

White Paper

Social Publishing with Drupal

**Building
Social Businesses
on the Web
with a Framework
for Content
and Community**

February 2010

Geoffrey Bock, Senior Analyst

Sponsored by



THE GILBANE GROUP

Gilbane Group Inc.

763 Massachusetts Avenue
Cambridge, MA 02139 USA

Tel: 617.497.9443

Fax: 617.497.5256

info@gilbane.com

<http://gilbane.com>

Executive Summary

Social publishing combines groomed and authoritative content, produced by an organization and emphasizing its core messages, with user-generated content that customers contribute via blogs, wikis, and social media tools. Drupal is an example of a social publishing platform, developed and maintained as an open source project, and delivered at an affordable cost.

Drupal is now deployed in major media companies, high technology firms, universities, magazine publishers, government agencies (including the White House), research groups, and non-profit organizations. Whether it is in a commercial, non-profit, or government setting, organizations rely on Drupal to project their presence over the web and to channel the interactive experiences that foster communities of contributors.

By leveraging Drupal's capabilities as a social publishing platform, organizations are able to reinforce their branded experiences and deliver relevant content to their customers and stakeholders. By exploiting Drupal as an open source project, developers supporting these organizations can easily enhance and extend Drupal's capabilities, and introduce innovative modes of interactivity that meet specific business requirements.

Drupal is an attractive investment with substantial business benefits. Organizations can keep their license and support costs modest by building on an open source project. Organizations can leverage the collective expertise of Drupal developers to solve immediate publishing problems. By relying on Drupal, organizations can stay abreast of the rapid technology changes when building competitive solutions for the digital age.

Social Publishing in 2012

It's 2012 and social publishing is more important than ever before. Let's imagine that your company designs and develops state-of-the-art home entertainment systems for price-conscious consumers. As a nimble leader in the consumer electronics market, your firm builds its brand and market share by delivering an intuitively controlled immersive environment, by rapidly incorporating innovative technologies, and by engaging with its customers.

In our hypothetical situation, you are the director for customer experience, responsible for ensuring that everything works together as expected, and that your customers have easy access to the relevant information. The web is the conduit to your current customers and prospects. You publish exclusively online.

When Publishing Goes Social

What does this entail? To begin with, you produce the multimedia clips and interactive product guides that are easily accessible from a browser embedded in your systems and over the web. You maintain the support knowledge base, and publish frequent updates about new components and enhancements. When customers have questions, they turn

first to your company's online resources. You carefully monitor customer results to ensure that they are finding the relevant information.

You engage communities of contributors to build awareness about your solutions. These community members are your loyal fans, boosters, and critics who voice their views about your products and services. You enable them to communicate with one another and share their perspectives. As the director of customer experience, you are continually listening to what they have to say and facilitating their efforts to discover relevant information.

Moreover, you learn from your customers and frequently incorporate their insights into your product plans. You provide the online venue for communication and collaboration where your customers reinforce your branded experience. With these efforts, you build the buzz, create the inter-personal connections, and support viral marketing campaigns for your systems.

Web Content Management plus Social Software

How do you do this? With social publishing, you combine web content management with social software to reinforce your messages and support your customers.

Groomed Content

To begin with, it is essential to publish the groomed (and authoritative) content about your company's products and services, and ensure that all of this information is current, accurate, and easily retrieved by your customers. No longer are you publishing only electronic documents and data sheets. Rich media clips and interactive guides are a core part of your content repertoire.

You are investing in a content platform, which includes a content repository and a predefined taxonomy. You rely on workflow to ensure that subject matter experts contribute and review information before you publish it online. You index your content using standardized tag sets to ensure that search engines can easily find relevant items.

With this platform in place you can manage all content types in a systematic manner. You can stay abreast of product changes by easily publishing new and updated information about your home entertainment systems and their capabilities.

User-Generated Content

At the same time, it is essential to integrate these authoritative sources with the user-generated content your customers contribute via blogs, wikis, comments, user submissions, and online forums. You make it easy for them to participate in the continuing conversations about your products and services by incorporating such capabilities as tagging, ratings, and social networking features into your social publishing environment.

Through these interactive resources, your customers are able to:

- Quickly and easily locate insights contributed by other community members
- Help one another solve problems
- Identify subject matter experts
- Discover new perspectives
- Build awareness about your products and services

As a result, you are able to keep in touch with your communities of contributors, and support their capabilities to share information over the web.

Managing Costs and Expectations

Finally you must manage your costs and expectations. As your firm grows rapidly and builds its branded experience, you need to keep a careful eye on your bottom line and optimize your investments in social publishing.

It's essential to manage both groomed and user-generated content over the web through a consistent interactive experience that meets your business objectives, at a price you can afford. It's important to build on the capabilities of an integrated platform. It is also important to invest wisely, and leverage the network effects from a community of contributors. For cost-effective development, you need a flexible and adaptable framework that supports your continually changing business environment.

A Framework for Social Publishing

What's needed is a modular approach for developing interactive environments that blends groomed content from authoritative sources with the user-generated content, produced through community experiences. We'll term this *a framework for social publishing*. This framework has several unique capabilities.

- It blends groomed content and user-generated content in ways that meet business objectives.
- It fosters connections among the content creators and the audiences who consume the information.

Moreover, this framework manages both content and context from two perspectives – content creation and content distribution. With a social publishing framework in place, groomed and user-generated content can be readily found and smartly distributed across the web by paying attention to the metadata and tags used to describe the various items.

Overcoming the Content Conundrum

Yet here's the challenge for managing content in the digital age. When relying on the web to build their brand and support their customers, companies need to overcome the content conundrum.

- If firms don't have great content to distribute over the web, customers and other end users won't be able to learn about their business solutions. It's essential to

engage the target audience with relevant interactive experiences that reinforce their brand and business.

- But even if companies do produce great content and involve their community of contributors, they still need to manage content production and delivery. Content producers need easy ways to publish information over the web. Findability is a key issue. Distribution is essential. Companies must maintain the environment, verify the accuracy of the published information, and facilitate online conversations through a seamless set of business activities.

As a publisher responsible for the content distributed over the web, every organization expects to involve its communities of contributors in its own way.

Capabilities for a Loosely Coupled Platform

How do organizations develop an interactive environment that captures the social publishing experience? What's needed is a flexible application environment that connects people with content and people with people.

A framework for social publishing encompasses four key capabilities: transparency, trust, collaboration, and meritocracy.

- **Transparency.** Organizations should focus on fulfilling a basic expectation, that all relevant information is readily available. Customers and other end users should be able to easily find and access what they need without any predefined barriers.
- **Trust.** The available content should be authoritative and verified by reliable and known sources.
- **Collaboration.** It's important to have many minds working together to solve problems. All the members of a community should have a voice in the conversations; the best ideas win.
- **Meritocracy.** While the best ideas win, there still needs to be a decision making process. In a community of contributors, recognized leaders should be able to influence and make decisions.

This framework leverages the essential standards of modern web site development. From a technical perspective, it's all about developing the resources to manage the connections among content and people, and increasing the value of information flowing through the community.

In particular, a framework for social publishing depends on a loosely coupled platform, where RESTful services are readily exposed and consumed. This framework blends great content with delightful social experiences. It incorporates a smart (yet seamless) interactive environment, supports key services for the semantic web, and includes easy, self-evident administration capabilities. It is modular and adaptive – developers can easily add new features to the content platform.

In short, this framework can adapt to the speed of business. Let's examine what the framework includes and how it works.

Great Content

A framework for social publishing manages all types of web content in a systematic manner, stored within a shared repository. Content creators within an organization, responsible for creating and updating groomed content, can easily add and modify content items on their own, without involving web masters or other information technology (IT) staff members.

Nor is the content limited to text, photos, and static images. Content creators can easily add (and modify) all kinds of rich media experiences including video clips (with sounds), interactive drawings, games, maps, and any other kind of online experience to the shared repository.

Granular Security and Remote Access

A framework for social publishing provides the system-level security, at a very granular (or object) level.

- The framework supports access rights and permissions to the underlying repository, to collections of content items stored within the repository (typically defined as folders within a hierarchy of nested folders), and to the content items themselves.
- The framework supports the definition of multiple roles. Developers have the capabilities to define their own sets of roles.

The groomed content is effectively organized, curated, and delivered in ways that both automated processes and end users can understand. Thus the content can be readily syndicated to external sources. With appropriate access rights granted through the system-level security, search engine crawlers and other automated processes can readily query the repository and extract predefined sets of items, including tag sets with links to content items or the content items themselves.

Rendering and Delivery

A framework for social publishing manages all aspects of the presentation and delivery of the varied content types. There's a single source for content, rendered to multiple resources and environments, using HTML, Flash/Flex, or other content-specific rendering formats.

Thus web pages look great and load quickly on a desktop or laptop browser – supporting the full panoply of web browsers available to developers. The same content can be rendered on digital home entertainment systems, smart phones, personal digital assistants (PDAs), tablets, and many other kinds of digitally connected devices.

The Publishing Process

The framework manages all aspects of the publishing processes to produce the groomed content. There is:

- Embedded workflow to ensure that the content is routed to predefined people and roles for review, updates, edits, and approval
- Granular security to ensure that only the people and roles with appropriate access rights can create, edit, approve, and publish content over the web

All aspects of the publishing process can be logged, tracked, monitored, and audited. For example, a managing editor for a web site can determine which editors are reviewing what content items and, with appropriate privileges, reassign the workflow requests.

Ensuring Transparency and Trust

As a result, a framework for social publishing ensures the authenticity and reputation for the groomed content published on a web site. It is essential to be able to verify trust in the sources. The content produced and rendered should be what the content creators and editors produced and published.

It is also essential to be able to demonstrate transparency in the sources. Once made publicly available (and published) to the target audience, the content should be tagged and rendered for easy browsing and discovery.

A Delightful Social Experience

A framework for social publishing builds and sustains a community of contributors – groups of people interested in a topic and able to share perspectives under the auspices of a web site owner, a person or firm termed a web publisher. This framework supports the full range of features for creating user-generated content – including threaded discussions, blogs, wikis, syndicated content (using RSS feeds), and shared editing resources.

Contributing Content to Conversations

Communities of contributors (people external to the formal publishing processes) can start a conversation and add content on any relevant topic, in any format that the site owner permits. Contributors can post their contributions and get the word out about their perspectives and experiences.

With social tagging, they can tag their content by the terms that others have used, or create their own tags. The framework provides the capabilities for sharing tag sets and also for managing the tagging process.

Nurturing the Community while Pruning Contributions

Of course, the web publisher fosters the community of contributors by hosting the site and maintaining the interactive services. Contributors can build community in several different ways. They can:

- Easily discover what other community members are contributing by following topics and threads

- Highlight and promote contributions with ratings, voting, and annotations
- Track contributions by predefined people and roles, such as following experts in terms of search activities, posts, profile entries, and colleagues
- Graph relationships to define and maintain social networks

Web publishers can integrate the groomed content with user-generated content by:

- Providing access to the predefined taxonomy of terms used to catalog the groomed content
- Using the social tagging capabilities to support ad hoc categorization
- On occasion, using the socially generated tags to add terms to the taxonomy

Furthermore, publishers can structure the flow of content and conversations by periodically reviewing the user-generated content and “pruning” the contributions – either by editing the contributed items or reorganizing how they are tagged and accessed, or both.

Supporting Inbound and Outbound Venues

As a result, a web publisher can transform a conventional publishing environment, featuring simply groomed content, into a social publishing environment that blends both authoritative sources with community-wide contributions. With a framework for social publishing in place, it is easy to develop and sustain an inbound community of contributors who engage with one another through a predefined site.

Moreover, community engagement can extend across the Internet. The framework can also support outbound communities of contributors, and push participation out to Facebook, Twitter, and other external communities available on the web. Both the groomed content and the user-generated content (cataloged by their appropriate tag sets) can be distributed to various external communities.

A Smart Environment for the Semantic Web

Beyond a familiar interactive environment, a framework for social publishing creates a smart environment for content delivery. The framework provides the essential resources for embedding semantics into content, and for delivering semantic web services (using technologies such as RDFa and OWL) across the Internet.

By managing and exposing the tag sets that describe content and links, the framework makes it easy for both end users and automated processes to categorize the underlying semantic relationships among terms, and to discover with increased precision and recall the content stored within a web site. The semantic relationships can be based on ad hoc mappings, taxonomies, and ontologies. These relationships are exposed as linked data.

Ad Hoc Mappings, Taxonomies, and Ontologies

In addition to managing the content items themselves, the framework also maintains associated tag sets that describe these items. Initially, the tag sets are based on categorization criteria maintained within a particular web site, including:

- The taxonomies tailored for the web site
- The community-wide tags produced through social tagging

Search engine crawlers and other automated processes use these tag sets to improve the quality of search results and the ability to discover content contained within the web site.

Furthermore, the framework supports semantic web services that describe the relationships among terms, often based on taxonomies or ontologies. These services:

- Allow automated processes running on distributed servers to locate, extract, share, and re-use content in a secure yet manageable fashion
- Give context to tags that describe content items and characterize the relationships among them

By managing the tags and links based on an understanding of the underlying semantic relationships among terms, the framework makes it easy for end users and for automated processes to discover content stored within the web site.

Exploiting Semantic Web Services over the Internet

The framework has the flexibility to easily exploit semantic web services available from remote web servers, and to incorporate the tag sets produced by these services. Thus the framework can categorize its content items and links according to standard ontologies maintained by known authorities. While publishers can manage both groomed and user-generated content within their own web sites, they can categorize and tag their content using semantic services running elsewhere on the Internet.

The list of known authorities exposing tags as semantic web services is growing steadily. Currently the authorities include:

- DBpedia, maintained by the Wikipedia Foundation for the categorization criteria of Wikipedia
- Open Calais and Times Tags API, maintained by Thompsons Reuters and the New York Times respectively, for categorizing news-related content

There are additional initiatives underway designed to describe generalized semantic web services, independent of established authorities. These initiatives include:

- Meaning of a tag (MOAT)
- Friend of a friend (FOAF)
- Semantically-interlinked online communities (SIOC)

When implemented, these initiatives promise to map the tags in a tag set according to different kinds of ontological criteria and make it easy to manage content for the semantic web. The end result will be a smart environment for social publishing where organizations can target how they publish content and build communities of contributors over the web.

Easy Administration

A framework for social publishing features an intuitive and consistent set of user experiences for easy administration. While an IT group designs and implements the capabilities of a web site, end users with little technical knowledge or skills are then responsible for the day-to-day administration of users, groups, content, categories, access rights and permissions, workflow definitions, or other activities required to support standard operations. Administrators are also the conversation gardeners, and are able to prune user-generated contributions as required.

Moreover, as part of the framework, the administrative capabilities are collected into one or more dashboards. The underlying functions are intuitively presented and structured for ease of use. For administrators, maintaining a social publishing environment is simply another set of user interfaces in a browser-based experience.

Harnessing Developers' Know-how

Social publishing lends itself to community-oriented development. With a loosely coupled platform based on a RESTful interfaces, developers can transform a site that supports conventional web publishing into an interactive environment, designed to build awareness around a theme and promote smart conversations among contributors.

How can companies, government agencies, and other types of organizations leverage the power and flexibility of a loosely coupled platform? It all depends on how best to harness developers' know-how, and nurture innovations for social publishing.

Certainly commercial software organizations can exploit the capabilities of a loosely coupled platform, and successfully add community-oriented capabilities to conventional web publishing environments. But building a platform for social publishing lends itself to an open source approach to software development.

With open source, developers leverage the power of the network and collaborate with one another to evolve innovative social publishing technologies. With open source, developers can 'scratch their own itch,' construct unique approaches to solve particular problems, test their innovations with colleagues and other contributors, learn quickly from one another, and contribute code back to the open source community. It is all together likely that a framework for social publishing will evolve much faster and more successfully through open source development.

How Drupal Makes a Difference

Drupal is one implementation of a framework for social publishing. It combines capabilities for web content management with social software tools into a seamless interactive environment. First released by Dries Buyaert in 2001 as GPL and now in its sixth major release, Drupal is developed as an open source project and nurtured through a distributed community of contributors (see www.drupal.org).

Drupal consists of a managed core, maintained by Buyaert and a core development team, plus additional community contributed modules. Currently, Drupal encompasses over 4,000 modules, developed and supported by a worldwide community of developers. Enterprise grade service and support for Drupal is available from Acquia (www.acquia.com), an open source software company.

Drupal combines the capabilities of a content management system with blog, wiki, forums, and custom forms, as shown in Illustration 1. It provides a consistent set of APIs to support user management and access controls, as well as unified services for managing tagging, taxonomy, ratings, comments, and search services.

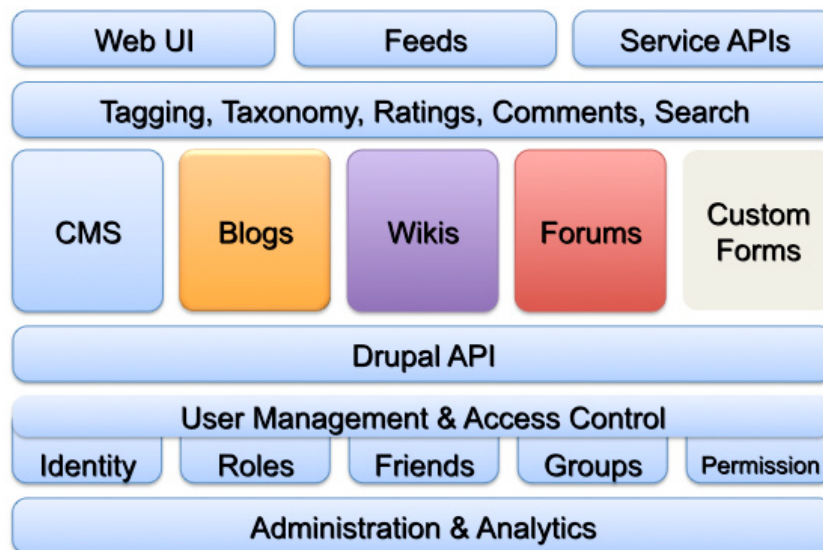


Illustration 1. Drupal provides an integrated platform for social publishing.

Drupal in Action

Drupal is being adopted by organizations large and small. Major media companies, high tech firms, universities, magazine publishers, government organizations, and research teams rely on Drupal to project their presence on the web, deliver content to their customers, and manage the interactive applications that foster community connections. Organizations running Drupal value its flexibility and its ability to seamlessly integrate with an underlying enterprise infrastructure.

For example, a media firm needed to add support for production quality streaming video to its site. Its existing web content management system, developed in an earlier

and less interactive era of the web and featuring text-oriented articles, could not easily support video clips at the scale required for a satisfactory customer experience. The media firm used Drupal to develop an interactive experience for video that was indistinguishable from its existing web pages.

As another example, a software company sought to build its brand by underwriting content development on several topics, and by sponsoring an interactive environment for communities of contributors to discuss various issues and develop action plans. Drupal provided the platform and support for social publishing -- managing the web content, capturing and pruning community contributions, and making the content easily findable and usable by search engines and other third party services. The end result is a branded experience that fosters the firm's overall business goals.

As a third example, now a well-known Drupal site, the White House values the flexibility of the underlying content infrastructure for publishing information at Whitehouse.gov.¹ The President's staff relies on Drupal to support its digital media initiatives. Text, images, and video clips are seamlessly delivered in a scalable and responsive environment. They are well organized and easily browsable and searchable. For instance, as events unfolded in the first twenty-four hours following the earthquake in Haiti on January 12, 2010, the White House published not only the news releases and video clips of President Obama's public statements, but also information about how people could donate money to the disaster relief efforts.²

Let's examine how Drupal delivers these kinds of business benefits.

Managing Content on the Web

Drupal is designed from the ground up to manage content for interactive distribution and delivery over the Internet. It encompasses a loosely coupled platform, based on a RESTful architecture. It manages any type of content running on the web – including web pages, multi-media clips, engineering drawing, maps, gaming experiences, and other digital formats now being invented.

Nodes and Themes

Drupal defines content as sets of nodes, or self-contained chunks of information. In addition, Drupal separates content from presentation.

- Each node includes a title, a body, and links to that body's type, authors, and publication metadata. This information is defined as XML tags and is directly

¹ For an analysis of Drupal at the White House, see <http://techpresident.com/blog-entry/whitehousegov-goes-drupal>.

² See <http://thelede.blogs.nytimes.com/2010/01/12/gleaning-information-from-haiti-online/?hp> for a chronology of breaking news and responses by world leaders on January 12th, as the scale of the tragedy became known.

addressable as a REST URI. Nodes are organized into collections of nodes, defined as multiple items stored within folders on a web site.

- To render content, Drupal includes a series of themes, comprising one or more CSS pages and one or more PHP-based templates. A theme defines the presentation of a node or a collection of nodes displayed within a device (such as a laptop web browser).

As a result, developers can both define and present content in a very flexible and adaptable way, relying on predefined tags that describe the nodes. Adding a new content type is simply a matter of defining the additional metadata (and tags) that describe a specific node type. Developers then manage both content and tags, and expose them as RESTful interfaces to remote web services.

Shared Repository with Role-based Security

Drupal removes the web master from content management. Drupal stores content within a shared repository. Content creators, editors, and other end users check content items into and out of the repository. Drupal detects the interactions and can be configured to ensure that one individual, modifying content stored within the repository, cannot overwrite another person's contributions. End users can publish content on a web site, without requiring any assistance from web masters or other IT staff members.

Key to its success, Drupal features extensible, flexible role-based security.

- All content items stored within the repository are secured with access rights. Drupal supports granular security down to the node level.
- Drupal assigns access rights to both the individual nodes and the containers that collect the nodes (such as folders in a file hierarchy).
- Drupal defines roles and, when enforcing a login process, authenticates users to roles.

As a result, Drupal grants permissions to content based on roles and content types in a granular, flexible, and scalable fashion.

A Drupal administrator can configure a Drupal web site with the roles and permissions that correspond to the way the organization works. For example a group of writers in a marketing department can have editing privileges for all the text-oriented content they produce, and can only access and view the standard product descriptions developed by writers in the engineering department.

Drupal modules can include their own roles and permissions, providing additional flexibility for the underlying security. The roles and permissions for blogging can be defined entirely differently than updating the home page for the corporate web site.

Illustration 2 provides an example of the range of permissions and roles available within Drupal.

Permissions	anonymous user	authenticated user	administrator	blogger
blog module				
create blog entries	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
delete any blog entry	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
delete own blog entries	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
edit any blog entry	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
edit own blog entries	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
comment module				
access comments	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
administer comments	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
post comments	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
post comments without approval	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Illustration 2. Drupal features a flexible and extensible set of permissions and roles. Drupal developers can define all the permission sets and roles that meet the business situation.

Integrated Workflow

Drupal features integrated workflow capabilities, based on users roles and permissions. A workflow designer defines the workflow by describing how a particular type of content needs to move through different stages or workflow states. Then the designer maps how individuals in different roles work with the content in particular states. Each person in a role interacts with his/her predefined permissions and workflow states. Finally, the workflow supports actions and triggers. It can automatically send notifications when workflow states are changed or perform other predefined actions, specified by a business rules engine.

Thus for example, a company can implement a press-release review workflow for publishing press releases on the company web site.

- A writer can draft the press release and send it to the communications manager for review and approval.
- Once the communications manager approves the press release, it can be forwarded to the legal department, where any one of a group of in-house attorneys can approve it for publication.
- The press release is finally forwarded to the publisher for posting on the appropriate section of the web site.

The Drupal workflow triggers alerts. The workflow manages all of the interactions as the content flows from one state to another, and as people in predefined roles perform their actions.

Business Outcomes

In short, Drupal manages web publishing in a flexible fashion, designed to meet the needs of business situations. Authors, journal editors, photographers, musicians, graphic designers, videographers, producers, interactive game designers, and other kinds content creators can use their familiar editors and capture tools, to add content to Drupal and manage it within a shared environment. Drupal can adjust the access roles and permissions to the way that people need to work together and share information.

Drupal automates the operational processes of digital publishing. It can ensure workflow and enforce retention rules – so that content is only published after being approved by predefined parties and then archived when it is no longer current. It intelligently categorizes content and makes it findable. It can project the distribution of content and the associated descriptors (or tag sets) to remote processes and thus enable wide-scale content syndication. And keeping abreast of the ever-changing world of media, Drupal can render content to browsers, PDAs, smart phones, tablets, and the variety of other digital devices gaining popularity for connecting to the web. And most important from a business perspective, knowledgeable end users (rather than technical specialists or web masters) can use Drupal to manage and maintain the business-related publishing and production processes.

Support for Social Media Tools

Drupal provides a wide range of modules that support the full panoply of social media tools – including blogs, wikis, threaded discussions, multi-media sharing, social tagging, voting, rating, and the like. It can manage not only content but also people and resources to develop the social graphs for social networking sites and to build directories of people and skills. Content managed by Drupal can be syndicated via RSS feeds.

Social media tools and techniques are continually evolving. Drupal modules developed by one set of open source developers are readily extended and enhanced by other community members to support additional social features. Drupal makes it easy to stay on the cusp of the social media revolution as open source developers are able to rapidly incorporate the latest technologies and social innovations into Drupal sites.

Drupal easily integrates groomed content, often managed by workflow processes, with user-generated content, produced by social media tools. All content types are nodes, with predefined metadata that are maintained by a common set of processes.

With appropriate access rights, editors can prune and organize the user-generated content, and create synergy with the groomed content, to reinforce the branded experience. Developers have flexibility to manage and blend the various kinds of content, regardless of the source, as best meet their business needs.

Thus businesses can easily support communities of contributors – their customers and partners who want to add their own perspectives to the groomed content. Using Drupal, businesses can structure the flow of user-generated content to build interactive communities that support business objectives.

Semantic Publishing

Drupal includes innovative capabilities for semantic publishing, and provides the ability to encode meaning into content available over the web. In addition to producing content for publication on web sites (featuring powerful full text search capabilities) and for syndicating content through RSS feeds, Drupal manages the tag sets used to categorize content.

Drupal can use RDFa as the underlying tag description standard. RDFa enables the content published on a Drupal site to gain a wider context, beyond the initial repository, and explain relationships to any other pool of knowledge or remote process that can be electronically referenced.

Drupal can interact with content extraction services running remotely on the web, such as DBpedia or Open Calais. These services auto-tag content by predefined taxonomies and ontologies. Each service encodes and encapsulates a particular area of expertise. Thus, a content extraction service automatically analyzes a content stream, extracts the key terms, maps them to predefined tag sets, and returns them to a Drupal site in RDFa format. The Drupal site can then expose these tags to search engine crawlers and other remote processes running on the web.

The end result is content that is widely distributed across the Internet in self-describing formats, leading to:

- Improve the capabilities of search engine crawlers and findability
- Improve content reuse
- Give context to tags for greater automated processing

Semantic web services provide the building blocks for increasing the value of content managed by Drupal and for building the semantic web.³

With semantic web services, businesses can widely syndicate and distribute their content across the web, highlighting the themes and contexts that they deem important. Businesses can improve how they optimize their content for remote search services, as well as how they can push their presence to partners and content aggregators.

Open Source Development

As an open source project, Drupal provides developers with the tools and resources for developing social publishing environment. Drupal makes it easy to enhance modules and support new technologies.

³ For instance, check out how to leverage Google Snippets, initially for improving the quality of Google's search results and then more broadly for encoding semantics into content streams. See: <http://www.google.com/support/webmasters/bin/answer.py?hl=en&answer=99170>

Adding Support for RDFa

For example, Drupal allows for the definition of disparate content types. Blogs, articles, reviews, forum posts, and so forth, can be added directly from the administrative interface. The Content Construction Kit (CCK) is used to extend these content types with the necessary data fields (text areas, drop down selects, file uploads, images etc.). Then Taxonomy module is used to add the proper tagging vocabularies to the content types so that the right metadata can be applied to the right content.

When developing content for the semantic web, it's essential to add the tags that describe the semantic relationships among terms, based on a predefined ontology. RDFa is a W3C standard for defining the terms of an ontology. Drupal developers have extended CCK to add support for tags defined by RDFa. Using a new module, RDF CCK, Drupal developers manage the RDFa tags on content nodes, and thus expose an additional set of tags to crawlers and other semantic applications.

For Drupal, semantically-aware content becomes just another set of custom content types. Business adopting Drupal can stay abreast of the continuing innovations to the underlying technologies of the web.

The Drupal Community Development Processes

Drupal evolves in an organic and distributed fashion, based on business needs and technical requirements. All types of open source developers – ranging from solo programmers to small, highly productive, development teams, and large groups -- develop modules for Drupal.

Drupal relies on a community development process for its ongoing growth and evolution. Developers not only build what they need but also contribute their modules back to the Drupal community. These community-contributed modules are then often incorporated into additional Drupal sites, tested in multiple environments, and updated with bug fixes and enhancements by various contributors (including the original developers).

As a result, Drupal developers can focus on the technical details of solving specific technical problems and addressing particular business requirements. They can then share their modules with a larger community of developers, who in turn can extend the initial set of features and functions to solve additional issues. As a module becomes more popular, multiple developers can make contributions and enhance it. Through the Drupal community development process, more eyes and more usage often result in faster innovations and lower development costs.

How Social Publishing Transforms Business Operations

What are the business implications of Drupal? Let's return to our initial example: how a nimble firm, designing and developing state-of-the-art home entertainment systems for price-conscious consumers, benefits from a framework for social publishing.

By leveraging the unique capabilities of Drupal, we believe that you transform your customer-facing operations. You semantically enable your content stream and harness an ever more powerful web infrastructure. When exploiting a framework for social publishing, you build your business and reinforce your branded experience over the web.

With Drupal as a social publishing platform, you are able to develop and deploy a seamless publishing environment that engages your customers and involves your community of contributors. You are able to maintain the transparency and trust for your groomed content – the multimedia clips and interactive product guides published over the web, as well as your product knowledge base maintained by your support organization. Through blogs, wikis, and other kinds of social media, your customers are able to communicate and share information with one another, build community, and help each other solve problems.

Moreover, through the semantic web, you reinforce your brand and your customer connections. Both your groomed and user-generated content is accessible through an integrated content infrastructure and is part of a unified user experience.

Not only can you create, manage, and deliver many different types of content over the web, but your customers can easily find what they want and need, and tell one another about their successes and problems. As you support your community of contributors, you enable them to spread the word and build your brand in the marketplace.

Finally, your investments match your results. By building on an open source framework, you keep your license and support costs modest, particularly when considering the value you extract from your social publishing environment and from engaging with your community of contributors. You leverage the collective expertise of the Drupal development community to solve your immediate business problems, and also to help you stay abreast of the rapid pace of technology change when building solutions for the digital age.