FEATURE
**EPUB 3:
NOT YOUR
FATHER'S
EPUB**

NISO REPORTS
**THE EVOLUTION
OF ACCESSIBLE
PUBLISHING**

OPINION
DRINKING
THE E-BOOK
KOOL-AID
IN A LARGE
ACADEMIC
LIBRARY

ISQ

SPECIAL EDITION: VIEWS OF THE E-BOOK RENAISSANCE

SPOTLIGHT
**THE CHALLENGE FOR
STANDARDS IN THE
E-BOOK SUPPLY CHAIN**

FEATURE
**E-BOOK
PUBLISHING
FOR
UNIVERSITY
PRESSES**

IN PRACTICE
**E-BOOKS
AND THE
PUBLIC
LIBRARY**
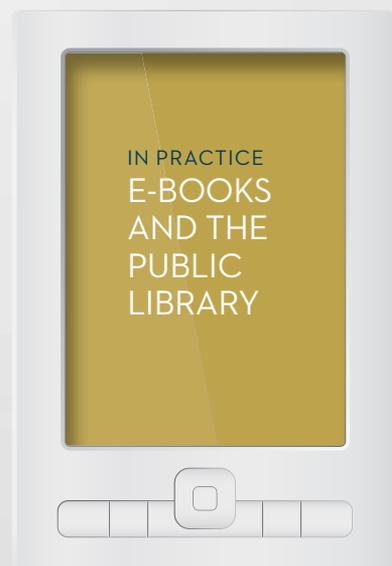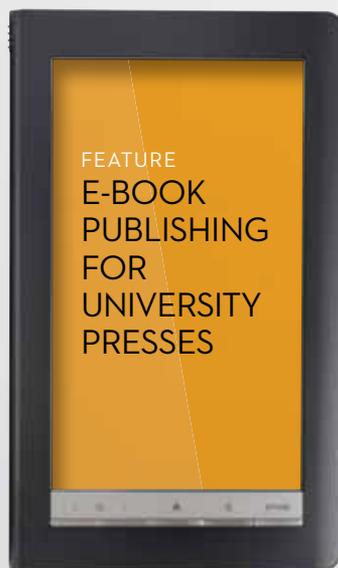
NISO
How the information world
CONNECTS

NR
[NISO REPORTS]
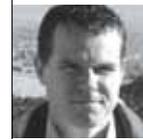
Matt
Garrish

Markus
Gylling

MATT GARRISH AND MARKUS GYLLING

The Evolution of Accessible Publishing: Revising the Z39.86 DAISY Standard

As you read through the revision of the DAISY standard, *Authoring and Interchange Framework Specification* (NISO Z39.86-201x) and its profiles, rest assured it's no coincidence that the markup examples are drawn predominantly from works of Charles Darwin. This specification has undergone a radical transformation since the Working Group began its work in the Fall of 2008, and the result represents a significant evolutionary leap forward in accessible content production.

Darwin's *On the Origin of Species* was consequently selected as the primary source for examples as a quiet nod to the principles of adaptation and evolution that this specification has embodied over the years. This new revision represents a whole new way of looking at the parallel-publishing model in particular, and at content model creation in general, but wouldn't have been possible if not for its predecessors on the road to universal accessible publishing.

of the specification quickly became the de facto standard for talking book production by libraries and organizations serving blind, dyslexic, and other print-disabled readers.

Although a very effective and time-tried specification (it remains in use by many producers to this day) and the one that ushered in the age of digital talking books, the community creating and using these books also had a need

to generate other formats from their text data. This task of reformatting documents was often a repetitive one that involved a combination of machine and human intervention. Producers were increasingly looking to their DAISY text files as the source for these conversions, to leverage the cost and effort that had already gone into converting the original documents they represent into digital form. But while HTML is fine for the Web, it didn't take

A brief history of Z39.86

Accessible format production has come a long way since the first *DAISY Digital Talking Book (DTB)* specification was developed in 1997. That early format combined the HTML and SMIL (Synchronized Multimedia Integration Language) standards to create a synchronized multimedia experience that was ahead of its time, and after a few early revisions the 2.02 version

TIMELINE OF THE DAISY STANDARD
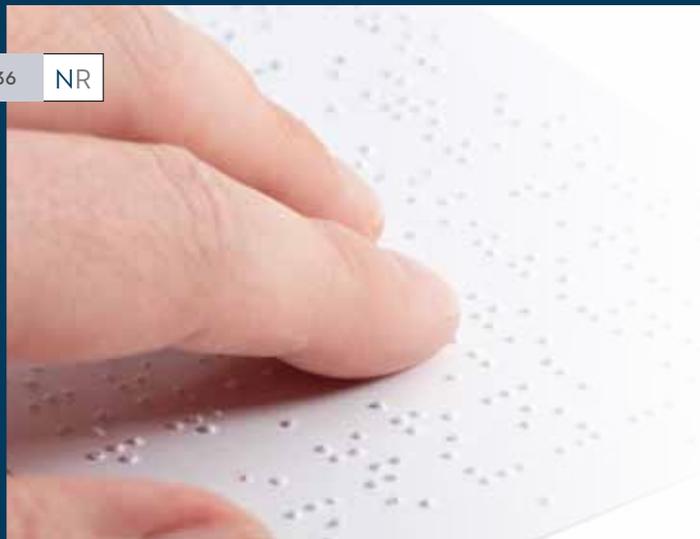
First DAISY specification — 1997

DAISY 2.02 (DAISY 2) — 2001

ANSI/NISO Z39.86-2002 (DAISY 3) — 2002

ANSI/NISO Z39.86-2005 — 2005

NISO Z39.86-201x (Draft for Trial Use) — 2011

While the markup that it provided proved efficient in many authoring contexts, it was insufficient in others; the requirements for formats like braille and large print were still not adequately addressed for all producers.

long before it became clear that a more structured format with better facilities for targeting outputs was going to be needed to enable multi-format production. (Attempting to generate print braille using the small tag set HTML makes available can prove no small feat, for example.)

From this need was born the ANSI/NISO Z39.86-2002 standard and its subsequent 2005 revision. The new text component of these versions of the standard was defined by the DTBook DTD, which built on the original HTML core but added significant improvements for structural and semantic fidelity. As DTBook was deployed by producers of accessible content across the globe, it clearly showed how producers could benefit from XML-based single-source production, and how end users benefit from textual content that is well-structured and semantically coherent.

The wind of change

But although DTBook again improved the production landscape, it brought forward with it the specification's legacy of talking book production and the Web. While the markup that it provided proved efficient in many authoring contexts, it was insufficient in others; the requirements for formats like braille and large print were still not adequately addressed for all producers. Its *book-centrism* and limited mechanisms for adaption and specialization additionally meant that it was inadequate to handle all the document types and regional requirements of producers.

Meanwhile, in the end user context, DTBook as a distribution format anticipated browser vendors moving to accommodate display and rendering of arbitrary XML grammars. This shift—which was seen as a given a decade ago when DTBook was originally created—never materialized, proving a major complication for the visual rendering of talking books. Further, the DAISY distribution format (as of Z39.86-2005) lacked several highly requested features such as

interactivity and better support for East Asian languages.

By 2008, it had become clear that the text component had to be thoroughly revised and cleanly separated from the talking book format if it was going to meet the multi-format production needs that the community was clamoring for, and the distribution format needed to be re-aligned with industry standards.

Evolution in action

The first decision made in undertaking the revision of the 2005 standard was to adopt the principle of *separation of concerns*: to split the incongruous parts in order to isolate and better tackle the problem domains. A new XML authoring standard would be developed to address the accessible text production needs of the community, while a distribution format—a more linear continuation of the previous standard—would focus on talking book production.

The next critical decision in designing the new text standard was that the past would not be a guide to the future. A radical departure was instead needed if the DAISY Revision Working Group was going to be successful. To this end, it was decided that creating a specific markup grammar was not going to be the primary goal of the revision.

While this might seem like a strange objective for a text standard, the group had to invert the production problem and look at it from a fresh perspective. The single monolithic format approach had so far failed to address the needs of the broad community DAISY serves, providing neither the structural and semantic richness nor the flexibility to accommodate the wide array of formats producers had to be able to generate. To begin developing yet another such standard would be to head down an evolutionary dead end.

To fully realize the benefits of a parallel publishing model, a true master source was needed that provided a content model that wasn't hampered by the formatting inherent in being embedded in a specific output or being designed for

a single use. It was envisaged that this new specification would define a common framework in which to develop new grammars, allowing the standard to be adapted to any document type it had to address, instead of the other way around. The framework would specify the technologies to use and define a universal markup core for all documents, but would stay silent about how to structure any given type of document: the structure would be left to be defined by profiles created according to the rules laid out in the framework.

This approach would provide the increased flexibility that producers were requesting to allow markup to be tailored to their unique needs. Understanding that the community shared the same core markup requirements and that the incompatible differences largely related to output production requirements, a common framework was seen as a means of allowing producers to work collaboratively on profiles that fit their shared needs, or to strike out on their own but in a manner that still kept their core markup in line with the wider community. This consistency was going to be key to adoption in a community moving in the direction of a global library, where knowing the differences in markup between any two profiles, producers could still easily exchange and transform their documents to account for the discrepancies.

But this model is intended to be useful beyond just accessible publishing. Too often, the only solution when marking up new document types is to either: a) create a whole new markup model from the ground up; or b) find the closest fitting language and hack a solution over top of it. The new Z39.86 model encourages new profiles to be developed by any interested parties for the benefit of the whole community, sharing knowledge and enhancing the existing foundation to expedite the process. Although initially targeted at the accessible publishing community, the markup is designed to capture the full structure and semantics of the information resources being described, allowing any formats to be generated from them. Adoption of Z39.86 beyond its traditional base is key to making publishing open to all, and it is hoped that all organizations with similar cooperative markup needs and goals will benefit from the work put into this specification.

The next goal of this revision was to move the DAISY standard away from the legacy DTD approach that had persisted from its earlier HTML days. The XML ecosystem has largely outgrown DTDs, and their lack of native namespace and datatype support alone made them an incompatible choice for the direction the group was heading. W3C XML schemas were also discounted as the right technology for defining the lexical constraints on markup models. Although more progressive than DTDs, their unique particle attribution limitations in a document context (where like-named elements in block and phrase contexts are common) were deemed to be too limiting to make them a viable choice.

The Data Schema Definition Languages (DSDL) framework was instead turned to as a model for the future. Combining RelaxNG schemas for the structured markup together with Schematron assertions for enforcing finer markup logic provided exactly the balance of power and flexibility that was going to be needed for the modular framework that was planned.

Next generation of markup

Knowing how the group was going to implement the standard still left a long road ahead to build it. A model framework had to be constructed, rules for creating profiles defined, and working implementations developed that proved the framework was more than just an elaborate theory. The DAISY Revision Working
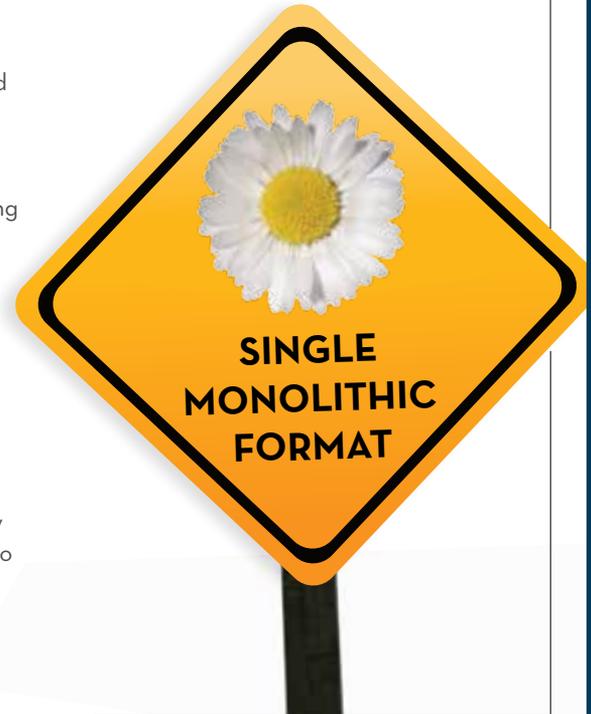
SINGLE MONOLITHIC FORMAT

The single monolithic format approach had so far failed to address the needs of the broad community DAISY serves, providing neither the structural and semantic richness nor the flexibility to accommodate the wide array of formats producers had to be able to generate. To begin developing yet another such standard would be to head down an evolutionary dead end.

book profile

periodicals profile

document profile

These profiles were targeted at the most prominent information resource types the community handles: a book profile for general works of fiction and non-fiction, a periodicals profile for news and magazine articles, and a general document type for word processing and similar documents found in office environments.

Group spent the next two and a half years filling in these blanks.

The Abstract Document Model, the theoretical model underpinning the specification, was developed to define the basic requirements all profiles had to adhere to. This model defines the common structure that all Z39.86-compliant profiles must implement (i.e., the root element and metadata and body content containers). Existing document definitions were analyzed in developing this model, and from this research a common layering of structural elements became apparent: sectioning, block, phrase, and text. These layers were then formalized into the model to ensure that markup is always structurally consistent across profiles.

To facilitate the modular, plug-in architecture of the framework, a set of core modules was also developed to accompany the specification (i.e., the set of pre-defined components that could be drawn on when building new profiles, reducing the work involved in creating profiles and ensuring greater consistency between them). These components allow the rapid development of new profiles, as they can be included in any new markup model and tailored to the needs of the resource being described without having to be completely rewritten.

RDF (Resource Description Framework) metadata was also given a prominent place in the new specification. All profiles must include a minimal set of RDF support for header metadata, and hooks into the document structure through metadata attributes are also provided. An RDF profile must be defined for each markup profile, which contributes to the consistency of prefix naming across documents and simplifies implementation for document creators. The Working Group also undertook to create an extensive structural vocabulary of properties to augment the markup with additional semantic meaning (one that can address both mainstream and accessible publishing needs).

To prove that this model could work for real-world production, a catalog of profiles was developed in parallel with the specification using the technologies and rules outlined in it. These profiles were targeted at the most prominent information resource types the community handles: a book profile for general works of fiction and non-fiction, a periodicals profile for news and magazine articles, and a general document type for word processing and similar documents found in office environments.

The profiles were built using the same common module pool, but the content models they define are uniquely crafted to the resources they define—proof that this new approach was working as designed. Sample documents were likewise created using these profiles to ensure that the content models were rich enough to support real production. After much review by the working group members and organizations and the release of three public working drafts, the profiles have now been made available for test use by the community as part of the current review of the specification to gain additional feedback.

Building a better DTB through EPUB

Having discussed text at length, the question so far left unanswered is what happened to the digital talking book portion of the specification. Originally envisioned as a Part B distribution format, work on this specification was suspended as it became apparent that the new EPUB 3 revision was open to incorporating even more of the essential functionality of DTBs, with the goal of turning it into a specification accommodating readers of all abilities.

Rather than create a competing specification, principals in the DAISY Consortium began working in earnest with the International Digital Publishing

Forum (IDPF) to pool their resources to forge a joint standard, one that could address the cross-cutting requirements of both constituencies: a single e-book format that meets the needs of all readers (no more delays producing accessible versions), and a single recognized accessible format that publishers can produce and distribute, thereby reducing their costs.

No longer a mixture of e-book and DTB technologies, this new revision of EPUB has seen the DTB accessibility components more fully integrated into the specification:

» The Navigation Control Center for XML Applications (NCX)— the menuing system for talking books—has been reformulated as an XHTML document to simplify its processing and rendering and to improve its international language capabilities.

» The subset of SMIL used for synchronization of audio and text content, now called a Media Overlay document, lives outside the content markup. The provision of audio and text synchronization has generated substantial interest from mainstream publishers, proving that it is not a feature of interest only to print-disabled users.

» Support for Text-to-Speech (TTS) markup has been integrated, allowing producers to enhance the content with pronunciation and prosody instructions.

The DAISY Consortium consequently anticipates adopting this new revision of EPUB as the distribution format for its members once the specification reaches recommendation status.

The End?

The Z39.86-201x *Authoring and Interchange Framework Specification* was recently released as a Draft Standard for Trial Use and a six month review is currently underway. The DAISY Revision Working Group anticipates being able to submit the specification for approval by NISO voting members and then ANSI after the trial closes on September 28.

While work will continue on the profile catalogs and core modules long after the specification itself becomes a standard, the anticipation is that this revision of the standard will provide a solid base on which the community can build their text production systems for many years to come.

The EPUB 3 family of specification documents are set to be released as formal recommendations during summer or early autumn of 2011, and their adoption by mainstream and accessible producers should be a swift process thereafter if early buzz is any indication. (See separate article on EPUB on page 4.)

But all things must evolve to stay relevant, especially standards… l NR l doi: 10.3789/isqv23n2.2011.08

MATT GARRISH <matt.garrish@bell.net> is an independent consultant who has done work for both the DAISY Consortium and the International Digital Publishing Forum (IDPF). **MARKUS GYLLING** <markus.gylling@gmail.com> is Chief Technology Officer at the DAISY Consortium.

DAISY Revision Working Group Members

Markus Gylling
DAISY, TPB (lead)

Josh Altherr
GH Braille

Ole Holst Andersen
Nota

Marisa DeMeglio
DAISY

Christian Egli
SBS

Matt Garrish
Invited expert

Boris Goldowsky
Center for Applied Special Technology

Leona Holloway
Vision Australia

Kenny Johar
Vision Australia

Dennis Leas
GH Braille

Sam Ogami
California State University

James Pritchett
Learning Ally (formerly RFB&D)

Kathryn Randall
Vision Australia

Per Sennels
Huseby

Christian Waldfogel
SBS

Z39.86 Authoring and Interchange Framework Specification
www.daisy.org/z3986/2011/Z3986-2011A.html

Z39.86 Profile Catalog
www.daisy.org/z3986/2011/auth/profiles/

Z39.86 Feature Catalog
www.daisy.org/z3986/2011/auth/features/

Z39.86 RDF Vocabulary Catalog
www.daisy.org/z3986/2011/vocab/

Z39.86 Core Modules
www.daisy.org/z3986/2011/auth/cm/

DAISY Consortium
www.daisy.org

International Digital Publishing Forum
www.idpf.org

EPUB 3.0
idpf.org/epub

RELEVANT LINKS