The Security and Privacy Implications of Using Social Networks to Deliver Healthcare

Carrie Gates
CA Labs
Islandia, NY
carrie.gates@ca.com

Matt Bishop
University of California Davis
Davis, CA
bishop@cs.ucdavis.edu

ABSTRACT

Healthcare technologies have tended to focus on electronic health records and devices (e.g., devices within the home for patients or handheld devices for nurses and physicians), and the interaction between the two. However, no one to date has investigated how social networking technologies might be used to provide an assistive environment for patients who participate in group therapy. In this paper we propose such an environment and go on to discuss the privacy requirements and security implications in developing an appropriate support mechanism.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

design,management,security

Keywords

healthcare technologies, security, privacy, social networks

1. INTRODUCTION

Research in healthcare technologies has focused on protecting patient privacy, the integrity of the mechanisms used to report and analyze patient data, mechanisms to ensure that data and the results of the analysis are available, and that legal compliance requirements have been met. The focus from a patient perspective has been on individuals. Further, the focus has been on the collection and analysis of physical traits, such as heartbeat, or the taking of medication. What has not been addressed are the requirements from a psychological perspective. For example, there are many patients who require group therapy as part of their treatment. This can include patients with mental illnesses that are not severe enough to keep them hospitalized or under constant care. Another such group might be recovering addicts who attend meetings such as Alcoholics Anonymous.

From a technology perspective, social computing and social networking technologies have demonstrated their utility in facilitating human contact and connectedness, even when individuals are geographically distant from each other. Families keep in touch using Facebook, for example when their daughter is a military officer serving in a foreign country. Colleagues keep up-to-date on each others' professional lives using sites like LinkedIn. From the early 1980s until very recently, USENET newsgroups provided forums for the exchange of information on technical and non-technical matters, and in some cases places to discuss personal problems.

These "human contact mechanisms" provide support for people in informal ways. The ability to see one's child, to know that a friend still has a job, often provides peace of mind. But sometimes more direct support is needed, as in a therapist's office for depression or a doctor's office to learn how others cope with a serious illness. Traditionally, this type of support is done through physical meetings, with people present.

We propose to leverage social computing technologies in order to support patients requiring group therapy. Specifically, we suggest that existing social networking technologies can provide a basis for holding "virtual" therapeutic meetings, which members attend over the Internet. Further, the flexibility and ubiquity of the Internet enables patients to continue attending meetings when traveling. Finally, if a member is having a particularly difficult time, therapy can be provided by a number of therapists spread throughout the world, eliminating the need for a "night shift" or stand-by therapist. In fact, a therapist could monitor a member's online behavior, looking for indications of problems much in the way that an intrusion detection system looks for indications of attacks.

Research in healthcare technologies has focused on technologies to solve problems in newer, better ways. Secure protocols for the exchange of health related information, protocols to protect privacy when in a pervasive environment, and systems that block information not relevant to the treatment of patients have been discussed in the literature. This paper focuses not on how to implement the particular technology, but rather what considerations should drive the use and development of technology to provide this support.

The rest of this paper begins with a discussion of what support groups look like, and might look like, on social net-

works. Privacy requirements are central to the provision of appropriate medical and counseling services, and those derive from the nature and purpose of the groups. Similarly, some security considerations arise: being able to prove membership without revealing identity to anyone outside (and, in some cases, inside) the group; and determining whether criteria for group membership are met. We conclude with a look to the future of these groups.

2. SUPPORT GROUPS ON THE NETWORK

Support groups on the network fall into three main classes, organized around control: self-moderating groups, facilitated groups, and moderated groups.

A *self-moderating group* has no leader; the group members act to support and help each other as they feel appropriate. There is no vetting to join the group; anyone who can access it can participate. Some such groups work well. For example, many weight-control web sites offer "communities" where members can post their struggles and other members can provide help and support. Others, notably many USENET discussion groups, have imploded because some of the members have been destructive, belittling others' efforts and posting derogatory comments. This drives the people who are trying to get support, and give support, to other arenas. For our purposes, the key characteristic is that all members can see all messages to all members.

The second type of group is a *facilitated group* that has a facilitator; the facilitator helps the group members interact, but does not direct or provide insight. The facilitator, or the facilitator and members of the group, decide the membership. In some cases, anyone may come to the meetings, and the group may accept them if they wish to stay (for example, Alcoholics Anonymous works this way); in other cases, prospective members may have to be invited or be voted in. Thus, the facilitator and members of the group can discourage the type of malicious behavior that members of a self-moderating group may take. For our purposes, the key characteristic is that the facilitator may delay messages in order to send her own messages, thereby ensuring her messages arrive first. This allows the facilitator to smooth over interactions that require intervention.

The third type of group is a *moderated group* that has a moderator acting as leader; the moderator controls interaction, providing insight and direction to the group members. The moderator decides the membership of the group. All communications go through the moderator, who can decide to block a message, return it with commentary to the sender, forward it with commentary to the full group, or merely pass it along. This corresponds to a group in which the moderator exercises control over what the group members see, and how they interact, as a group. In traditional groups such as USENET newsgroups, the moderator simply acts as a filter, deciding which messages to post and which not to post. In therapeutic groups, the moderator may draw out members' comments and thoughts, and interject ideas, suggestions, and other comments of her own.

When these groups exist virtually, there are several element that distinguish them from groups that meet physically. First, the membership may be much broader; for example, people who cannot leave their homes (for reasons such as acute agoraphobia) can participate. Secondly, members may be geographically dispersed. A virtual group may consist of people from France, French Polynesia, and Tunisia; or people from California, Maine, Alaska, Hawaii, and Texas. Third, care may be available at any time of the day or night. Rather than having a group facilitator or moderator on call 24 hours a day, or rotating moderators some of whom must work an undesirable night shift, facilitators and moderators can be chosen so that they work during their desired hours, by selecting some who are in other parts of the world.

The ability to consolidate geographically distant people means that treatment and therapy can be cost-effective. For example, suppose several people scattered throughout the United States have a particular psychological condition that can best be treated by group therapy, with other who have the same condition. Without virtual meetings, it is unlikely that group therapy will occur. With virtual meetings, group therapy becomes possible. Similarly, if the follow-on to medical treatment is best handled in a group setting where patients can support one another, the geographic separation inhibiting physical meetings will not inhibit virtual meetings.

The selection of moderators external to the country in which a member, or members, live raises challenging issues. First is the issue of *information governance*. Suppose the members of the group are American, and one moderator is not. This is a moderated group in which the moderators are providing psychological treatment, so the moderators need access to the medical records of the members of the group. It is unclear under what conditions, if any, United States law applies to the moderators not in the United States. Equally unclear are the ramifications if it does not.

An equally interesting issue is the selection of moderators. Given different cultural backgrounds of the moderators and the members of the group, the moderators would need to be trained to understand how the culture in which the members live affects their particular problems. The ability to bring the perspectives of different cultures and forms of treatment may be a benefit; it may, on the other hand, inhibit treatment. This issue is critical to the success or failure of on-line support groups. However it is outside the scope of this paper.

A number of different existing social networking forums can be used to provide support for those requiring it. In what follows, we use Facebook as our example social networking site. Other sites exist, such as MySpace, Twitter, and World of Warcraft and other gaming communities that may be more appropriate depending on the particular purpose of the group and the nature of its membership.

In order to implement a group support system, there needs to be a balance between privacy (or anonymity) and known information. By its nature, participation in an online group requires the sharing of some form of identity, even if that identity is limited to a username or handle. As in the real world, the accuracy or information known about a particular identity will vary based on the requirements of the group. For example, participation in a group providing support for people on a diet might only require a pseudonym, and no

one need ever know the true identity or any personal details about the participant. In contrast, participation in a group that provides support to people with mental health issues might require that at least the attending therapist know the identity of the participant, along with details relevant to his circumstances and disease.

3. PRIVACY REQUIREMENTS

Privacy requirements vary depending on the nature of the group, the membership, the leaders, and the condition or conditions being treated. Further, the requirements typically differ between the privacy members require when dealing with other members, and the privacy members require when dealing with leaders. We note that medical and psychological treatment groups rarely lack leaders, so we do not speak to self-moderating groups in this context.[1]

Privacy requirements affect three types of data. First is information that can lead to the identification of an individual; this type of information may not be private, if the group does not accept anonymity. Second is information about the particular patient's condition. Again, in a therapeutic setting, this may not be private, especially if other members can infer details from the group therapy. Third is information unrelated to the other two types. This may or may not be private, depending on the needs of the patient.

Consider first identity information. The simplest case is information uniquely associated with that individual, such as a name and address or (in the United States) a social security number. This information can simply be withheld from the group, and (if appropriate) shared only with the leaders. The more complex case is information from which the identity of the member can be derived. For example, suppose the member is the only night watchman at the Wicket Factory. Should she reveal to group members that she works nights, and at a later time that she works at the Wicket Factory, deducing her identity becomes simple. In general, the complex case requires knowing what external information (that is, information not obtained from group interactions) is available, and how that information relates to identity. This problem relates to the data sanitization problem, and the approach suggested by Bhumiratana and Bishop [1] might prove fruitful here.

Information about the patient's condition affects the patient's ability to join the group. This information must be shared with the leader(s), or whoever makes the decision to allow or disallow the patient to join the group. For example, a member needs to demonstrate that she has the conditions that the group is meant to address before she is allowed to join the group. Additionally, the leader needs to decide if this potential member not only meets the conditions, but should be allowed to join the group based on other considerations, such as how severe her condition might be, or how much benefit she is likely to gain through group therapy. However, the information should not be disseminated further without the consent of the patient. This becomes a problem in originator-controlled access control (ORCON) [6].

[1]The requirements for such a group would be the same as for one with leaders, except that the considerations for leaders would not hold.

Governance issues surround the release of information to facilitators, and these issues impact the privacy of members. Although ORCON implies that the member will control the dissemination of the information, laws, customs, and other matters[2] may override that control. More specifically, given that the support group may need to have leaders available 24 hours a day, 7 days a week—and, indeed, this is one of the advantages of having an on-line support group—the group leader may not necessarily be physically in the same country as the member. In such a case, the laws governing the collection, use and retention of the health data of a member involve two legal jurisdictions with possibly conflicting laws. It is unclear which set of laws will apply—or whether *both* will apply. And it is equally unclear whether the member will understand this. Thus, one aspect of using social networks for group therapy and treatment is an understanding of the laws involved, and an ability to ensure all members understand what privacy, if any, they have.

Conversely, the group leaders, whether facilitators or moderators, may want to preserve their privacy from members of the group. This would give the leaders privacy, preventing members of the group from contacting them outside the social network. Even within the social network, they may desire anonymity, separating their personal *persona* from the *persona* they use in the group, in order to prevent harassment or other undesirable consequences. This raises converse issues, and the problems with identity stated above apply equally well here.

Social networks such as Facebook typically capture communications and save them so they can be viewed at a later time. This means that dialogue or conversation in group sessions are not ephemeral. They can be saved, and replayed later. In addition to the privacy issues raised above, that this recording can be made raises an issue of trust. No longer can one make a "passing comment"; the comment will be saved for later perusal by those for whom it is intended (and, possibly, other bystanders).

Underlying these issues is trust: the members trusting the group leaders, and the leaders trusting group members. The trusts are of different types, because the roles of the leaders and members are different, and their powers and responsibilities are different. Belief logics capture this notion, because ultimately belief—in identity and in abilities—is what drives trust in this environment.

Trust is a factor not only in privacy, but also in security considerations. It is to those we now turn.

4. SECURITY IMPLICATIONS

In addition to privacy considerations, various security considerations arise. They are based on the differing characteristics of the three types of groups.

First, consider privacy. We consider "privacy" to be a subset of security, specifically that part of security enabling a person to control the dissemination of personally identifiable information (PII). This is essentially a problem of imple-

[2]For example, in the United States, disclosure to insurance companies to obtain payment.

menting originator-controlled access control, something that is required for digital rights management (DRM), among other controls. The privacy and security issues inherent in allowing users to control their own information, both in a health context and more generally, have been discussed in other papers (e.g., [5]).

However, in many jurisdictions, privacy controls cannot be absolute. Consider a facilitated or moderated group, with one or more leaders (facilitators or moderators). If the leader is not someone to whom confidentiality applies (such as a medical professional)[3], the leader may have a legal requirement to report certain types of messages, such as threats to harm oneself or another. In some jurisdictions, *even if* legal confidentiality applies, if a patient informs her psychiatrist or psychologist that she intends to harm a third party, the professional must use "reasonable care" in order to protect the third party [3].

Even in groups in which the members do not enjoy the privilege of legal confidentiality, reversing anonymity can be important. A simple example occurred in 2009, when a woman posted in a Twitter group belonging to actress Demi Moore her intention to commit suicide. The police in the woman's home town were contacted and intervened; the woman was hospitalized [4].

This implies that, in some circumstances, group leaders must be able to violate privacy. It implies more, though—that the ability to attribute statements in the medium used to support the group is accurate. For example, if Alice and Bob are members of the group, statements made by Alice should be attributable to Alice and not to Bob, regardless of the nature of the group (or of whether "Alice" and "Bob" are pseudonyms). Attributing to *people* as opposed to network (IP) address is a non-trivial problem, and indeed in some specific cases may be inappropriate. Thus, the type of attribution desirable for the group must be considered; once that is understood, then the implementation may be undertaken. Both are non-trivial problems [2].

Group membership may be constituted in many ways. Some self-moderating groups may have members who simply show up occasionally, much as unmoderated USENET groups work. Others, including most facilitated and moderated groups, would require that attendees be identified and confirmed as a member of the group, even if their identity is not revealed. This can be accomplished using two forms of credentials.

The first form uses a single credential identifying the subject as a member of the group. If anonymity is supported, the credential can be a *persona* credential issued by the group leader (or designated credentialing authority) and attesting that the person to whom it is issued is a member of the group. The subject field would be meaningless; the issuer field would provide the assurances, just like the *Persona* certificates in the certificate hierarchy used for privacy-enhanced electronic mail [7]. Otherwise, a certificate identifying the person and issued by a trusted, authorized certification authority would suffice.

The second form uses dual credentials and is appropriate when the leader must know the members, but the members can remain anonymous to the other members. The first credential is a *persona* credential as discussed above, and is used to sign messages to the group. The second credential identifies the person explicitly, and contains the serial number (or other identifier) of the *persona* certificate. This one is given to the leader when the member joins the group. Thus, the leader can determine who is sending messages, and act accordingly; but the group members cannot determine the identity of the sender. Note that group members *can* identify multiple messages sent by one person as having been sent by one person. This emulates what happens in groups that meet physically.

A second question is the *bona fides* of the leaders (facilitators and moderators) of the group. Considering that, under some circumstances, members of the group will impart very sensitive information, there must be some means for them to know that the group leaders will honor their confidences (within the limits of the law, of course). It is not possible to establish technical mechanisms to prevent information from being wrongfully disclosed, but it *is* possible to use techniques such as watermarking to determine whether the leader or another group member leaked the information.[4] As an aside, it is important that the group members believe that the leaders are keeping their personal information in confidence in order to establish the trust relationship necessary for successful therapy, and we note that this is trust is formed from personal interactions rather than through technical security mechanisms.

A related issue is the sharing of information between leaders, which is an issue that does not have an analog in the physical world. Support groups that employ social networking media will have multiple leaders due to requiring support through different time zones and their 24/7 nature. In addition to the online discussions, which any of the leaders will be able to access, leaders may have additional information about various group members (as discussed above) and so this information needs to be shared among the leaders using a channel that is related to the group, yet not accessible by the group members. This requires that appropriate authentication and access control mechanisms be in place. It might be the case that a leader feels that he needs to leave a comment about a group member for other leaders, in which case the other leaders will need to have the ability to confirm the identity of the leader leaving the comment.

Credentials that provide cryptographic keys enable messages to be sent with integrity. If secrecy is required, any of a number of group sharing cryptographic schemes can be used. A more interesting question is raised by the way Facebook works, specifically in that it preserves messages once they are sent for all to see. While the theory of protecting these messages is well understood, implementation errors or user interface problems may compromise their integrity. Worse, the user interface may confuse users so they are unclear

[3]The exact class of people to whom legal confidentiality applies is established by the laws of the appropriate jurisdiction. As noted above, *which* jurisdiction is appropriate depends on many factors—and even then may not be clear.

[4]Basically, each message is marked differently and then signed. If a message is leaked, the mark identifies the recipient.

about *what* is being protected, and unknowingly set their protections to allow past messages to be changed. Hence, the user interface should be simple—for example, not allow messages to be changed or deleted by anyone. Further, it should not be possible for the access control on a message or set of messages to be changed so that an individual outside the group can see the message.

Should the leaders be able to change or delete messages? This depends on the nature and purpose of the group. In a moderated group, the moderator should be able to block messages considered inappropriate, and discuss them with the sender. In many cases, the moderator should control the transmission of messages to ensure that no inappropriate messages are sent; in other cases, the moderator may only be able to see "side-bar" conversations and comment on them; in still other cases, the moderator should be able to block or terminate side-bar conversations. In those cases where a message is deleted or there is side-bar conversation, the other leaders of the group will need to have access to this information so that they have full context for any later discussions with the associated group member or members.

A facilitated group may be like a moderated group, but without the facilitator being able to block messages. The facilitator can comment on the messages, and urge the sender to send follow-ups to correct the message.

In both cases, the communications medium must be reliable and available. The "medium" includes the repository of messages so members and leaders can see them when needed. An additional requirement may be that at least one leader (or, in a self-moderated group, members) be available at any time, to help other members. Here, the advantages of a global Internet come into play, as noted above.

5. CONCLUSIONS

Social on-line networks have grown in popularity, complexity, and capability. Facebook was launched in 2004, and now has tens of millions of users. MySpace and LinkedIn have demonstrated similar explosive growth. Generally, these networks have expanded social circles, focusing on friendships (Facebook, MySpace) and professional and business contacts (LinkedIn). The older USENET network, essentially a global bulletin board system, has groups that were intended to provide support to people with various problems.[5] So, there is precedent for on-line groups to provide support for people; from there, it is only a short leap to providing medical and psychological assistance on-line in groups.

The advantages of these groups were mentioned above. But there are potential problems in protecting privacy and providing security commensurate with the needs of the group members and leaders. These considerations are crucial for the groups to succeed, both in the sense of being able to obtain members and in the sense of helping those members as much as would groups that are physically together. These issues must be considered both as the rules for the group are being constituted, and as the group evolves over time, because as time passes the need for new rules may arise, and the type of privacy and security desired may change.

[5]We note that the USENET groups were typically self-moderating, and when "flamers" came along to stir up members of the group, those who needed support simply left.

The issue of governance, mentioned earlier, is central to understanding the issues and framing solutions. Governance describes who has control of information, who has access to information, and how people can access that information. Governance of a group that meets physically raises information management issues that are essentially local. Governance of a group that meets virtually, and has members scattered around the globe, introduces very different complexities and challenges.

The goal of this paper was to raise these issues. The problem is not how to solve any given issue. Indeed, although some are easy to solve, such as message integrity and concealing simple identification information from group members, others are not, such as concealing complex identification information from group members and controlling the dissemination of medical information. The problem is *which* issues need to be solved, and this will depend upon the purpose of the group, the method of treatment that the leaders and members use, and other societal and legal issues. Once a particular group decides what those issues are, they can then determine what mechanisms—technological and procedural—can provide a solution acceptable to the members and leaders of the group.

Acknowledgement: Thanks to Barbara Langer for helpful conversations.

Matt Bishop was supported by grant CCF-0905503 from the National Science Foundation to the University of California at Davis. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

6. REFERENCES

[1] B. Bhumiratana and M. Bishop, "Privacy Aware Data Sharing: Balancing the Usability and Privacy of Datasets," *Proceedings of the 2nd ACM International Conference on Pervasive Technologies Related to Assistive Environments* (June 2009).

[2] M. Bishop, C. Gates, and J. Hunker, "Sisterhood of the Traveling Packets," to appear in *Proceedings of the 2009 New Security Paradigms Workshop* (Sep. 2009).

[3] *Tarasoff v. The Regents of the University of California*, 17 Cal 3d. 425 (1976)

[4] "Demi Moore uses her Twitter to help stop woman's suicide attempt," *New York Daily News* (Apr. 3, 2009); available at http://www.nydailynews.com/gossip/2009/04/04/2009-04-04_demi_moore_uses_her_twitter_to_help_stop.html

[5] C. Gates and J. Slonim, "Owner-Controlled Information," *Proceedings of the 2003 New Security Paradigms Workshop* (Aug. 2003).

[6] R. Graubert, "On the Need for a Third Form of Access Control," *Proceedings of the Twelfth National Computer Security Conference* (Oct. 1989).

[7] S. Kent, *Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management*, RFC 1422 (Feb. 1993).