Chapter 2
Empowering the Distributed Editorial Workforce

Steve McNally

Abstract The scope of Content Management Systems has expanded to include the entire ecosystem surrounding the creation, management and consumption of content. The definition of the editorial workforce has expanded, too, to include participants within and outside of traditional publishing organizations. Requirements for efficient interfaces, intention- and contextually-aware semantic tools in support of content creation, editing and interaction across multiple devices has grown in importance as they enable working effectively with the expanded scope of content and people interacting with the content being managed. Participants in the ecosystem each have their own sub-networks of friends, professional and academic contacts to whom they should be able to easily turn for insights, feedback and general participation. It is this enabling and leveraging of the network effect[1]: that helps achieve the growth and scale that is mutually beneficial to creators, editors, publishers, advertisers and community members.

2.1 Introduction

The people who use Content Management Systems (CMSs) each use the tools that make up the platform not only to manage content but also to build an experience. Those who experience these creations—the people formerly known as the audience [4]—are participants in the experience: they read, comment, and share what has been created, edited and curated.

Traditionally, CMSs were built around specific processes rather than around people. While reasonable systems need to be built around sets of tasks, workflows and rules, peoples' jobs usually encompass a number of disparate tasks. With regard to

Steve McNally

Senior Group Director, Editorial Tools & Product Development, Forbes Media, New York, USA, e-mail: smcnally@forbes.com

[1] http://en.wikipedia.org/wiki/Network_effect

web content creators, editors, curators, these tasks span pages, sections and entire sites.

The modern CMS allows creators, editors and curators to interact with pages, sections and the whole site just as the general community does: Creators write posts, articles or comments. Editors can modify content of their own creation, sometimes that of others, and they can also package and promote content—slot it onto a home page, e.g., or into featured areas in a header or sidebar. Curators note that content is worthwhile by reading it, interacting with it, or sharing it. All community members view the same pages, but the creators, editors and curators have the ability to manipulate them.

Traditional roles are expanded. The editorial workforce is no longer centralized. Creators, editors and curators can be staffers, part-timers, or members of the community. Positive participation is rewarded with additional authority and capability.

Through it all is the on-going balancing act between high-function and intrusion. Much content and metadata regarding that content's relevance is presented within the context of what the creator or editor is viewing and doing. This relevance is contextual. It can always be requested explicitly. Opportunities are constantly sought to make contextually relevant content and information ambiently available so it shows up when it is natural for it to do so.

The modern CMS supports a community of participants. These participants interact with the content to varying degrees based on their permissions. All participants can be creators, editors, and curators. Following are details and use cases to demonstrate how this is accomplished.

The contribution of this chapter is intended to present works completed, works-in-progress and concepts yet to be realized with regard to the modern CMS. We will discuss the philosophy of tools designed to fit into organic workflows and to make relevant information available ambiently to creators and editors. Next, we discuss the roles involved in the creation of a decentralized editorial workforce empowered to attract, engage and grow an audience. Finally, we detail methods by which we engage, measure and reward credible, reputable contributions to the system and how the combination of all of the above streamlines and amplifies the process of content creation, management and interaction.

2.2 CMS Expectations: Philosophy and Tools

The modern CMS strives for frictionless publishing and promotion of content. The system is geared toward creating the shortest distance between creators' ideas and their publication. Removing layers between concept and publishing—between content and community—changes the mindset of working within the system: It transforms "tasks" into "work-style". It deepens and widens the pool of potential participants in the process.

Removing the need for creators and editors to go to a special place—a backend tool outside of the normal course of events for the majority of participants—creates a

new norm for creators and editors. It removes some of the work and provides an on-going perspective. The mindset of "working" changes as it becomes "looking at the site, seeing what's coming in, and seeing what's worth promoting and packaging."

The system enables a distributed editorial team comprised of staffers, part-timers, and engaged community members. They are each granted a role authorizing them to participate in specific, pre-defined tasks. They are given incentives to participate.

The tools create efficiencies in the creation, editing, and curation—selection, packaging and promotion—of content throughout the system.

2.2.1 In-Line Editing

In-line editing—also referred to as surf-and-edit—allows editorial team members to view site content as the general community does while also allowing them to update that content. Additional functionality beyond viewing is presented to authenticated members based on their role and authorization level. In-line editing provides context and perspective of how the content being edited will appear once those edits are published.

Use Case 1: Editor

A member with the role of *Home Page Editor* will see additional links when visiting the home page. Story titles will be highlighted to note they are editable. Excerpts and images will have *Edit*, *Add New* and *Promote* link actions next to them. A *Re-arrange* link allows editors to drag and drop home page stories into a new order or to insert other content not currently present into the home page (see Fig. 2.1).

Working as a content creator or editor within the same context that community members are interacting with the content provides an on-going contextuality. From a tools perspective, subsuming the create and edit tools within the context of the final product removes a layer between editor and community: editors need not log into a separate area, traverse a file system or series of directory folders to find the file to edit and later publish. Rather, they work on the content in-line and in the same context the public will experience that content.

Practical advantages of role-based in-line editing also include mitigated training requirements. This enables a significant expansion of potential participants: as everyone is capable of visiting site pages that interest them, granularizing editorial capabilities down to the level of topics, sub-topics, and even keywords means publishers can recruit workforce members interested solely in "scratching their own

Fig. 2.1: In-line editing

itch"[2]: Community members care less about what the Publisher, e.g., wants; they are interested in what they want themselves. Some additional participation can be expected from the community when Publishers provide community members with what they are interested in seeing. Once interests and roles are defined, community members can be targeted down to the individual, and individuals can be grouped along lines of historical interests and activities. Members and groups can be catered to organically. Authorized community members can view, edit and curate topics, keywords or subjects in which they have vested interest. Individual itches can be scratched while still serving the good of the overall community.

Members of the community, through their own activities related to that content can become packagers and promoters implicitly via their engagement with the content and explicitly via rewards systems (more on that to come in Section *Roles*, *Rights* and *Reputation*). This empowers and engages community members. It provides a highly curated experience featuring quality content for the audience. It provides the publisher with a highly engaged and specialized workforce packaging and promoting content based on these individuals' subject matter interest and expertise.

2.2.2 Semantic Tools and Contextual Presentation

Providing basic categorical or taxonomic information for content allows that content to flow to appropriate areas of the system. Taxonomic metadata about content can be applied manually and explicitly by creators and editors or automatically and implicitly via algorithm.

Explicit taxonomic information includes when creators and editors manually add keywords and categories. Information can be added explicitly via first-hand knowledge of creators and editors or may be suggested by tools like Zemanta (http://

---

[2] http://en.wikipedia.org/wiki/The_Cathedral_and_the_Bazaar#Content

www.zemanta.com) or OpenCalais (http://www.opencalais.com) which process content as it is being created as well as after publication.

Implicit taxonomic information generally includes post-publish processing of content including the extraction and handling of known categories and entities.

Use Case 2: Creator

Implicit information is added when occurrences of *Robert DeNiro*, *Lady Gaga* or *Steve Jobs* appear in content being created. These names are automatically marked-up and linked to profile pages from The Internet Movie Database (http://www.imdb.com), Gracenote (http://www.gracenote.com) and Forbes' Billionaires' Lists based on pre-defined rules (Fig. 2.2).
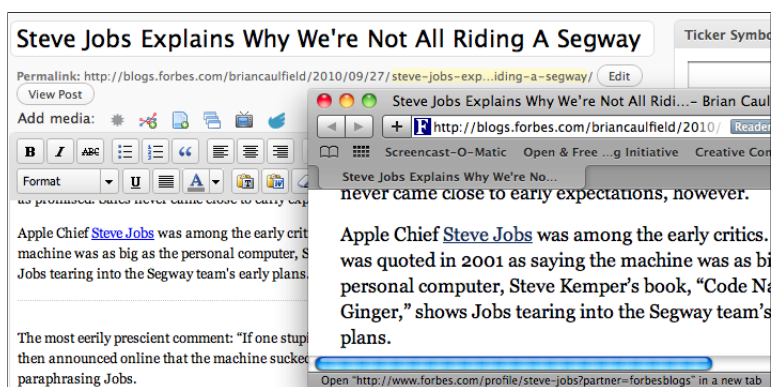

Fig. 2.2: Names are automatically marked-up and linked by semantic entity extraction

Extending the taxonomy provided by, e.g., the Open Directory Project (http://www.dmoz.org), to include publisher-specific information regarding entities and their appropriate linked pages allows publishers to extract those entities and link them to resources that create opportunities for additional revenues, Search Engine Optimization (SEO) benefits, and direct traffic while giving community members the opportunity to delve more deeply into subject maters of interest. The object being availing of cross-promotional opportunities, within and outside of the current content and site, while mitigating the need for explicit taxonomic work on the part of creators and editors: defined entity extraction rules, e.g., will link celebrity names, automatically, to a content partner specializing in Hollywood; the names of popular cultural figures will link to a sister-site specializing in music, culture and the arts, and the names of business figures will link to people and company profile

information elsewhere on your own content network. Creators and editors will have the opportunity to change these automatically generated links; but, by default, unlinked entities will be discovered and linked to predefined sources without the need for further action on their parts.

2.2.3 Curation Tools

Curation is the act of noting certain content as worthwhile. All roles and community members can curate content, and curation can be explicit or implicit. First, explicit curation includes up-voting posts or comments—such as using the *Like*-button from Facebook (http://www.facebook.com)—leaving a comment, or sharing content via Facebook, Reddit (http://www.reddit.com) or elsewhere. Giving money to posts or comments with the help of micropayment services like Flattr (http://flattr.com) would be another example for explicit curation.

Second, implicit curation includes clicking through to an article or comment, spending active time on that content (e.g. time on page plus scrolling past pre-configured checkpoints within the content). It can include navigational path information: posts arrived at, e.g., via related or promoted links from another post are implicitly noting the value of both that link and target content.

Explicit and implicit curation actions are recorded and can be weighted and used to present contextually appropriate content and functionality to creators, editors and curators throughout the community. Curation can also include selecting content for additional promotion—giving a story a position on the home page, e.g.

Curation tools include (1) recommending a post or comment, (2) selecting content for inclusion on the home or a section index page and (3) selecting content for other packaging including Editor's Picks, special feeds or newsletters. Curation tools are contextual. They are presented to members according to the scope of their role and per any pre-existing content categorization.

Use Case 3: Editor

A member with the role of Technology Editor is permitted to categorize new content as *Technology*. Within the Technology section, that Editor can promote a story to appear on the home page or above the fold on the section index page by noting it as an Editor's Pick (Fig. 2.3).

While writing a new article, opportunities to see relevant keywords, archival content, and trending topics are presented. This can be in a sidebar, an overlay or via contextual-click.

Fig. 2.3: Assigning a post to the Technology section

Use Case 4: Creator

Within the course of writing a paragraph, suggested links to related images and news sources appear in the sidebar of the editing tool. Via an explicit request, creators can reach out to Aardvark (http://vark.com), LinkedIn Answers (http://www.linkedin.com/answers), or a specific Twitter (http://twitter.com) list or tag for verification or amplification of their topic-in-progress (see Figs. 2.4 and 2.5).
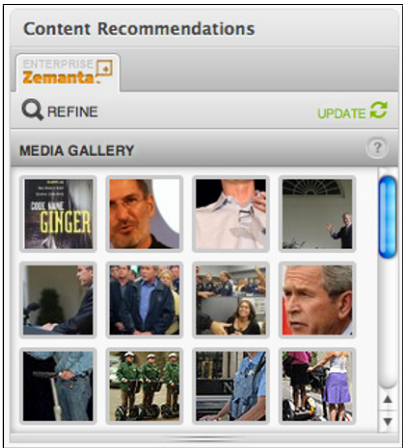

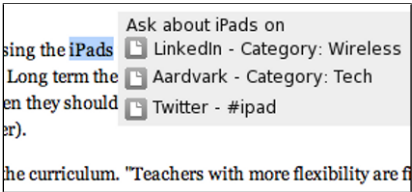
Fig. 2.4: Recommendation of related images

Fig. 2.5: Request of further information for a term

Reaching out to the community with "Here's what I know. Here's what I don't know. What do you know?" [3] helps accomplish more work and involves the community more deeply in the process. This deepens sense of community—it is a virtuous cycle of increasing engagement.

Once the new content is published, a list of community members who are active commentators or curators of similar content can be presented to creators. The opportunity to promote relevant content to the relevant people—presuming community members have given permission [2]—presents creators with the ability to find a community and helps the community find relevant content.

Based on what the creator or editor is working on, she may want more contextual information or may wish that contextual information to be even less intrusive. Creators have options for contextual presentation to best fit their work-styles.

The data are available as are the tools to get at and present it. The challenge is to present such data and information to the creator as organically as possible, to present related information naturally so that it is easily discoverable and utilized, but not so it overtly interrupts the creator's workflow.

2.3 Roles, Rights and Reputation

When building a CMS to serve a growing network of writers, editors and participants, *Who are you? What are you allowed to do here? What have you done previously?* are core to execution from business and technology points of view. Roles and Rights cover who you are and what you are permitted to do within the system. Reputation refers to the credibility and trust built up through historical activities and is the measure by which participants can be promoted through the system by gaining authorization to perform additional editorial tasks.

2.3.1 Definition of Role

A role is a defined list of tasks a community member is authorized to perform. Roles create a replicable model for adding new talent to the team. The team includes salaried staffers and others who are compensated via stipends, performance bonuses and other monetary and non-monetary rewards. Defining roles for participants provides the framework for creating a highly scalable distributed workforce. This scalability can include a one-person operation looking to recruit and manage wider community participation to large organizations looking to augment their editorial staff.

Everyone who visits and participates in the content created via the CMS is a member of the community. All roles are included under *Community Members*; they are differentiated by their authorization levels—in the degree to which they can interact with the content.

Reputation also provides a path for advancement through the system and incentives to work toward this advancement. These incentives are designed to drive deeper engagement and to advance the goals of the publisher as well as the community as a whole [5].

2.3.2 Basic Roles and Rights

1. Creator: Somebody who can publish and manage their own content
2. Editor: Somebody who can publish and manage content as well as manage others' content
3. Curator: Someone who can, explicitly or implicitly, note content as worthwhile
4. Community Member: anyone interacting with content on the platform

Being able to slot people into these defined roles allows for the scaling of a distributed workforce into the thousands. Measuring the performance of individuals in these roles means we can modify responsibilities based on individuals' actions. Reputation is a key component of managing these individuals in their roles; Reputation will be explored in detail in the following paragraph. Focus on community member-recruitment is based on collective needs and interests.

2.3.3 Reputation

Reputation is the history of a creator, piece of content or community member. Reputation's purpose is to make default decisions about quality based on historical activities (cf. [1]). Reputation will help inform re the best creators, editors, curators, community members and content. In this sense, *best* is not particularly subjective: Reputation is an amalgam of desirable behaviours and formulated weightings and scores based on these. Reputation can drive compensation. Reputation can be employed as a first-line filter: when there is a lot of content for editors to consume and process, sorting by reputation can provide an efficient first step.

The primary usage of Reputation is as an initial filtration system: content and community input can be sorted by and presented to editors by reputation order. Using these weightings, we will be able to present various tangible benefits or capabilities to those we have weighted. Another purpose is to drive participation in desirable behaviours throughout the system via rewards.

Desirable behaviours include actions like writing posts, leaving comments and generating page views. Participation in desirable behaviours will drive tangible benefits, e.g.:

- Creators with better reputations may have their content promoted more frequently
- Community members with higher reputation scores may have their comments displayed ahead of comments with members with lower reputation scores
- Content with better reputation may be promoted in more places; Content with better reputation may be displayed ahead of content with lesser reputation

Participation in these behaviours informs numeric weightings for each creator, each community member and each piece of content in the system. Reputation quotient is a score kept of an amalgam of desired behaviours, actions and occurrences.

Examples for reputation quotients for creators, community members and content are given in the following:

- Reputation quotient for a creator, e.g., includes how frequently that creator logs in, posts and comments on others' articles.
- Reputation quotient for content includes page views and comments received, referrals and quality of referral sources. These same positively affect the creator of that content.
- Reputation quotient for community members includes logins, comments left, comments up-voted and replied-to by others and curation actions completed.

A community member who has participated in enough desirable behaviours that they have achieved a high reputation score would have their subsequent comments, e.g., default to a higher score. A highly reputed community member could be made known to the creator(s) she follows.

Another tangible benefit for a creator or editor could be enabling them to curate the home page based on having reached a threshold through desirable behaviours.

In all, a reputation system helps keep management informed as the value grows across the network. It is also a tireless watcher of cross-network activity: 24 x 7 x 365.25, it is seeing and remembering all desirable behaviours performed. This combination of automation and human intelligence allows us to build a network of creators, community members and content with definable credentials. It gives us actionable insight into the relative value of our participants. The system will be extensible as to what can inform it—new metrics and desired behaviours—and what it informs—new roles and capabilities for those meeting thresholds.

It is designed do all this with minimal manual intervention: as we envision new roles – deputies, acolytes, network programmers for a day – we can create workflows that enable them automatically with reputation as the core of a self-healing, spontaneously-creative system than can help us scale our network and business.

Use Case 5: Editor

The Technology Editor logs in to update section home page. There have been 30 new posts and 90 new comments since she logged in last. Presenting her with posts from creators and community members with the highest reputation quotients first is one way to prioritize which content she should promote and use to refresh the Technology section home page (Fig. 2.6).

Arming intelligent and engaged community members—including staffers and the more distributed workforce—allows content management decisions to be made based on science *and* art. That community members, staffers and the wider team also have reputation scores associated with them means their actions and behaviours can be measured and managed: Publishers can ensure needs are being met to the mutual

New Post

3 hours ago
21 views

23 New Items | Sort By:
☐ Creator Rating
☐ Topic Rating
☐ Publication Date (Current)

Echoes of Abacus: JP Morgan to Face SEC Probe Over CDO Transaction with Hedge Fund

The massive explosion that was the bursting of the housing bubble has left a toxic cloud that still hasn't blown away. On Monday, a report of an SEC investigation targeting J.P. Morgan over the creation of a complex mortgage instrument echoed the Abacus investigation where Goldman Sachs paid a record [...]

AGUSTINO FONTEVECCHIA
Moral Hazard

New Post

3 hours ago
538 views
1 comment

This Week in Credit Card News

Provided by LowCards.com For Many Over 55, Debt Defers Dreams The recession has forced many Americans age 55 and over into unemployment. It has stripped away years of their savings or dramatically reduced incomes during their final high-earning years. The ranks of older bankruptcy filers have been swelling rapidly. More than two-thirds [...]
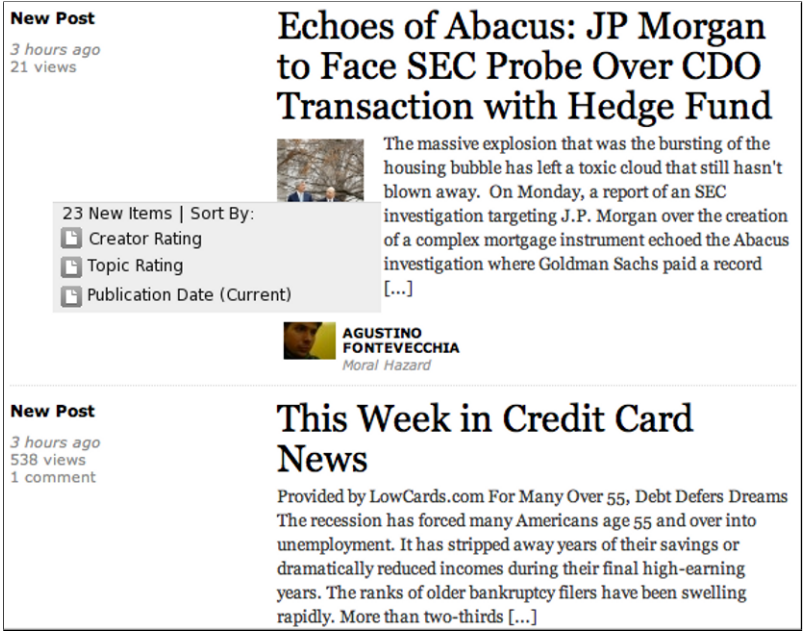
Fig. 2.6: Prioritisation of posts

benefit of all; Editors can spend their limited resources first on the most highly-reputed content; Members of the community can decide to read and engage with most highly reputed content and creators. All members are provided incentive to engage more deeply with content throughout the system.

A lingua franca for reputation—a standard protocol that captures desired behaviours—would enable the portability of reputation across CMSs, sites and communities. There has been work to create a standard activity stream protocol (`http://activitystrea.ms`); that protocol could be extended to provide scores and weightings against those activities. Diaspora (`http://www.joindiaspora.com`) and the DiSo (`http://diso-project.org`) projects both have goals of distributed, portable activities and reputation. Additional sites could use those concepts, add their own desired behaviours, and authorize community members within their own CMSs accordingly.

2.4 Conclusion

Much of this CMS work is available for implementation today. The pieces around creation, roles, rights, reputation and the semantic tools certainly exist and can be

improved upon and integrated to suit individual organizations' needs. Execution comes to definition, adoption, customization and usage.

For more traditional organizations, especially, there is a learning curve: less-formal editorial workflows — as there is where creators have less direct editorial oversight and there is little to no direct approval process -– necessitate a ceding of direct control of the messaging. Publishers, Public Relations, Corporate Communications, and Marketing departments are used to strong, direct control. Trust in, and monitoring of, how the wider community manages content quality has to be learned through experience. But the benefits of distributed editorial systems should be clear: the ability to massively distribute workload and the creation of ecosystems that are mutually beneficial to publishers, creators, editors, and community members.

References

1. Farmer, R., Glass, B.: Building Web Reputation Systems, 1st edn. O'Reilly, Sebastopol, CA, USA (2010)
2. Godin, S.: Music lessons (2008). URL http://sethgodin.typepad.com/seths_blog/2008/01/music-lessons.html
3. Jarvis, J.: Denton goes to the bench (2007). URL http://www.buzzmachine.com/2007/12/14/denton-goes-to-the-bench/
4. Rosen, J.: The people formerly known as the audience (2006). URL http://archive.pressthink.org/2006/06/27/ppl_frmr.html
5. Yahoo Developer Network: Reputation (2010). URL http://developer.yahoo.com/ypatterns/social/people/reputation/index.html

Biographical Notes

Steve McNally creates platforms and processes to empower the distributed editorial workforce at Forbes Media. Previously, Steve oversaw technology architecture, software development and site operations as CTO for True/Slant (which Forbes acquired in June 2010). He came to True/Slant from Condé Nast, where he spent eight years, most recently as Senior Director, Product Development. In that role he managed the build-out of Parade.com's core publishing platform. Prior to that, Steve managed channel development for About.com, where he designed, built and launched an affiliate ad network. Steve also worked for Prodigy Internet and IBM where he built knowledge management, business intelligence, communication and collaboration tools for internal teams and partners.