preferably not do internal software development projects at all. However, we appreciate RCE's desire to take technical progress in their own hands, and we encourage them to continue to play this dual role, with an eye to collaboration and open licensing.

Recommendations

15. The synergy between the big hubs needs to be improved. RCE can be more pro-active in seeking cooperation. Concrete examples which would improve inter-hub collaboration on a national level are: increasing participation in OpenSKOS, higher priority for persistent identifiers, higher priority for explicit licensing policies and an intensified collaboration with Geonovum on geospatial services.
16. On an international level, RCE can improve its representation in (the many projects related to) Europeana. Although on an operational level the approaches taken within Europeana might not always be the same as those taken by RCE, on a larger scale we think the long term European vision of Europeana matches the national vision of RCE.
17. RCE's dual role requires careful balancing, especially when promoting software developed by commercial third-party developers, or when RCE is directly involved in establishing quality certification programs. Components of a common

digital infrastructure that are developed, need to be available under an open license.

3.2 RCE's Relation with DEN

It is important to have a clear and shared understanding of the separation of tasks between the two organisations. In the past, DEN used to have a role as a "watchdog" organisation for data and api standard adherence and other IT quality issues. RCE seems to have taken over part of that role, but it is unclear how much support that transition has in the heritage field, and to what extent RCE is in the right position to play that role.

Recommendation

18. A watchdog organization with tasks such as setting up software certification programs, should be a truly independent organisation that represents the entire field, across all hub boundaries.

3.3 RCE in relation to the organisations it represents

There is broad agreement with RCE's observation that many organisations in the heritage field are too small to have sufficient in-house IT knowledge to implement effective access to their collections, and that these organisations need help, for example by offering the required tools as a web service which is implemented and maintained elsewhere. However, there is less consensus about what role a government organisation such as RCE should play in this, and what such services should look like. For example, the ErfgoedSuite is perceived by some as too complex, resulting in a slower uptake. It is not always clear to what extent the organisations represented by RCE share the RCE vision and to what extent they are willing to commit resources to help RCE to realise its vision.

Recommendations:

19. RCE should look into ways to "educate" the organisations in the field, and to aim for more direct involvement of these organisations, both on the management level, and (especially) by direct involvement in the IT services that are under development.

3.4 Business models

- The different parties have rather divergent and sometimes vague ideas on how to pay for the implementation and maintenance of an national, cross-hub infrastructure as proposed by the RCE view. Several alternative business models have been discussed during our interviews:
- Usage model: users of the 3rd layer (eg app-builders) pay for costs of the 1st and 2nd layer
- Public utility model (paid from tax income): the government

pays for 2nd layer because it is a generic infrastructure that yields economic growth and reduces costs of other public administration tasks. Examples of cost reduction for public administration that were mentioned include: cheaper for small heritage organisations to expose their collection, cheaper management of monuments, cheaper and faster decision processes on the city-council level (e.g. on historical regions), cheaper and more effective creation and curation of scientific and educational resources, and cost reduction for central government for other digital heritage projects

- Membership model (i.e. paid by the organisations that need the infrastructure). Currently these would be the hubs, potentially augmented by government subsidy as in the SURF model.

A related issue is that current copyright law is blocking much content-reuse. It is particularly ineffective in a “linking” scenario. While the authors of this report see an important role for the Dutch national government in this issue, there seems to be little consensus on this issues among the people we interviewed.

Recommendation

20. It is crucial that all hubs, in collaboration with the Dutch government, develop a vision on how to fund the infrastructure that links the Dutch heritage collections, its governance model, and the licensing of both its contents and its software.⁷

⁷ As authors of this report we refrain from a specific recommendation among the above options.

4 Internal Perspective: RCE as a Heritage Organisation

As is clear from the above, RCE has very high ambitions. In recent years, RCE has gone through a phase of designing, building and prototyping. Now seems to be the time to start a transition to a next phase. We discuss organisational choices to be made by RCE.

4.1. Management vision and work floor reality

During our interviews, praise for the vision of RCE's management was often followed by concern about the capability of the organisation to execute that vision. People both inside and outside RCE expressed concerns about a possible gap between the long term vision and the day to day concerns of the staff that need to execute that vision. Whether true or not, there is a risk that such a perception could grow into a self-fulfilling prophecy.

4.2. Production versus lab environments

RCE's ambition requires an organisation and funding that can execute short-lived, light-weight, flexible software projects that work towards a specific goal in a specified period. RCE current organisation makes it difficult to execute fast innovative projects, to quickly develop proof-of-concept prototypes and demonstrators in-house, or to quickly test and evaluate innovative technology developed by other parties.

The current technical and organisational skills inside the RCE organisation are understandably focussed on reliability and uptime of long-running services. Ideally, RCE needs to develop an organisation structure in which it would be possible and even stimulated to develop and test new technology in an environment that is independent from the current long-running production-level services. That is, new services need to be developed and/or tested without risking uptime, reliability, or other quality aspects of existing services.

21. We recommend, if at all possible, to set up an R&D team with skilled developers who have experience in modern agile development techniques and lead by project leaders that have experience in managing short-lived development projects. Preferably this is done in collaboration similar labs at other hubs such as Sound & Vision and the National Library

4.3. Level of technical expertise

In addition to a flexible project organisation, a highly skilled technical staff is necessary to realize RCE's IT ambitions. IT has become part of RCE's "core business", and can no longer be outsourced completely. Skilled IT personnel is needed to run the pilot projects discussed above, but also to collaborate with third party software vendors on a more equal level. Note that the latter is perceived as a problem in many governmental organisations, and especially in the cultural heritage field.

While the main ambition of running pilot projects should be to develop innovations that become essential parts of RCE's future infrastructure, such integration is not the only goal of pilot projects. As a national hub and center of expertise, it is essential that RCE keeps up to date and has hands on experience with a wide range of technologies that are, or might become relevant in the field. This is particularly important now that RCE's "three layer architecture" is increasingly adopted by OCW and by other players in the field. Given the nature of the collections RCE deals with on a daily basis, it is not surprising that geospatial modeling is a central ingredient in RCE's middle layer, and should therefore be included in RCE's skillset. Practical expertise is valuable, even if a pilot project ends with the conclusion that the results are not (yet) suitable to be transferred to practical applications.

22. We recommend that to enter the envisioned next phase, RCE should further increase the IT skills of RCE staff at the different levels in the organisation, either through new hiring, if at all possible, or through training.

5 Recommendation summary

We strongly support the design principles of the RCE architecture, its basis in a vision about an open, networked world, its three-tier model and its emphasis on open Web standards. We also support RCE's link-centric, store-once model and think that in the long term, this has better potential than the centralised aggregation model.

Substantial progress has been made by RCE in gathering support for this architecture, and in building the general infrastructure and generic tool support.

The time has now come to bring this vision to fruition. Important next steps should be

- (a) demonstrating the added value of applications of the infrastructure in concrete use-cases, developed in close collaboration with the stakeholders, and
- (b) starting discussions with the national government, the five main hubs and end-user organisations to come to a sustainable business model that can fund the longer term development and maintenance of a digital heritage infrastructure that can cross the boundaries of the individual organisations involved.

These next steps towards a national networked digital infrastructure for cultural heritage will also place new demands on RCE as an organisation. Our recommendations from this report can be summarised in the following three maxims:

Further strengthen support for the users of RCE data and services, that is, the organisations using level 3 of the three-tier model. Development activities of RCE infrastructure should increasingly be driven by a strong pull from level-3 stakeholders, especially those who have a clear financial interest (for example decision makers that require fast and effective management information on archeological sites or monuments).

In the same vein, RCE can capitalise even more on the grass-roots potential of the many enthusiastic people working in the organisations that own the collection data (i.e. using level 1 of the three-tier model). With a clear vision on the needs of the level-1 stakeholders, RCE can become a national center of expertise that can push best practices and lessons learned to smaller heritage organisations.

Invest in RCE's internal organisation: RCE has ambitious goals for creating a non-trivial IT infrastructure. We recommend that RCE invests both in its human capital and in its organisational structure to fit this ambition level. The challenge for RCE will be to strike a balance between developing prototypes to demonstrate advanced innovations, and managing production-level infrastructure. RCE should set up a project organisation aimed at initiating and over-

seeing fast, light-weight and short-lived innovative pilot projects that can operate, to a large extent, independently from the daily production-level activities.

Further strengthen the partner network: It is crucial that activities around RCE's digital infrastructure are perceived as activities where RCE acts with sufficient collaboration and support from the other hubs and the national government. Collaboration with a variety of commercial partners is crucial to avoid vendor lock-in. It is important that RCE pro-actively seeks concrete collaborations with the other hubs, both with respect to the activities it organised by RCE itself, as in national and international collaboration projects organised by others.

In conclusion, building a networked infrastructure also requires operating a networked organisation, both networked with users (actual and potential) and with peer organisations (national and European).

6 Acknowledgements

We thank all interviewees for their open and frank cooperation, and for their feedback on earlier versions of this document. Thanks to Shelley de Jong-Versluys from RCE in giving us access to all the documents we needed. In addition, Lizzy Jongma from the Rijksmuseum and Joop Vanderheiden from RCE provided us with valuable information.

7 Appendix: Consulted literature

- Beleidsbrief Modernisering Monumentenzorg, <http://www.cultureelerfgoed.nl/sites/default/files/u10/beleidsbrief-momo.pdf>
- Rapportage Virtuele Informatie Centra, http://www.abcultuur.nl/Eelco_Bruinsma_Test_Virtueel_Informatie_Centrum_Naturalis.pdf
- “Impactanalyse visie informatievoorziening Cultureel Erfgoed”, Geonovum, January 2011
- Draft of “Voorstel voor een digitale infrastructuur voor cultuur, media en wetenschap”, by Bram Voermans and Gert-Jan Willighagen, November 2011.
- “De modernisering van de monumentenzorg”, Rijksdienst voor het cultureel erfgoed, 1, winter 2010
- RCE Thesauri, Versie: 0.2 (internal draft for discussion only, dd. 7-2-2012.
- Documentatie RNA API 2, Trezorix (Sander van der Meulen), dd. 25-02-2013
- “Werk Document Cluster Erfgoedthesaurus: Omschrijving taken en werkzaamheden van het cluster”, dd. 25 januari 2013.
- “Geactiveerde businesscase project Kennis Infrastructuur Modernisering Monumentenzorg”, Nico Verbeij and Kees Hendriks, 31-5-2013
- Slidedeck “RCE Informatiebeleid” , “2011-05-24 Informatiebeleid.pptx”



Colophon

Cultural Heritage Agency of the Netherlands
P.O. Box 1600
3800 BP Amersfoort
The Netherlands
info@cultureelerfgoed.nl
www.cultureelerfgoed.nl