Paper I: Open Wireless Networks on University Campuses

Open Wireless Networks on University Campuses

Although they raise privacy issues and entail security risks, open wireless networks enhance system usability and expand access to a university's nonsensitive system resources.

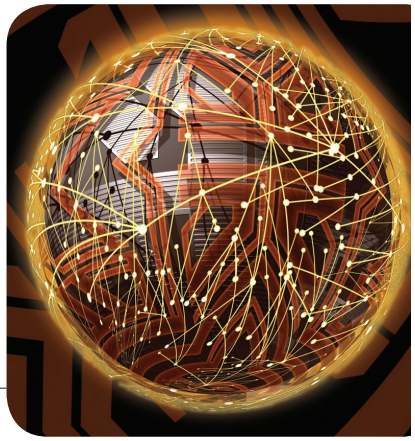

Kjell J. Hole, Lars-Helge Netland, Yngve Espelid, André N. Klingsheim, Hallvar Helleseth, and Jan B. Henriksen
University of Bergen

Universities are particularly interested in wireless communication networks because they let students using wireless-enabled mobile terminals download lecture slides, watch educational audio and video programs, and take online practice quizzes at any time and from anywhere on campus. This can both reduce paper handouts and simplify paperless assignments and submissions. Wireless networks can also strengthen teamwork among students and faculty, making it easier for them to email each other with preliminary results, use chat channels to discuss problems, and readily access information resources during problem-solving sessions.[1] When students and faculty can access all the information they need via their own mobile terminals, universities can even consider retiring their expensive computer labs.

To achieve these benefits, usability is crucial—that is, university information systems must make it easy for users to access resources and achieve their goals. Typically, universities employ individual authentication to give users system access. Such authentication mechanisms tend to reduce system usability, however, because users view them as both an intrusion and an obstacle to completing their primary tasks. Also, individual authentication often requires users to remember multiple log in names and passwords and sometimes to possess particular authentication devices.[2] Finally, authentication requires a trade-off between a user's privacy rights and an administrator's need to protect access to resources. On wired networks, university IT departments have long been capable of collecting information on authenticated users, including what they do and when they come and go. Introducing a campus-wide wireless network that uses authentication makes it even simpler to track user movements and activities on campus.

As an alternative, universities can opt for an open network, granting wired and wireless users access to the network infrastructure without any form of authentication. Open networks let users freely surf the Internet and access library catalogs and other services that offer nonsensitive information; sensitive services, such as email, would still require authentication. An information system's degree of openness largely depends on the number of different devices that can access the system, the number of services available without authentication, and the network's availability.

Open networks increase system usability, but they also raise privacy issues and increase the risks of illegal downloads and various attacks (and the negative press coverage that can result). We discuss these benefits and drawbacks here, using a simple model in which campus-wide mobile terminals communicate over wireless links with services on the university's wired infrastructure. We also discuss how administrators can mitigate open-network risks and vary their network's openness to reduce these risks while still increasing usability.

System authentication

Individual authentication establishes an understood confidence level that an identifier—such as a name—refers to a particular individual. Many specific authentication techniques exist, including the traditional approach, based on passwords or passphrases; two-factor authentication, based on something the user knows (such as a

password) and has (such as a hardware security token); and authentication based on public key cryptography.[3] Here, however, we deal only with the general authentication process, which occurs in two phases:

- an identification phase, during which the user presents an identifier; and
- an authentication phase, during which the system establishes the required confidence level.[4] If the resulting confidence level is high, the authentication is strong.

Individual authentication differs from device identification, which seldom leads to strong individual authentication. For example, administrators can identify a user's mobile terminal by its unique Media Access Control address. They can then associate this MAC address with a specific individual, which is relatively easy if the IT department registers the MAC addresses of all new terminals given to students and faculty. However, the resulting individual authentication is weak for two reasons: first, it's easy to fake (spoof) a MAC address, and second, the terminal's owner can always claim that someone else was using the terminal during a particular time period.

Although we deal mainly with the user authentication process here, it has ties to user authorization as well. Authentication establishes who an individual is, whereas authorization determines what an individual is allowed to do. Organizations typically have an authorization policy that determines how authorization decisions are made. In our model, we assume that authorization occurs and that the authorization policy requires individual authentication of all potential users. Clearly, such mandatory authentication can reduce user privacy because it makes it possible for administrators to build a behavioral dossier on every user.

Network boundary authentication

Figure 1 shows an example university information system in which mobile terminals communicate with wireless access points to access servers on the wired backbone. The system authenticates users at the network perimeter prior to their accessing any offered services. The solid black line indicates the authentication boundary; resources within this boundary are available only to users who have authenticated themselves correctly.

Figure 1 shows the system from the mobile users' viewpoint. The red lines indicate possible information flows between entities, not how developers should implement the information system. Because network boundary authentication applies only to wireless users, service access is simpler from the Internet than from the wireless network. Also, the sensitive email service requires all users to enter an additional service-level authentication.

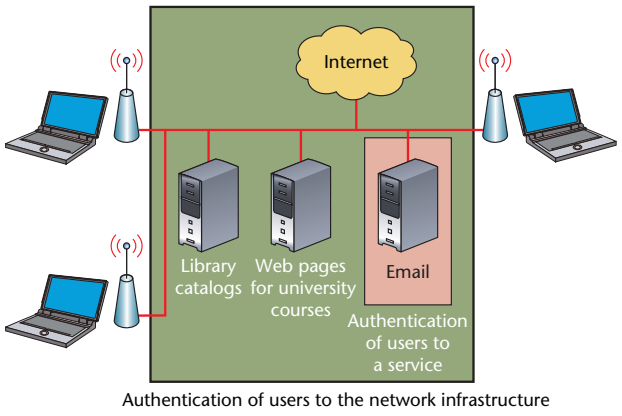

Figure 1. An example university information system. Resources within the network boundary are available only after user authentication.

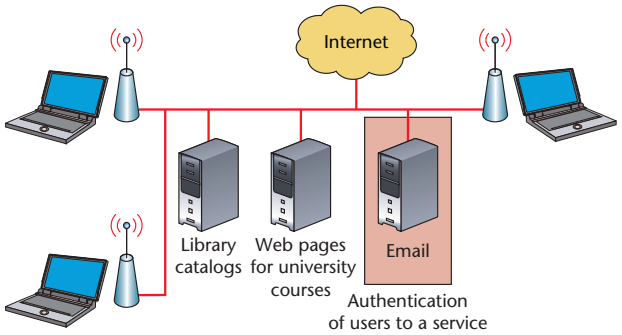

Figure 2. University information system with service-level authentication. Users must authenticate themselves to access sensitive services, such as email.

Service-level authentication

Figure 2 shows an example of an open network based on Figure 1's information system. In this case, there's no authentication at the network perimeter; users who can communicate with the wireless access points have unrestricted access to library catalogs, the Internet, Web pages containing course information, and other nonsensitive services.

Because users must authenticate themselves to each sensitive service, administrators can adjust the authentication's strength to fit a particular service's security needs. A service with highly sensitive information requires strong authentication, such as authentication techniques based on a hardware token or a public-key infrastructure, whereas those containing less sensitive information can use password-based authentication. To determine how to best select and implement appropriate authentication techniques for a particular system, we recommend that you consult the literature.[3,4]

Although Figure 2's system lacks individual authentication at the network perimeter, it still supports a defense-in-depth strategy. The perimeter includes security mechanisms—such as firewalls and intrusion detection systems—while the internal defense includes security mechanisms on the client and server machines where the information resides.[5] Administrators can also deploy monitoring techniques[6] to verify that all entities (including users), behave according to established security policies.

System usability

Deploying a wireless campus network increases system usability for insiders, such as students and permanent or visiting faculty members, as well as for outsiders, such as invited guests and the general public.

Usability for insiders

When a campus network has an open network perimeter, usability increases for students and faculty because they don't have to remember passwords or carry authentication tokens to access nonsensitive services. This can lead to increased use of information resources and, thus, a better learning environment and increased knowledge among insiders.

To achieve high usability in a campus network with boundary authentication, administrators must provide a well-functioning authentication mechanism for all wireless devices—including laptops, PDAs, and mobile phones. This can be a daunting task: popular wireless devices run multiple versions of different platforms (Windows, Mac OS X, Linux, Symbian OS, Palm OS, and so on), and each has its own unique challenges. Although open campus networks reduce this problem, giving insiders wireless access to sensitive services that require authentication remains a challenge.

Usability for outsiders

Public universities are important democratic institutions with a responsibility to make essential information easily available to the general population. Open wireless networks can give citizens effortless access to university library catalogs and other important information. Universities should avoid an authentication scheme's added complexity because people generally have rather limited understanding of security mechanisms and find it hard to authenticate themselves to a network. Strong individual authentication also requires that users provide information upfront to enable the authentication process, which further lowers the network's usability and adds to network operation costs.

At our university, employees must request temporary accounts from the IT department to give short-term guests wireless network access. Visitors participating in more loosely organized activities can't get wireless access unless they happen to know an employee who can set up their account in advance. Most first-time visitors don't even realize they need an account to get wireless access.

Offering an open wireless network is a much more user-friendly alternative, and such networks can even help narrow the digital divide. For example, universities with open networks in developing countries can offer citizens with little computer training easy Internet access via low-cost laptops and open wireless networks.

Risk analysis and mitigation

We define a *vulnerability* as a design flaw or system bug, and a *threat* as an adversary with the capabilities and intentions to exploit a vulnerability. (Although it's common to group adversaries into rogue insiders and outside attackers, this distinction obviously makes less sense on an open network.) A system's total *risk* is a function of its exploitable vulnerabilities, its threat severity, and the value of its information assets.[7]

To illustrate open network risks, we'll analyze the risk associated with wireless access of Figure 2's open information system. We'll also determine how removing authentication from the network's perimeter increases the risk, as well as how to mitigate it.

Illegal downloads

As has been widely reported, students and university employees sometimes engage in illegal music and movie downloads. Naturally, removing individual authentication from a network perimeter increases this temptation, so universities deploying open wireless networks must therefore present users with rules governing network usage and reserve the right to prosecute any user who causes economic loss or damage to the university's reputation through illegal downloads.

The promise of legal action against perpetrators will limit, but not eliminate, misuse. To further discourage illegal download activity, IT departments can monitor and log network traffic. As we describe later, it's also possible to "filter out" many illegal downloads by allowing traffic only on certain network ports.

In addition, some countries maintain child pornography filters that universities can deploy—for example, the Norwegian police and Norwegian ISPs jointly maintain a child pornography filter that warns ISP customers when they try to access such sites. The ISPs regularly update the filter using a domain names list provided by the police (in May 2007, the filter contained 4,235 domain names).

Terminal-to-terminal attacks

When students and faculty use mobile terminals, attackers can attack any terminal over a direct wireless link from their own terminals or via a wireless access point without going through the wired infrastructure. Attackers can also attack an open wireless net-

work from the wired infrastructure. They might, for example, install an unauthorized (rogue) access point that runs malicious software.

Today, however, mobile terminals that lack personal firewalls and updated antivirus software have likely suffered many attacks already, even if the user has connected only to wired networks. At this point, most terminals should already have the necessary protective software needed to access an open wireless network.

Rogue access points constitute a serious security problem because they can give attackers unauthenticated access to a university's information assets. Although this problem is reduced in open networks, a university's IT staff should still search for and remove rogue access points.

Attacks on local networks

When we remove Figure 1's network perimeter authentication, it makes it easier for people to attack services on the wired infrastructure using spoofing attacks, in which one entity illegitimately poses as another to gain access to restricted information. Introducing rules and threatening legal prosecution doesn't significantly reduce this risk—instead, we must introduce security techniques for sensitive and nonsensitive services.

When nonsensitive services are available without authentication, attackers can introduce false services—such as providing fake lecture notes or bogus research papers. Administrators can mitigate this risk by installing antivirus software and firewalls on their servers, and by running auditing programs, recording all user activities. Regularly reviewing audit logs can also help detect illegal activity and identify attackers. Finally, the IT department should be prepared to quickly reinstall and secure a Web server if an attacker subverts its defenses and modifies Web page content. Universities already offer Internet-based services, and thus should have experience and competence in both securing systems and dealing with security breaches.

To ensure that only legitimate users have access to sensitive services, universities need strong authentication. Because it's easy for anyone to sniff passwords on unencrypted wireless links, password-based authentication requires end-to-end encryption between the mobile terminals and the server. Administrators can mitigate spoof-attack risks on sensitive services by requiring users to access the services via encrypted virtual private networks (VPNs), Secure Shell (SSH), or SSL. The steps for mitigating risks on nonsensitive services also alleviate those linked to sensitive services. Highly sensitive services—such as those processing sensitive medical information—shouldn't be on any Internet-connected network.

We evaluated one campus network that used a VPN solution that didn't encrypt network traffic. As a result, it was easy to sniff usernames and passwords on the wireless links. Once attackers had an employee's username and password, they could download a central password file—containing thousands of password hashes—to a local machine. Using a password cracker, they could then obtain several hundred usernames and passwords. Because one of the vulnerable passwords belonged to an IT department network engineer, they could even escalate their privileges. As this example illustrates, university networks require a robust design that employs both strong authentication and strong encryption to protect sensitive information.

Anonymous attacks on remote networks

When administrators remove authentication from a university network's perimeter, attackers can use the open network to carry out anonymous attacks on information assets anywhere on the Internet. Such anonymous attacks include music and software piracy, identity theft, denial-of-service attacks, spam and phishing, and attacks on remote machines.

When analyzing this risk, it's important to realize that a university-based open wireless network doesn't represent a major new possibility for attackers to gain anonymous Internet access. Many cities now have numerous small wireless networks that private citizens and small businesses own and operate. For example, the city of Bergen, Norway, has many wireless networks (based on IEEE 802.11a/b/g standards) that are considered "open" because they don't use the Wired Equivalent Privacy (WEP) or Wi-Fi Protected Access (WPA) security protocols. Although the percentage of open networks in Bergen has apparently decreased since a 2004 study,[8] a 2006 survey by one of the authors indicated that roughly 40 percent of these networks were still open. Furthermore, the total number of wireless networks in Bergen has continued to increase since 2004, giving attackers many more open networks to choose from. The situation is similar in other cities around the world.

Also, attackers aren't limited to using open networks to obtain anonymous Internet access. A famous example here is Tor (tor.eff.org), an anonymous Internet communication system that uses a network of computers (proxies) to reroute users' Internet traffic. Multiple layers of encryption inside the Tor network protect traffic from eavesdroppers, letting individuals and organizations share information without compromising their anonymity. Tor was primarily created to anonymize Web browsing and publishing, instant messaging, and other applications that use the TCP protocol. Although it enables anonymous access that's clearly valuable to many individuals—such as rape and abuse survivors who want anonymous access to helpful chat rooms and Web forums, and journalists, who want to communicate safely with whistle-blowers and dissidents—it also lets attackers carry out anonymous attacks.

Using Tor and a mobile terminal with a spoofed MAC address, attackers can easily conceal their identities on a university's open wireless network. This approach is particularly effective for several reasons:

- The wireless network owner can't determine which remote system attackers are accessing.
- The wireless network owner can't determine the content of communication between attackers and the remote system.
- The remote system's owner can't determine the attackers' originating network.
- If attackers reveal their identities to someone on the Internet, that person still can't determine the attackers' home network.
- With so many active terminals on campus, it's difficult for the wireless network's owner to identify the attackers' mobile terminal.

The anonymity isn't completely fail-safe, however; the application layer can leak information. As an example, if an attacker uses the same cookie during two Web-browsing sessions—one via Tor and the other not—the Web server can match the two sessions and determine the attacker's Internet address. Administrators can also use advanced browser plug-ins and scripts to access private information on an attacker's computer and send it to the Web server. However, attackers can disable this troublesome software and thus protect their anonymity.

Network monitoring can reduce—though not eliminate—the risk associated with anonymous Internet access. Attackers are more likely to choose an open wireless network owned by a private party or small business rather than attempting to exploit an open university network monitored by a large IT department. In any case, the main threat isn't an attacker using spoofed MAC addresses and Tor to mount an anonymous attack. What an attacker really wants is a legitimate user's identity. So, the main vulnerability is in using weak (password-based) authentication to control access to sensitive services. As we described earlier, it's possible for attackers to steal numerous online identities. Once attackers assume a privileged identity, such as an IT department employee, they can hide their activities with relative ease.

Press coverage

If an incident occurs, a university should cooperate with the press to reduce the negative impact. The IT department should explain what happened and outline the steps it will take to better secure the system. A university that actively values and builds trust in this way is far less vulnerable to loss of reputation than one that tries to hide the bad news.[9]

Ultimately, attempting to maintain a posture of system infallibility shows a limited understanding of security. It's far better to realize that intruders can enter business systems and earmark resources for procedures that deal swiftly with attacks when they inevitably occur. Currently, however, the marketplace doesn't consider disclosure of breaches a good security policy. For commercial vendors, disclosure entails a short-term financial penalty; on average, attacked companies lose 2.1 percent of their market value within two days of an announcement.[10]

Whatever the general perception, disclosing security problems doesn't translate to poor security. Universities can be key players in correcting this misperception by disclosing security breaches and valuing open security processes.

Privacy concerns

The meaning of "privacy" depends on the context. We define privacy as an individual's right to decide when and how sensitive personal information should be revealed. A university's information system contains considerable personal information, including medical information, social security numbers, annual salaries, student grades, and information about disciplinary actions. Traditionally, most students and employees haven't worried about privacy issues, despite having little or no knowledge about how their personal information is collected, processed, and stored. Recent press reports describing large information thefts have, however, raised questions about the privacy level afforded university students and employees. Clearly, an open network shouldn't make it easier for attackers to steal personal information.

Introducing campus-wide wireless networks that use individual authentication makes it easier to gather information about user movements and activities, but this raises new privacy concerns, particularly when administrators combine wireless network information with information collected from the wired infrastructure. Because they know users' identities, administrators can combine data from many network sessions and build accurate user profiles that include a user's preferred whereabouts. The risk of such tracking is alleviated when universities eliminate authentication on the wireless network's perimeter.

IT departments naturally tend to prioritize network monitoring over individual privacy because they're charged with stopping system misuse. As a result, they might choose monitoring techniques that reduce user privacy to unacceptable levels. To curb this tendency, the department should openly disclose its network monitoring techniques—such openness builds further trust, which again reduces user privacy concerns.

Another risk here is that attackers will invade mobile users' privacy by sniffing the wireless network traffic or accessing information stored on a user's device. Again, all mobile devices must incorporate basic

security mechanisms to protect local data and to encrypt transmission of sensitive information. When users authenticate to the network infrastructure, they're often forced to use encryption. With open university networks, this decision is left to users. Hence, to mitigate this risk, universities must educate users, informing them that they're responsible for their own security and offering advice on necessary security precautions.

Legal issues

In March 2006, the EU parliament adopted Directive 2006/24/EC to track EU citizens' Internet communications. With regard to Internet access, email, and telephony, EU member states must now retain data to trace and identify a communication's source and destination, as well as to identify

- a communication's date, time, and duration;
- the type of communication;
- the communication device; and
- the location of mobile terminals.

The directive offers a detailed specification of data to be retained—for example, collected data must be stored for at least six months, but not more than two years, from the date of the communication. Furthermore, the directive also states that no data revealing the communication's content may be retained.

Currently, there's much uncertainty as to the directive's full impact. Privacy advocates claim that the directive makes it illegal for any entity, including a university, to give users access to the Internet without satisfying the requirements. Although each member state can postpone application of the directive until 15 March 2009, after that date, EU universities might be unable to operate campus networks without mandatory user authentication.

In the US, the Federal Communications Commission wants universities to obey the Communications Assistance for Law Enforcement Act, which was originally written to force telephone companies to open up their digital lines to law enforcement agencies. While universities with networks that exclude the general public are likely to be exempt from CALEA, the situation remains unclear for universities with networks that give Internet access to the public at large. If CALEA applies to a particular campus network, then the university will have to introduce technology to facilitate wiretaps on the network. The university might also have to introduce mandatory user authentication to further facilitate the wiretapping and avoid any future legal actions from the US government.

Clearly, the future legal status of open wireless networks in the US and EU is uncertain. Universities planning to introduce open wireless networks should therefore consult with lawyers to assess their legal risks.

Negative responsibility

Negative responsibility involves what you didn't do but could have done. If an information asset gets hacked, university IT department managers might worry that they'll be blamed for not requiring all users to authenticate to the network infrastructure. As a result, proposals for new open wireless networks must be backed by solid analysis of the system's security properties. If previous security analyses aren't well documented, this can be a tremendous task. In such cases, it might be more convenient to maintain the "default belief" that threats from attackers and other system abusers make introducing an open network too risky. However, as our analysis here shows, universities are ill served by buying into the default belief without carrying out their own security analysis.

Network openness

When designing and implementing an open wireless network, university IT departments must carefully consider several factors, including captive portals, port filtering, and the level of network accessibility.

Captive portals

The captive portal technique forces Web browsers to display a special Web page when users request Internet access. Hence, prior to granting access, universities can ensure that users view the open network's usage rules—and confirm that they accept them—using a "catch and release" captive portal.

The selected portal should meet several requirements, including that it works on different platforms and gives users uninterrupted network access over long periods. Portals requiring pop-up windows in the browser should be avoided; many users find pop-ups annoying, and many terminals don't support them. Finally, the portal should be accessible through smart-phone browsers, as users are increasingly accessing wireless networks from such phones with wireless communication capabilities.

The drawback of captive portals is that not all devices can run a Web browser. The Vocera Communication Badge, for example, is a screenless device that enables instant two-way voice communication over a wireless network. For some universities, the openness reduction entailed by captive portals might be unacceptable.

Port filtering

Typically, letting people in the general population use a university's open network isn't a problem because their numbers are usually small compared to the number of university users. However, if such individuals started downloading mass quantities of data, they could compromise network performance for students and faculty.

File-sharing applications can undoubtedly cause network capacity problems on a wireless network.

Closing the network ports that these applications typically use can help address the situation. If the wireless network has few open ports, it'll be less attractive for people trying to download large data sets. (It's impossible to stop all downloads; users can always turn to options such as port 80.) The University of Bergen's Department of Informatics allows traffic only through SSH (port 22), HTTP (port 80), and HTTPS (port 443). Without port filtering, networks obviously obtain a much higher degree of openness.

Accessibility

An open network's service area is the total indoor and outdoor area from which a mobile terminal can communicate with at least one network access point. The service area should cover all buildings in which university employees and students work, including faculty and administration offices, classrooms, dorms, and libraries.

If the outdoor service areas are too extensive, the open network is likely to interfere with other wireless networks run by private individuals and businesses. Hence, the IT department must work with neighboring wireless network owners to avoid interference problems. The resulting service area's size—in particular, the portion available to the general public—influences the degree of openness. If the area is too small, then not enough people will have access to the network, rendering its openness inadequate.

An open wireless network can improve students' education and make important information more easily available to both faculty and the general public. It can also help universities better focus their security initiatives where they can do the most good. Still, deploying open networks creates an ethical dilemma because they can give attackers anonymous Internet access. In our view, the legitimate privacy requirements of guests, students, and university employees are a powerful argument in favor of open networks. Monitoring network traffic is acceptable, assuming the IT department informs users about the activity.

Because our practical experience is with wireless networks in a university setting, we've focused on them here. Nevertheless, much of our analysis is highly relevant to open networks in general, and more work is needed to determine the best way to implement them. □

Acknowledgments

We thank the reviewers for their suggestions, which significantly improved our article.

References

1. R.B. Kvavik and J.B. Caruso, *ECAR Study of Students and Information Technology, 2005: Convenience, Connection, Control, and Learning*, Educause Center for Applied Research, vol. 6, 2005; www.educause.edu/ers0506.
2. L.F. Cranor and S. Garfinkel, eds., *Security and Usability*, O'Reilly, 2005.
3. R.E. Smith, *Authentication*, Addison-Wesley, 2001.
4. S.T. Kent and L.I. Millett, eds., *Who Goes There?*, US Nat'l Academies Press, 2003.
5. P. Kuper, "A Warning to Industry—Fix It or Lose It," *IEEE Security & Privacy*, vol. 4, no. 2, 2006, pp. 56–60.
6. R. Bejtlich, *The Tao of Network Security Monitoring*, Addison-Wesley, 2005.
7. A. Jones and D. Ashenden, *Risk Management for Computer Security*, Elsevier, 2005.
8. K.J. Hole, E. Dyrnes, and P. Thorsheim, "Securing Wi-Fi Networks," *Computer*, vol. 38, no. 7, 2005, pp. 28–34.
9. S. Bibb and J. Kourdi, *Trust Matters*, Palgrave Macmillan, 2004.
10. H. Cavusoglu, B. Mishra, and S. Raghunathan, "The Effect of Internet Security Breach Announcements on Market Value: Capital Market Reactions for Breached Firms and Internet Security Developers," *Int'l J. Electronic Commerce*, vol. 9, no. 1, 2004, pp. 69–104.

Kjell J. Hole *is a professor in the Department of Informatics, University of Bergen, Norway. His research interests include risk management and application security. Hole has a PhD in computer science from the University of Bergen. He is a member of the IEEE and the IEEE Computer Society. Contact him at kjell.hole@ii.uib.no.*

Lars-Helge Netland *is a PhD student in the Department of Informatics, University of Bergen, Norway. His research interests include trust management and risk analysis of software systems. Netland has an MS in computer science from the University of Bergen. Contact him at larshn@ii.uib.no.*

Yngve Espelid *is a software developer with Bouvet ASA. He has a PhD in computer science from the University of Bergen. Contact him at yngve.espelid@bouvet.no.*

André N. Klingsheim *is a PhD student in the Department of Informatics, University of Bergen, Norway. His research interests include network and application security. Klingsheim received an MS in computer science from the University of Bergen. He is a member of the IEEE Computer Society. Contact him at klings@ii.uib.no.*

Hallvar Helleseth *is a software developer with Avenir AS. He has an MS in computer science from the University of Bergen. Contact him at hallvar@ii.uib.no.*

Jan B. Henriksen *is a senior engineer in the Department of Informatics, University of Bergen, Norway. Henriksen is an automatics engineer with a varied computer science background; his research interests include computer networks and security. Contact him at jan.berger.henriksen@ii.uib.no.*