



## How to De-Risk the Creation and Moderation of User-Generated Content

### Executive Summary

Since the inception of the Internet, Web sites have enabled easy creation and distribution of user-generated content (UGC) to global audiences. The tremendous growth of UGC, following the advent of Web 2.0, has highlighted the need for Web sites that can more proactively alert organizations to the existence of mal-content (i.e., bad content).

How a Web site moderates its UGC is an essential part of its online brand identity. While some Web sites allow an open sharing atmosphere where almost anything goes, others ensure that UGC not only meets the highest standards but also reflects positively on the brand. The persona of a Web site is reflected in its approach to and consistency with determining which UGC is acceptable vs. not acceptable. To generate traffic, some Web sites may be very minimally moderated, often at the expense of overall public perception. Web sites with established, strong brands known to be family- and age-appropriate often require a more sophisticated approach toward UGC moderation.

While the explosion in UGC poses numerous opportunities, it is not without risk. With the ability of individuals, groups and machines to publish inappropriate, irrelevant or copyright-infringing materials, online companies must ensure that this content does not reflect poorly on the brand, negatively impact loyal customers or damage the bottom line. To mitigate UGC risks, content moderation must continuously

adopt and leverage scalable, effective and cost-efficient options to continue providing an enriched, socially-enabled user experience.

This white paper discusses the growth of user-generated content, the challenges of effectively moderating UGC content and how to think through approaching these challenges to provide the best long-term UGC moderation solution.

### Growth of User-Generated Content

*4.1 million minutes of video are uploaded to YouTube everyday ... six billion images per month are uploaded to Facebook ... 40% of images and 80% of videos [created] are inappropriate for business."*

Broadly speaking, UGC is any material piece of content that a user creates or leverages from existing sources and uploads or shares on a Web site for others to view (see Figure 1). UGC comes in many different forms, including short-text content such as tweets and forum comments; long-text posts on blogs and profiles; and multimedia material such as images, audio, video and applications (see Figure 2). Such content often further manifests itself as targeted or non-targeted online display ads, search engine results and archived Web content, as well as tags, posts or hyperlinks on various Web sites.

The online industry has adopted numerous business models that create, capture and deliver

## Typical Ways Users Generate UGC




Figure 1

## Types of UGC

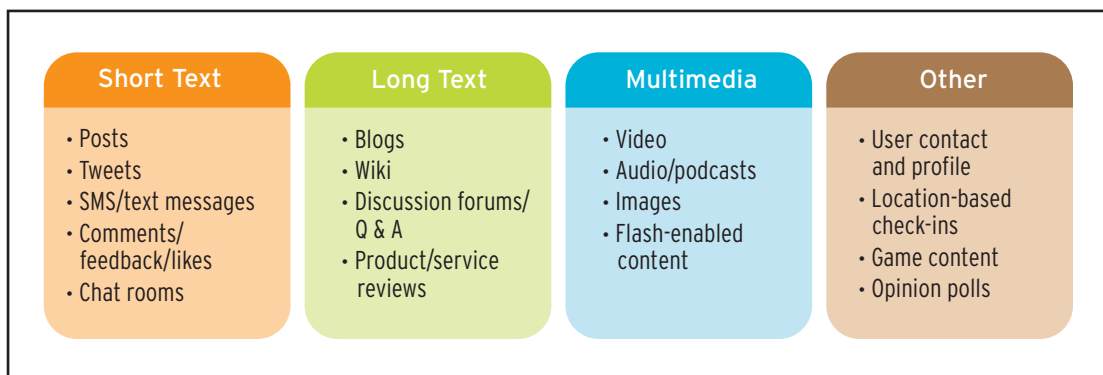


Figure 2

**The list of companies building communities around UGC to buoy their bottom line is large and growing, much like the number of users who create and consume UGC.**

business value. In recent years, the social phenomenon has become embedded within the DNA of many online-only companies. Web sites such as YouTube and Facebook rely on UGC (and the resulting visitor traffic) to attract and sell advertisements. Mobile and gaming companies provide platforms to sell third-party-created applications and virtual elements and collect a transaction fee for each sale.

Other companies, such as professional networking Web sites, sell access to user-created profile information. Various online retail and media companies thrive on user opinion, reviews and feedback to enable social and viral marketing/selling, as well as an enhancement for helping business partners sell their own product and service offerings.

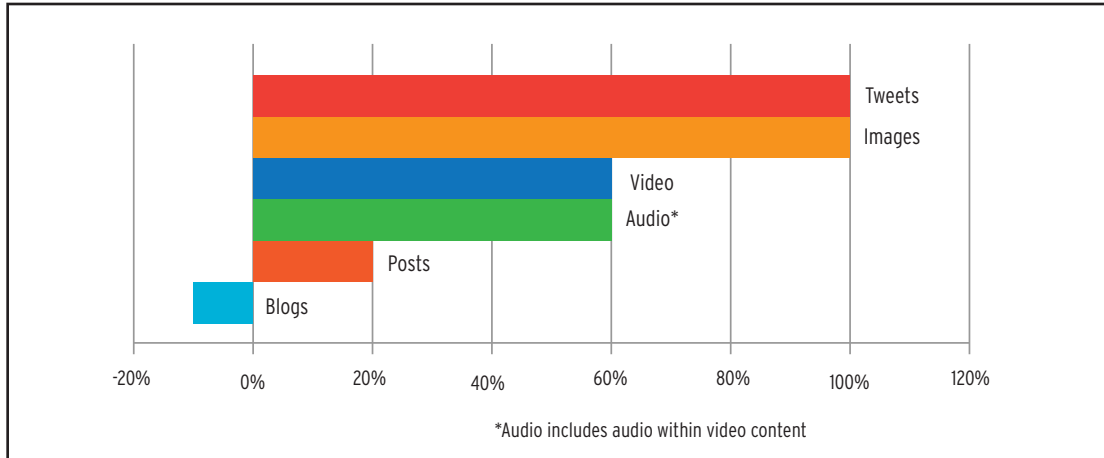
Some Web sites have been able to monetize UGC content beyond just selling advertising, such as

offering subscription access to user-generated business profiles, user-provided data and the like. The list of companies building communities around UGC to buoy their bottom line is large and growing, much like the number of users who create and consume UGC.

The monetary benefits of UGC are obvious: Content stickiness builds customer loyalty, sustained platform traffic that achieves critical mass, and top-line growth for companies that are able to monetize subscription services and page views via advertising or affiliate relationships. Non-monetary incentives include status-building with those who like and/or follow site content, network and relationship-building with viewers and affiliated sites, and content sharing/communication with viewers' colleagues and friends.

Users may also have the opportunity for financial incentives by creating UGC for Web sites that leverage crowdsourcing as a specific approach to content generation, which is the sourcing of tasks (in this case, the generation of content) to a group or community of people. Business models

## Projected Annual Growth of UGC (2011 - 2013)



Source: Cognizant research  
Figure 3

and new technologies such as mobile devices and cloud computing have made UGC creation and publishing more convenient, thereby amplifying UGC growth. Overall, UGC creation and distribution have grown astronomically (see Figure 3).

### Need for Content Moderation

It is crucial to ensure that only appropriate UGC is posted on the Web site by screening and filtering for mal-content. Failing to do so could severely impact user traffic, company brand and the bottom line. The huge growth and pervasiveness of UGC within companies' core online user experience poses potentially complex challenges and heightens unnecessary exposure to risk.

Strict policies must be set and applied to govern content authenticity, originality, privacy, political/social correctness and legalities, both locally and globally. Such policies should permit and promote

UGC and sustain viable social network interactions. Inappropriate content includes, but is not limited to, profanity, sedition, violence, bare skin, false and outdated information, spam and other inappropriate content.

Three main methods exist for Web sites to moderate content (see Figure 4):

- Automated moderation, using computer applications and algorithms.
- Community moderation, leveraging the online community to self-moderate content (such as flagging or volunteer administration).
- Human moderation, whether by a dedicated staff or crowdsourced.

Additionally, within each method, there are various ways to determine whether content is unacceptable. In many cases – in addition to validating

## Thinking Through Content Moderation

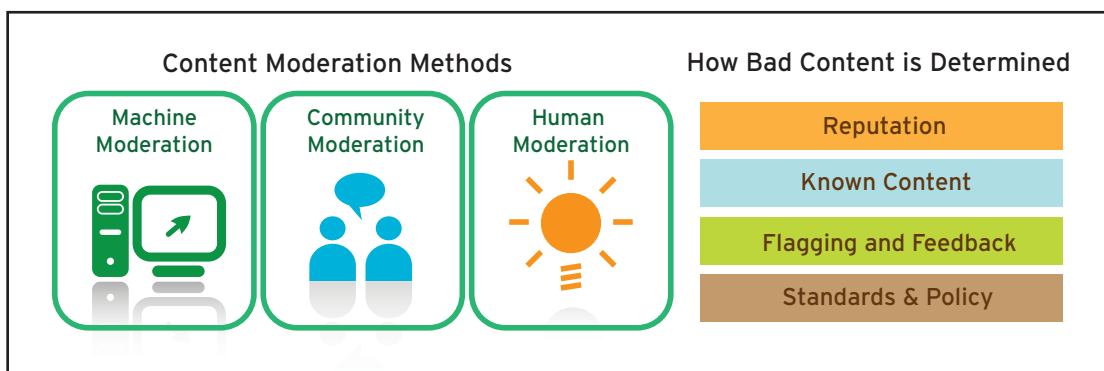


Figure 4

copyright or determining inappropriateness of UGC – content may also need to be moderated for quality, structure and relevance.

**While real-time moderation appears ideal, the associated costs may not be viable for all kinds of UGC; thus, a diligent analysis of the quality and cost tradeoffs must be performed.**

Lapses in content regulation can result in costly lawsuits from either original content rights-holders or offended Web site visitors in countries where these laws apply. Such lapses can additionally result in Web site traffic reduction, loss of advertisers and subscribers, as well as a poor user/buying experience that may severely impact future earnings. In 2007, YouTube was sued for nearly \$1 billion by Viacom for publishing copyrighted material.<sup>2</sup> Facebook has often been criticized for publishing posts and providing a platform to user groups that are politically or culturally sensitive.<sup>3</sup>

### Challenges in Content Moderation

There are many challenges to determining not only the optimal content moderation strategy that corresponds to the Web site's identity, brand and visitors but also how to put it into effect.

#### Cost, Time, and Quality Tradeoff

Organizations must be diligent when choosing the right mix of real-time moderation, pre-moderation and post-moderation of UGC on their Web sites, as well as whether or how to apply a combination of machine-automated and community- and/or human-moderated approaches. While real-time moderation appears ideal, the associated costs

may not be viable for all kinds of UGC; thus, a diligent analysis of the quality and cost tradeoffs must be performed. The dimension of moderation time further complicates the analysis, as specific UGC-like tweets and blog posts require real-time publishing to align with user expectations.

When to moderate is not the only question. How to implement moderation that is scalable and cost-efficient is just as important. For instance, should you be reactive or proactive (e.g., should you conduct post-moderation for all UGC or just for the content reported on or flagged by users)?

#### Inefficient Moderation Techniques

A large percentage of content moderation costs can be attributed to process inefficiencies. Choosing the appropriate moderation technique is also critical. Effective methods include automated algorithms, such as Bayesian filtering and pattern detection of blacklisted words and phrases, color tone and user/location profiling. However, most automated techniques do not moderate every piece of content but only samples of it; this can lead to mal-content leakage.

Consider scenarios where too few image samples of a video piece are taken or when the script is in English but is merely a translation of another language. In these cases, automated moderation is insufficient, and either human or community moderation is also required. This is easier said than done.

Human moderation, although effective, can be highly inefficient if one has to continue moderating the same UGC in different formats or if multiple moderators must continually track previously

### Cost Estimates by Content Type

Content Type	Estimated Average Size (per piece)	Estimated Moderation Time (per piece)	Approximate Machine Moderation Cost (per 1,000 pieces)	Approximate Manual Moderation Cost (per 1,000 pieces)
Video	6 min (100 MB)	1.7 min	\$2.61	\$277
Audio	6 min (5 MB)	1.4 min	\$0.13	\$230
Images	500 KB	0.4 sec	\$0.013	\$0.70
Text	200 words (200 KB)	1 min	\$0.005	\$167

Source: Cognizant research  
Figure 5

moderated UGC. The biggest challenge with human moderation, however, is the lack of scalability, which is an issue given UGC's proliferation and the high cost of sustaining such operations.

Figure 5 estimates the average cost of moderation, assuming a modest rate for a human moderator. Depending on the moderation rules and policies that need to be applied, the cost of certain types of moderation may be significantly higher.

#### Holistic Moderation

Organizations must moderate not only the content but also the users, as mal-content is often the result of user ignorance or lack of awareness. Much mitigation can be achieved through user-friendly and upfront communication of policies and guidelines. Some Web sites allow the possibility of self-moderation through tagging, filtering and warnings. Ideally, the more upfront moderation with simple computerized checking and labeling, the less downstream volume and impact to content moderation processes after submission.

#### Localization Challenges

Web sites may often need to serve global users across countries and locales. In order to determine the appropriateness of text-based UGC, moderators must be able to understand the language in which UGC is written, as well as the content's localized context and intent. Furthermore, perception of the content's inappropriateness may change based on the acceptable norms of the locale in which the Web site and the owner reside. Content that is deemed appropriate in the U.S. may be perceived as highly inappropriate in areas within Europe, Asia and the Middle East. In certain parts of the Middle East, for instance, any degree of skin revealed on a woman is unacceptable. In certain Asian countries, the color red may be perceived as inappropriate. Acceptable speech in different countries may also vary, especially pertaining to elements of politics or religion.

Likewise, perception may even vary in demographic groups within a single geography. In a growing number of cases, the requirements for content moderation are starting to become more "hyper-localized" to focus on UGC from specific regions or populations, thereby mandating the need to create separate rules and guidelines for each locale or user group. The personalization and relevance of more hyper-localized Web sites and content will not only increase the demand for hyper-localized content moderation but will also

result in expanding the amount of UGC content created, thus increasing the demand for content moderation even further.

While the challenges involved in moderating UGC are multifaceted and complex to navigate, it is vital to choose the right combination of moderation techniques. These combinations will be governed not just by the dimension of accuracy but by the total cost of operations, as well as moderating time required.

#### The Right Content Moderation Approach

Content moderation has grown into a discipline that requires expertise in pattern detection and analysis. Although there are numerous software-based solutions in the market, they do not address the custom needs of particular businesses. With the right level of investment in moderating content, Web site operators can create an optimal strategy that maximizes customer satisfaction while minimizing abuse and impact on the company's brand and reputation (see Figure 6).

With the increasing growth of UGC, as well as the technology, cost and scale needed to moderate it, achieving an optimal long-term solution requires detailed strategic planning and execution. A variety of options currently exist to help Web site operators protect and optimize their investments and reduce increasing moderation costs. Many organizations have outsourced their content moderation operations to reduce costs and enable more scalable and predictable business outcomes. Others have implemented both custom and standardized technology options to replace existing technology or to cut development and maintenance expenses.

Some companies have also experimented with content moderation crowdsourcing solutions to replace human moderation, with mixed success. While seemingly leveraging an unlimited

**The biggest challenge with human moderation, however, is the lack of scalability, which is an issue given UGC's proliferation and the high cost of sustaining such operations.**

**Ideally, the more upfront moderation with simple computerized checking and labeling, the less downstream volume and impact to content moderation processes after submission.**

## Content Moderation Decision Framework

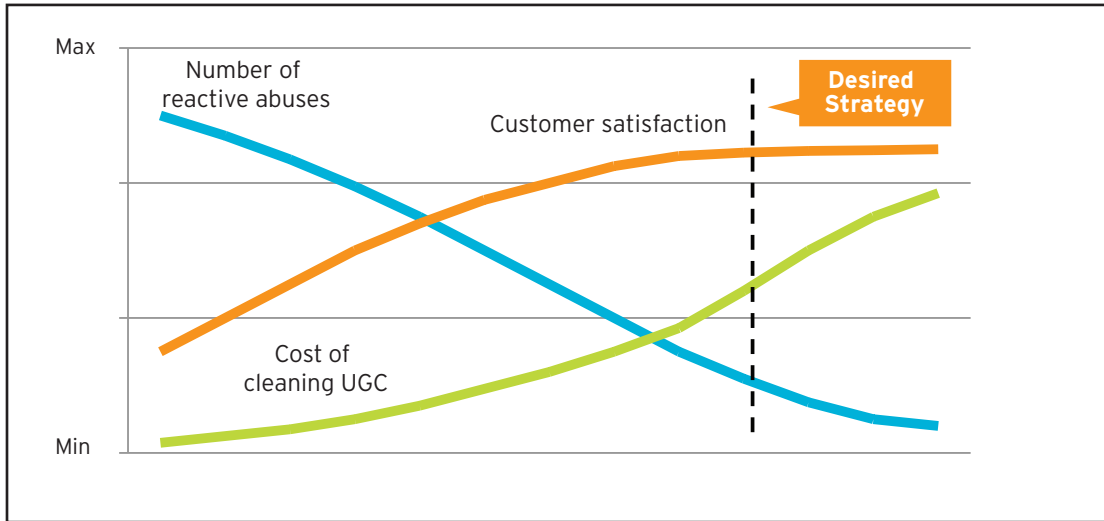


Figure 6

number of resources at minimal cost may seem ideal, crowdsourcing to date still produces poor moderation quality. Other Web sites have gone purely with community moderation to reduce costs, but this also may produce mixed results, as the moderation of UGC can become overwhelming even for members.

Thus, many solutions are available for content moderation and evaluation, implementation and

management. However, finding the right content moderation solution could be a difficult endeavor without solid strategic advice and a well-thought-out approach, leveraging industry best practices customized for specific needs, as well as a clear understanding of the objectives and ecosystem of your Web site.

### Footnotes

- <sup>1</sup> Quotes from Digitalrecognition.net and various sources, Nov. 15, 2011.
- <sup>2</sup> Anne Broache and Greg Sandoval, "Viacom Sues Google over YouTube Clips," C-NET, March 13, 2007. [http://news.cnet.com/Viacom-sues-Google-over-YouTube-clips/2100-1030\\_3-6166668.html?tag=mncol;txt](http://news.cnet.com/Viacom-sues-Google-over-YouTube-clips/2100-1030_3-6166668.html?tag=mncol;txt)
- <sup>3</sup> "Criticism of Facebook," Wikipedia, Nov. 1, 2011. [http://en.wikipedia.org/wiki/Criticism\\_of\\_Facebook](http://en.wikipedia.org/wiki/Criticism_of_Facebook)

## About the Authors

*Abhijeet Khadilkar is a Director within Cognizant's Business Consulting Practice, where he advises technology companies on sales enablement and business transformation. Abhijeet can be reached at [Abhijeet.Khadilkar@cognizant.com](mailto:Abhijeet.Khadilkar@cognizant.com).*

*Tom Pai is a Manager within Cognizant's Business Consulting Practice and advises technology companies on customer experience, online operations strategy, mobile marketing strategy, social media enablement and enterprise business challenges. He has more than 12 years of experience working for such companies as Deloitte Consulting, Accenture and IBM, serving a number of clients including Cisco, Adobe, GameStop and Wellsfargo.com. Tom can be reached at [Tom.Pai@cognizant.com](mailto:Tom.Pai@cognizant.com).*

*Shabbir Ghadiali is a Manager within Cognizant's Business Consulting Practice and is focused on operations enablement of new business models, including cloud and mobile computing. He also specializes in online retail, channel strategy, sales and service operations. Shabbir can be reached at [Shabbir.Ghadiali@cognizant.com](mailto:Shabbir.Ghadiali@cognizant.com).*

*The authors would also like to acknowledge the contributions of Charlie Goldenberg, a Vice President within Cognizant Business Consulting, and Matt Sugerman, a Senior Executive within Cognizant's Technology Business Unit.*

---

## About Cognizant

Cognizant (NASDAQ: CTSH) is a leading provider of information technology, consulting, and business process outsourcing services, dedicated to helping the world's leading companies build stronger businesses. Headquartered in Teaneck, New Jersey (U.S.), Cognizant combines a passion for client satisfaction, technology innovation, deep industry and business process expertise, and a global, collaborative workforce that embodies the future of work. With over 50 delivery centers worldwide and approximately 130,000 employees as of September 30, 2011, Cognizant is a member of the NASDAQ-100, the S&P 500, the Forbes Global 2000, and the Fortune 500 and is ranked among the top performing and fastest growing companies in the world. Visit us online at [www.cognizant.com](http://www.cognizant.com) or follow us on Twitter: Cognizant.



### World Headquarters

500 Frank W. Burr Blvd.  
Teaneck, NJ 07666 USA  
Phone: +1 201 801 0233  
Fax: +1 201 801 0243  
Toll Free: +1 888 937 3277  
Email: [inquiry@cognizant.com](mailto:inquiry@cognizant.com)

### European Headquarters

1 Kingdom Street  
Paddington Central  
London W2 6BD  
Phone: +44 (0) 20 7297 7600  
Fax: +44 (0) 20 7121 0102  
Email: [infouk@cognizant.com](mailto:infouk@cognizant.com)

### India Operations Headquarters

#5/535, Old Mahabalipuram Road  
Okkiyam Pettai, Thoraipakkam  
Chennai, 600 096 India  
Phone: +91 (0) 44 4209 6000  
Fax: +91 (0) 44 4209 6060  
Email: [inquiryindia@cognizant.com](mailto:inquiryindia@cognizant.com)