AUSTRIA INSTITUT FÜR
EUROPA- UND SICHERHEITSPOLITIK

AIES-STUDIEN

David B. Skillicorn / Christian Leuprecht

Beyond the Castle Model of Cyber-Risk and Cyber-Security

Nr. 4
Juli 2015
Maria Enzersdorf

Prof. David Skillicorn is Adjunct Professor in the Mathematics and Computer Science Department of the Royal Military College in Kingston, Ontario.

Dr. Christian Leuprecht is Associate Fellow at the AIES. He is associate professor of political science at the Royal Military College of Canada and at Queen's University.

AIES
Austria Institut für
Europa und Sicherheitspolitik

Schlossgasse 6
2344 Maria Enzersdorf
Österreich
Telefon +43 (0)2236 411 96
Fax +43 (0)2236 411 96-9
www.aies.at
office@aies.at

Contents

Executive Summary

The predominant metaphor for secure computing today is defence in depth: higher, better layers of walls. This article explains why that approach is as outmoded for cybersecurity today as it became for physical security centuries ago. Three forces are undermining the castle model as a practical security solution. First, organizations themselves tear down their walls and make their gateways more porous because it pays off in terms of better agility and responsiveness – they can do more, faster and better. Second, technological developments increasingly destroy walls from the outside as computation becomes cheaper for attackers, and the implementation of virtual walls and gateways becomes more complex, and so contains more vulnerabilities to be exploited by the clever and unscrupulous. Third, changes in the way humans and technology interact, exemplified (but not limited to) the Millennial generation, blur and dissolve the concepts of inside and outside, so that distinctions become invisible, or even unwanted, and boundaries become annoyances to be circumvented. A new approach to cybersecurity is needed: Organizations and individuals need to get used to operating in compromised environments. The article's conclusion operationalize this strategy in terms of a paradigm shift away from a Castle Model and towards a more nuanced form of computation and data assurance.

Introduction

Walls have a dubious history as tools of defence. The Roman Empire disintegrated despite Hadrian's Wall. The Great Wall of China became irrelevant once China's elite, confronting a peasant rebellion, invited in those same Mongols the Wall had been meant to keep out. Its modern incarnation, the Great Firewall, has Chinese spoofing IP addresses to circumvent it. The Maginot line failed to keep the Wehrmacht out of France. The Berlin Wall could not isolate East Germans from the lure of a better life, and was eventually dismantled. The border between the United States and Mexico remains porous, various barriers notwithstanding. Fencing has not prevented migrants from swarming Ceuta and Melilla. The "Castle Model" of cybersecurity is as alluring as these physical defences but, as we shall show, creates an equally false sense of security.

This article problematizes the relationship between risk and security in cyberspace. Hitherto cybersecurity has been taken as a "given," a banal fact of the digital world in which we live. By contrast, understood as a social process, cybersecurity as a mundane vernacular becomes untenable: operationally, conceptually, and theoretically. That cyberspace is yet another example of the way the success of modernity is changing society is uncontroversial; not so the transformation of order in the form of unseen consequences of cyberspace where our institutional resources are unable to cope. This article posits cyberspace as yet another example where the success of modernization is creating problems it is unprepared to solve with present-day institutions. It shows how uncertainties in cyberspace are manufactured, and the inability to control them, and the feedback loops and consequences they engender.

The imbrication of the digital and materials worlds is well established (Sassen 2002). So is the Internet as a productive system (Foucault 1991, 1998) that transforms ordinary citizens (Bauman and Lyon 2013). It is a socializing agent that creates citizens and consumers through a process of discovery and participation: it allows adolescents to share, discuss, influence and learn interactively from each other and from the medium. In cyberspace, activities such as communication, commerce, and entertainment, are conducted, mediated, and learned socially in a way that differs fundamentally from physical space (Lee and Conroy 2003: 1709). The "digital drift" that follows from the individualization that cyberspace enables "individuals to both 'embed' and 'disembed themselves in a variety of criminal activities and lifestyles off- as well as online. [These changes] impact upon the level of individual commitment to criminal activities and lifestyles as well as upon the degree and forms of interdependence and reciprocity implicated in the accomplishment of crime" (Goldsmith & Brewer 2015: 2). The micro-social dimensions of criminality, especially those mediated by the Internet, thus warrant closer attention (Resnyasnsky et al. 2012). The article's preoccupation with cybersecurity is one corollary of the qualitative change in cybercriminality.

In cyberspace individuals do not merely observe and model routes of data production; they engage in a process of discovery and participation that gives rise to "technologically mediated sociality" (Tufekci 2008: 21; Pariser 2012). Individuals divulge information about themselves in ways and to an extent that is quite unprecedented in physical space, they play with and within the boundaries of the software, react to and resist the impulsions written into the codes, set up revelatory profiles, hold back or give info (Beer 2009: 998; 2014). In fact, cyberspace's "impotentiality" encourages citizens to do anything while diminishing the ability "not to" do anything and everything (Agamben 1991).

Security in the physical world involves social processes. The sociology of surveillance has long shown the same to hold for security in the digital age. Yet, neither surveillance studies nor critical theory has explicitly pondered the social processes that cause an individual to be in/secure in cyberspace, nor the implications that follow. Individuals enter, explore, exploit, and exit cybserspace. It is their nascent, emergent, tentative behavior, and the social processes that ensue, that generates cyberrisk in the first place. Luhmann, Giddens, and Habermas are renowned for observing how risk is related to decision-making, but those decisions are also creating largely unintended consequences for others (Leydesdorff 2010). By virtue of its interconnectivity, cyberspace is the prime example where deciders and those who are affected by the consequences have little ability to participate in decision-making. Once user, interface, executables, platforms are situated in the cyber-habitus, it becomes clear that the current paradigm of cybersecurity and its provision is fundamentally flawed.

Organizations with security concerns normally frame the issue as a dichotomy: "inside" versus "outside". What happens inside the organization is permissible; what happens outside is considered to be, at least potentially, harmful or dangerous. This framing applies to countries and their governments (see, for instance, a recent review of US cybersecurity policy by Harknett and Stever 2011), to government departments, including the military and security components, to businesses, to other kinds of organizations, and even to households, where land is delineated by property lines, and houses by lockable doors and windows. The difference between "inside" and "outside" is a delineation that defines the two sides. No organization can exist as an island. Boundaries must inevitably have gateways that permit resources and information to flow in and out. This separation into inside and outside can also exist recursively within the organization. For example, departments within an organization can have their own "inside" and regard (at least in some sense) the rest of the organization as "outside". This explains, for example, the persistent difficulty of sharing intelligence among organizations within the same government.

This metaphor of "inside" and "outside" – euphemistically known as defence in depth – is better called the Castle Model (Frincke and Bishop 2004), since it replicates the mindset of the medieval castle: strong (often layered) walls preserving the integrity of the inside against any form of attack from the outside – and the ability to impose strict controls over movement in and out (but with a curious blind spot to movements within). As in physical castles, walls in cyberspace are costly to build and impede the movement of digital goods, services, and information between the inside and the outside. When these "castles" fail to nest properly, difficult issues present themselves that hint at the fraying of this view of the world. Businesses were once contained inside national borders; the rise of multinational corporations, with their own boundaries that intersect national borders, creates difficult issues that reveal themselves in, for example, the problems that national governments have collecting taxes they are owed.

Quigley and Roy link the approach to cyber security and typical reactions to a security breach to cultural theory. They argue that governments favour a "hierarchist" approach to security. This method emphasizes the importance of structure, rules, and fairness. Any departure from the hierarchy and rules signifies a risk that may not be overcome if members of the hierarchy have inadequate training or skills (Quigley and Roy, 2012). While the hierarchic structure can be beneficial for organization, it leaves little margin for error. Additionally, the government's hierarchist approach focuses on control; this approach can be intimidating for private sector partners. Were a cyber security crisis to occur, governments require flexibility and partners.

Although all boundaries differentiate inside and outside, they can make this differentiation in multiple ways. Organizations have boundaries in at least three important domains:

The first is physical – there are physical or geographical spaces that are defined to be inside the organization. When the organization is a country, this is its territory; when it is a business, this is its workplace (factories, offices, warehouses, and retail space). Boundaries that separate inside and outside in this domain are usually obvious: walls and fences; and gateways and doors to pass through them.

The second domain is temporal – there are times that, at least for businesses, are defined to be inside. We call them the working day. Boundaries in this domain are less obvious, but they are there nevertheless. In some businesses, employees must clock on and off; in others the maintenance of these boundaries is a management task, and employees are expected to seek permission when they will not be "inside" during the normal, expected times.

The third domain is the online world – there are computational and network resources that are considered as inside the organization; and a much larger set that is considered outside. The boundaries in this case are a set of electronic and computational wall technologies that are designed to stop data from moving in and out, except as allowed. The gateways now become more distributed and harder to see, which raises new issues.

Some of these wall technologies are:

- Antivirus software that examines incoming email and web traffic for the signatures of known attacks.
- Firewalls that embody rules about what other kinds of traffic is allowed in and out of the organizational network and individual systems.
- Anti-spam software that examines incoming email for messages that are not real communications.
- Authentication mechanisms such as passwords that allow only approved users to access the network and systems.
- Exfiltration detectors that examine outgoing data and block any (usually documents) that are intended to remain inside the network.

Authentication mechanisms sufficed for standalone systems. These other virtual wall technologies are the response to systems that are connected to the Internet; consequently, their internal content is potentially accessible to anyone on the planet. Even organizations that are not connected to the Internet, for example militaries and security and intelligence organizations that run their own air-gapped “closed” networks, have been forced to admit that they cannot really consider themselves as separate from the larger world. For example, ubiquitous cameras on laptops mean that data can be passed by pointing the camera of a computer on an outside network at the screen of a computer on an inside network; ubiquitous microphones mean that a computer on an outside network can listen to sounds made by a computer on an inside network (even at frequencies inaudible to humans).

Boundaries, and so the preservation of the concepts of inside and outside, have been dissolving under three main forces:

- Strong incentives to reduce boundaries because of the opportunities this creates for agile response to the environment and streamlined access from the outside; and the cost of constructing, operating, and maintaining boundaries;
- Technological changes to the way organizations structure their computational resources that make boundaries increasingly porous; and
- Changing human culture, captured most strongly in the so-called Millennial generation, for which boundaries are becoming irrelevant.

Organizations are tearing down walls from the inside

The first driver of change is the opportunities that having weaker boundaries create in a connected world, and the costs of putting boundaries in place and operating them.

Removing or weakening boundaries allows more flexible travel and use of human capital in the physical world, and new levels of sophistication in acquisition of information and coordination in the online world. For example, the Schengen area in Europe allows unfettered movement across national borders, making it easier for business interaction and tourism. More flexible working hours encourage greater workforce participation. Allowing employees to access email at home has ushered in a new level of business responsiveness. Making it possible for citizens to access government services from their homes, rather than having to visit a government office, has streamlined service delivery. Reducing or weakening boundaries has considerable upsides: flexibility, greater workforce participation, and responsiveness.

Also, creating and enforcing strong boundaries imposes considerable costs and delays. These boundaries have to be built and operated, a cost that is approximately proportional to how robust and secure they are. They also impose delays and costs whenever something has to pass across them. National borders create the need for visa and passport mechanisms, lines at borders to verify who may cross, and civil servants to administer the process. Security for buildings requires locks and keys, CCTV, and security guards to control entry and egress. Fixed working hours require time clocks (and those who check them) or management’s attention to tardiness.

Erasing such boundaries reduces the marginal costs they impose. As organizations face a more competitive world, where expectations of productivity, efficiency, performance and responsiveness increase, and where overheads continue to be squeezed, it is unsurprising that they feel pressure to reduce transaction costs by reducing boundaries (Pew 2010: 23). Many organizations have yet to come to grips with the impact this has on their conception of inside and outside, and the ensuing security implications.

Technological developments are destroying walls from the outside

The second driver of change is the increasing difficulty, even impossibility, of providing strong boundaries because of technological change.

In the physical world, the reduced costs of transport (private places, unmanned aerial vehicles, small submersibles) and the increased ease of forging documents (physical or electronic) is making borders more porous. The difficulty that the U.S. has in interdicting drug shipments and illegal immigration, despite having a strong border-security regime and having put considerable resources into it, illustrates this development. In the context of building security, keys are easy to copy, and even "high-end" technologies such as fingerprint readers and iris scanners are relatively easy to spoof.

In the cybersecurity domain, the virtual wall technologies discussed above are all becoming increasingly porous (McDougal 2009). Quigley and Roy found that these porous networks are allowing cybersecurity threats to flourish. Websense Security Labs found that over the span of a half year threatening websites had increased by an overwhelming 233%. Additionally, there was a more narrowed focus on data. Of the threats recorded, 37% involved stealing data (Quigley and Roy, 2012). These statistics from 2009 show that there is a growing, threatening presence in the cyber sphere, and frequent news stories indicate how these figures are increasing as technology advances.

When passwords provided access to a single system from a dedicated, connected device, it was easy to protect them. When passwords must necessarily pass over public networks, they cannot be robustly protected, even though they are encrypted. Standard attacks require only that every possible string (shorter than a given length) be encrypted using one of only a few standard algorithms and compared to the encrypted password to discover what the plaintext password is. The computational requirements to do this are, by today's standards, modest and can be rented from grid service providers for a few dollars. The only defence is to make passwords long, so that many potential strings must be encrypted by the attacker – but even a 15-character password is only a mild impediment, and humans begin to struggle to remember strings as long as this. New methods of authentication have proven difficult to build and operate reliably: multifactor authentication can be awkward to use, and biometric authenticators easy to spoof.

Virtual wall technologies also have two major weaknesses: (i) it can be hard to identify where the walls actually are; and (ii) the hardware and software that implements the virtual wall is almost invariably not built by the organization using it; rather, it is bought off the shelf. Paradoxically, then, technology meant to protect actually introduces new vulnerabilities.

The first weakness means that it is hard to know where a virtual wall is needed, and makes it easy to miss places where a wall might be necessary. For example, virtual private networks allow employees to use the organizational network from home as if they were physically connected to it. However, the connection between the home computer and the organizational network is now a vulnerability, even if it is encrypted (as the recent Heartbleed vulnerability dramatically showed) (CVE 2014); and the home computer has become, in practice, a part of the organizational network, together with any virus and malware infections it may have previously acquired. Exfiltration detectors can be defeated by first moving a document to a home computer and disseminating it from there. Other tools such as Microsoft's Remote Desktop allow similar functionality with even less protection. Despite the vulnerabilities created by the ability to connect remotely, many organizations feel compelled to allow telecommuting, and want their employees to be available 24x7 because it allows organizational responsiveness that increases the bottom line. As far as we are aware, there are no standard products that allow a remote computer to be incorporated into an organizational network in this way while preserving the full security that a computer physically located on the network would have.

The recent trend towards using clouds for storage and computation introduce similar vulnerabilities. If organizational data is stored in a cloud, that data is no longer clearly inside the organization. The process of transferring it from organizational systems to the cloud creates a potentially vulnerable channel; the data held by the cloud may be vulnerable to access by others, even if it is encrypted; and the data becomes a kind of hostage to the hosting organization. For example, a failure of their systems or an injunction served on them for an unrelated matter may prevent continuing access to the data in a timely fashion.

A somewhat similar vulnerability comes from allowing other organizations to access a given organization's network. There are strong incentives for this: business-to-business connectivity allows collaborative work to happen smoothly; just-in-time component delivery requires a supplier to be aware of not only how much of a component is held by the consumer, in real time, but also the rate at which it is being consumed, so that the optimal time for the next delivery can be planned sufficiently far in advance. A major data breach of the retail chain Target occurred at the end of 2013; the attack came via access granted to an HVAC supplier. Organizations that implement strong security themselves can be vulnerable because of the weaker security of these partner organizations that they regard as separate (outside), but are actually salients of the more secure organization.

Another category of vulnerability comes from the evolution of web browsers as tools, not just for the consumption of static information, but as portals for two-way information flow, and often control of other systems. The problem here is that all traffic involving a web browser travels over the same port: port 80. As a result, all sorts of different traffic, innocuous and potentially dangerous, flows over a single channel. Blocking it would cut off even the simplest web browsing, so it is almost invariably left unblocked. It is extremely difficult to parse the traffic stream that passes through this channel to block some kinds of traffic while allowing others. This is an extremely difficult task, and the bar is constantly raised as more and more services are piggy-backed on the ubiquitous browser-server mechanism.
Nor are the cyber and physical worlds decoupled. In 2008, U.S. military networks were successfully infiltrated by a worm on a USB device which had been dropped in a military parking lot in the Middle East; the U.S. Department of Homeland Security carried out an experiment where they dropped USB devices in various parking lots in the U.S. and found that more than 60% of them were picked up and plugged into computers (Bloomberg 2011). The walls of an organization's network may be as simple as the USB connectors on its systems.

Another vulnerability of the castle model of security is that it distracts attention from what is happening inside the walls. A major weakness is the ability of insiders to carry out attacks from within the organization. This is a particular problem, even for organizations with high levels of security, because of the prevalence of contractors who are treated as insiders, but may not have the organization's interests at heart. They do not usually have the same degree of loyalty because they are not subject to the same amount of hostage capital as permanent employees and may, therefore, have an incentive to prize individual gain in the short-term over long-term payoff for the organization as a whole. They may also not have been vetted to the same level as mainstream employees. Edward Snowden stands out as the prime example, partly the National Security Administration (NSA) for which he was a contractor is among the most secure in the world. Bradley Manning, a former uniformed member of the US Department of Defence (DOD) is another high-profile example. Why the two most prominent examples are both American, and why they both came from the US national security apparatus, is an interesting puzzle of its own.

The technologies that implement the virtual wall technologies are themselves a source of vulnerability. Very few organizations implement these technologies themselves. Indeed, to do so would require building their own hardware, and then a considerable amount of software. Instead, most organizations use off-the-shelf hardware and software systems that they configure by defining sets of rules of what is allowed and forbidden. Even if these rules are correct, the organization cannot know if there is a vulnerability embedded in the system that applies the rules. Worse still, some of these technologies have a mechanism that exists to allow them to be updated remotely. The organization using them may not even be aware that this mechanism exists and, because it is built into the wall, other wall technologies may not notice it.

Changes in human interaction are blurring the distinction between inside and outside

The third driver of change is the new attitudes to connectedness that have developed in a population that has discovered the Internet and cheap network access, and even more strongly in the generational cohort that has grown up with it. The so-called Millennials, the generation born between, roughly, 1984 and 2004, are now beginning to become the majority of the workforce, as Baby Boomers retire. These "digital natives" did not discover and learn network technology: it has been a ubiquitous background to their lives while growing up. Their attitudes to technology, work and organizations are having an impact on how organizations conceive themselves that is at least as significant as the effects of technology per se. Many of their attitudes are also held by earlier generational cohorts, but with lower intensity. Previous generations also use technology less fluently, and, therefore, with more variability.

Some of the characteristics associated with Millennials are:

- ★ They expect technical innovation as a matter of course; they have seen it happening steadily throughout their lives, they expect it to continue, and a significant portion feel a need to be on the forefront of technical change. Whereas previous generational cohorts included "early adopters", Millennials are early adopters (Deloitte 2012).

- ★ They depend on technology. Millennials are used to a world in which a personal communication device is always within reach, even when sleeping. This device can connect them to other individuals in their personal peer groups, and to the informational content of the entire Internet instantaneously and in an almost unlimited way. They expect connectivity everywhere, on public transport, on aircraft, in tunnels, and in meetings. Their sense of physical space is weakened by virtue of the fact that they carry a substantial part of their environment with them (Hershatter & Epstein 2010).

- ★ They interface differently to the world than previous cohorts, both in terms of perception and interaction. Their approach to knowledge tends towards just-in-time information gathering, rather than just-in-case learning. This poses challenges for the educational establishment; it also means that Millennials tend not to plan, even for events as simple as getting together with friends, converging on time and place in real time. Similarly, their relationships are simultaneously tighter and looser than previous generational cohorts (Pew 2010: 9): tighter because it is easy to remain connected, in a superficial way, to many people (it is hard to imagine Millennials coming to a high school reunion to find out how their classmates have turned out – they will already know, at least to some extent); but looser because, even when they are physically together, some part of their attention tends to be in cyberspace.

- ★ Their attention is not deployed in large blocks (in the way that previous generational cohorts at least claimed to do) but rather interleaved in smaller time slices. They are often accused of multitasking everything; there is some truth to this but probably not as much as previous generational cohorts believe.

- ★ They have been exposed to a much greater diversity of people and opinions. Their information sources are not just regional, not just national, but international by default. They can easily find text and video of people speaking other languages. They can encounter a wider range of opinions and contexts than any human could half a century ago.

- ★ They have developed new ways of interacting, effectively a new etiquette for communication, so that the possibility of constant communication with a very large circle of acquaintances does not become intrusive. For example, because personal communication devices are always close, it is considered rude to send text messages at a time when the recipient is probably asleep. So contrary to stereotypes, Millennials have, and are, developing ways of managing an always-on world.

These characteristics of Millennials have implications for their behaviors in an organizational context, implications to which organizations will necessarily have to respond. Many of these implications are positive and provide a springboard for organizations to become more effective. Others are negative and require organizations to find new ways of dealing with them.

Some of the positive implications of the Millennial worldview are:

- ★ They have discovered new ways of cooperating and creating that can be leveraged within organizations to build more holistic, dynamic and so responsive ways of working (Verdon 2012). As a concrete example, businesses whose products are digital, such as software or video, can use three shifts to get these products built more quickly – but these shifts take place in three different physical locations, each spaced eight time zones apart. Building products collaboratively this way requires detailed and regular interaction with members of other cultures, which Millennials are well-equipped to do (Myers and Sadaghiani 2010).

★ They are members of a much wider number of interlocking communities than previous generational cohorts (Statistics Canada 2007). As a result, they provide organizations with a greater, and more diverse, reach. Their membership in these communities is longer-lasting, effectively, for example, discouraging organizations from short-term drive-by marketing and encouraging long-term permission-driven marketing. It also provides them with a competitive edge, for example in job hunting (Pew 2010: 9).

★ They are sophisticated consumers of diverse sources of information. As a result, they are used to cross-referencing and triangulating information they are given, including that from within their organizations. Management strategies that involve holding back information will not be well received by Millennials, who expect to be told what is going on (Myers & Sadaghiani 2010).

★ They believe that technology increases productivity and efficiency (Pempek et al. 2012).

★ Despite the stereotypes of Millennials constantly checking their phones, they use time productively. In particular, they devote time to community activity in a way that previous generational cohorts do not. This is partly because the barriers to doing so have been lowered by technology; and partly because it can be done in smaller chunks (Shirky 2008). Computational tools remember context, reducing the effort of returning to a task in progress, and so enabling productive work to be done in smaller increments.

Organizations can, therefore, expect Millennials to be at least as productive as previous generational cohorts, but in novel ways that may require some adjustments. Their view of community is richer than that of previous generational cohorts, creating new opportunities for many kinds of organizations.

However, there are some negative implications of the Millennial worldview, and many of these are relevant to security. Some of these implications are:

★ They prefer broadcast channels (many-to-many) rather than the one-to-one or one-to-many channels provided by email (Fritzon et al. 2008). In a fundamental way, communication is conceived as a multilogue, a conversation, rather than as a dialogue. They have been called "ambient broadcasters" (Pew 2010: 17). Furthermore, the audience component of a communication is often not a coherent shared-interest group but something more ad hoc ("friends") (Jacobs and Diefenbach 2012). This creates a plethora of problems:

  There is a weaker match between content and receivers. A mailing list has some internal coherence that a group of friends or followers may not. Communications can be easily misconstrued, as a recipient does not necessarily have enough of the context to understand their full meaning.

  Dissemination is not controllable by the original sender, and the technology makes it easy to pass communications on, far beyond their intended reach.

  There are no gatekeepers to control what does and does not get disseminated – individuals decide for themselves (Johnson and Kaye 2010: 326).

  For organizations, this has the potential for public relations and security disasters.

★ Millennials, because they act in an interleaved fashion, do not have a strong sense of role, time, and place. Whereas previous generational cohorts might consider whether or not LOLcat emails were appropriate for organizational email, Millennials are less likely even to conceptualize that their work and leisure roles might require different decisions. In their lives, cyber and physical space blend in a way that is not the case for preceding generations (Harris 2014). From a security point of view, this means less sensitivity about whether, say, a potential organizational decision should be mentioned outside the organization.

★ Similarly, they are likely to distinguish less between being "at work" or not, being used to dealing with work issues outside of normal working hours. The idea of not making personal calls during business hours is totally foreign to them. They are willing to deal with non-work issues during working hours. In front-facing consumer-service businesses this already creates management issues. For the same reasons, they have less sense of being physically at work, and so might perhaps work on confidential organizational business at a local coffee shop, unaware of any security concerns. (Of course, this also means that they are likely to "work" even during leisure time, which can be to an organization's advantage.)

- ★ Their sense of privacy is different to that of most adults from the second half of the 20th Century (Pew 2010: 8). At the mundane level, they are accustomed to living their lives under the pervasive gaze of cultures of social surveillance (Bauman et al. 2014: 141-142) as exemplified by social media sites that disseminate their personal information widely within the social media framework, and also leverage it by selling it to other organizations. It is not yet clear whether Millennials do not realize that their personal information is not only widely spread but also archived for the foreseeable future, or whether they do not care, feeling that living life in the open is natural and appropriate (Accenture 2008; Fritzon et al. 2008). Millennials thus find the imposition of privacy and security irksome at best, and something to resist at worst.

- ★ Because of their use of personal devices and software that knows their location, organizations must take into account that their employees' locations are essentially public information. This is of particular concern, of course, for organizations such as police and armed forces (Hibbard 2011; Drapeau and Wells 2009).

- ★ Millennials will provide their own technology when employers are unable or unwilling to oblige (Accenture 2008). Where previous generational cohorts expected their employers to provide the necessary tools for work, Millennials are predisposed to short-circuit this process. For example, organizations that provide employees with smart phones to ensure their availability may replace these devices on a two-year cycle; much longer than the 6-month or shorter cycles that smart phone makers use. As a result, Millennials may just buy leading-edge smart phones in place of those provided. There are security implications when these (unauthorized and perhaps unrealized) devices are used for organizational activities.

- ★ Similarly, if the software tools provided by an organization are deemed inadequate by Millennials, they are perfectly comfortable acquiring others, perhaps open-source freeware and even installing them on the organization's systems. Again there are security implications.

Millennials and their attitudes, many of which are present in older cohorts albeit at lower intensities, represent challenges for organizations. Many of their characteristics are positive and represent considerable potential for new organizational paradigms. However, from a security perspective, these characteristics create potential vulnerabilities that organizations have perhaps not yet fully realized, and for which good responses are still unclear.

Conclusion: Data assurance in compromised environments

The Castle Model for organizational, and especially network, security is based on layers of walls that define, very strictly, what is inside and what is outside. This model, at least in the cyber domain, has never been very effective. We have suggested that three forces are eating away at this model as a practical security solution. First, organizations themselves tear down their walls and make their gateways more porous because it pays off in terms of better agility and responsiveness – they can do more, faster and better. Second, technological developments increasingly destroy walls from the outside as computation becomes cheaper, and as the implementation of virtual walls and gateways becomes more complex, and, therefore, contains more vulnerabilities to be exploited by the clever and unscrupulous. Third, changes in the way humans and technology interact, exemplified by the Millennial generation, blur and dissolve the concepts of inside and outside, so that the distinction becomes invisible, or even unwanted, and boundaries become either anachronisms or annoyances to be circumvented.

Moreover, the Castle Approach to cybersecurity is marred by a fundamental ethical problem: access to the model is a function of finances, as the degree of protection afforded correlates loosely with sunk costs invested. The rise of the cybersecurity industry is evidence to that effect (Zedner 2009; Gill 2006). The Castle Model thus directly reinforces the digital divide, and indirectly the digital divide's economic and social fault lines among individuals, households, businesses, geographic areas, class, race, ethnicity, and gender across the globe (Castells 2001; Norris 2001; Lu 2001; National Telecommunications and Information Administration 1995). Ethically, any model of cybersecurity that reinforces privilege and, arguably, power relations, is necessarily problematic: as with physical security, cybersecurity should not be parceled out by financial means.

What can be done in a world where the separation between inside and outside is so porous as to prevent hardly anything? Organizations still need to get work done without it being visible to the rest of the world, including their competitors and others whose interests are in opposition. It seems clear that the solution is not to "fix" the three forces that have driven us to the current situation. Organizations may not have consciously decided to weaken boundaries to achieve greater agility, but it has been successful nevertheless. While technology may provide some limited improvements in virtual wall techniques, it is clear that, as ever, the advantage is with attackers. And it is hopeless to imagine that Millennials, and their successor cohorts can be convinced to cut themselves off from the networked world just because they are "at work".

A new kind of solution is needed (Karas et el. 2008). Although still in its infancy, the most hopeful direction to protect data is a strategy known as *computing in compromised environments*. Its goal is to allow organizations (and individuals) to do useful and confidential things in cyberspace, even in the face of the issues we have been discussing. Techniques for computing in compromised environments must allow useful work to be done even if an attacker is already inside the castle. There may still be a role for walls, but only as impediments, and not as protection.

Ways to do this are the subject of active research, so it is only possible to give some flavor of the ideas under consideration, which include:

★ Operating in virtual castles. Virtual machines run on top of physical computers and can emulate the software that would normally run directly on top of the hardware. However, a virtual machine can be created as needed, and destroyed when its usefulness is over. Furthermore, each virtual machine can be configured randomly to be slightly different. This makes it difficult for an attacker to target the task the virtual machine is carrying out because a generic attack can no longer be used – they must first work out which variant is actually in use, and then develop and launch a customized attack. The time window in which this sequence must be carried out has to be smaller than the existence time of the virtual machine.

★ Operating with virtual software. Much popular software has known vulnerabilities that are compensated for by malware detectors and regular software updates. However, so-called zero day exploits – vulnerabilities that are not known to the software creators – remain a problem. It is now possible to create a piece of software to carry out some task using pieces of code found in other places in the system (a kind of software Frankenstein's monster). The advantage of such created-on-the-fly software is that it will be different each time; so, knowing a vulnerability in the official, static version of the software does not mean that any particular occurrence of the actual software will contain it.

★ Modelling at the level of behavior or intent rather than at the level of moving bits. Wall technologies tend to focus on what is crossing the boundary and passing through the gates. Once an attacker in "inside" there is often much less scrutiny. Behavior modelling tries to understand the *intent* of traffic and actions, so that activities whose individual pieces look innocuous can be detected at a more abstract level. This is the sort of traffic monitoring to which signals intelligence agencies are heavily committed.

★ Using secret sharing. Secret sharing allows two or more people to hold individual pieces of information that, on their own, are useless but that, when assembled, reveal some secret to one or more of them. As with safeguards against accidental nuclear launches, systems can be created so that any number of participants must share their piece for the entire secret to be revealed.

Secret sharing can provide an alternative to passwords. As a simple (and artificial) example, a system may provide a user who wants to authenticate with a latitude. The user's correct response is a country with an A in its name that lies on that latitude line. If the system generates the latitude value randomly, it takes a very large number of observations of the challenge-response pair even to begin to guess the rule – but all the user needs is a globe.

★ Use multiple versions of all files and use secret sharing to allow users to work with the true ones. Suppose there is a document that the organization does not want exfiltrated. The system creates multiple copies of this document, one the true one, and the others false. The false ones need not look artificial – the Frankenstein mechanism already discussed means that they can be created from pieces of true documents so that there is no easily automated way to tell, from the content, the true from the false.

Of course, users want to edit and read the true documents and ignore the false ones. Secret shar-

> ing can be used to identify which one is the true one. Suppose, for the sake of a simple example, that there is one true version and one false version. The system generates a fixed-length bit string *Y*. Offline the user is given the bit string that results from computing the exclusive-or of a secret bit string, *S*, and *Y*. When the user wants to access a file, the system provides *Y*, the user (offline) computes the exclusive-or of *Y* with the given string (*S xor Y*) which recreates *S*. If the parity (the number of one bits) in *S* is even, the true document is document 1, otherwise the true document is document 2. Knowing *Y* doesn't help someone else, even an insider, to know which version to exfiltrate; even if the user writes down (*S xor Y*) and leaves it visible, this isn't enough to identify the version to exfiltrate.
>
> This simple idea can be generalized to much larger scale and there are many ways to encode the partial secrets. Furthermore, the true version can be swapped around, can appear to be differently named for different users, and the secrets can be altered easily and cheaply.

These ideas are still in the early stages of development. However, they seem to hold more promise than attempting to continue to build higher and thicker walls, and persuade users not to dig through them, open the gateways from the inside, or circumvent them in other ways. The history of real castles is an object lesson of the weakness of the more-and-better-walls strategy; and of the failure to grasp the sociology of cyber security being posited in this article, and its operational, conceptual and theoretical implications.

Works Cited

Aas, K. F. (2013). Globalization & crime. 2nd ed. London: SAGE, chapter 7, pp. 172-192.

Accenture. (2008). Millennials at the Gates: Results from Accenture's High Performance IT Research. New York: Accenture Research USA.

Agamben, G. (1999). Potentialities. Stanford: Stanford University Press.

Bauman, Z,; and Lyon, D. (2013). Liquid Surveillance: A Conversation. Cambridge: Polity Press.
Bauman, Z., Bigo,D., Esteves, P., Guild, E., Jabri, V., Lyon, D., and Walker, R.B.J.. (2014). "After Snowden: Rethinking the Impact of Surveillance." International Political Sociology, 8(2): 121-144.

Beer,D. (2009). "Power through the algorithm? Participatory web cultures and the technological unconscious." New Media & Society,11(6): 985-1002.

Bloomberg Business. June 27, 2011. Human Errors Fuel Hacking as Test Shows Nothing Stops Idiocy. Accessed June 30, 2011. http://www.bloomberg.com/news/articles/2011-06-27/human-errors-fuel-hacking-as-test-shows-nothing-prevents-idiocy.

Castells, M. (2001). The Internet Galaxy: Reflections on the Internet, Business, and Society. Oxford: Oxford University Press.

Common Vulnerabilities and Exposures, MITRE. 2013. Heartbleed. Accessed April 1, 2013. https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2014-0160 .

Cordes, J.J. (2011). An Overview of the Economics of Cybersecurity and Cybersecurity Policy. George Washington University Cyber Security and Policy Research Institute, Report GW-CSPRI-2011-6.

Deloitte. (2012). Tech Trends 2012: Elevate IT for Digital Business; a Federal Perspective. London: Deloitte LLP Services. Accessed April 1, 2015. http://www2.deloitte.com/content/dam/Deloitte/us/Documents/technology/us-cons-tech-trends-2012.pdf

Drapeau, M., Wells, L. II. (2009). Social Software and National Security: and Initial Net Assessment. Center for Technology and National Security Policy. Washington, DC: National Defense University.

Foucault, M. (1991). Discipline and Punish: the birth of a prison. London: Penguin.

Foucault, M. (1998). The History of Sexuality: The Will to Knowledge. London: Penguin.

Frincke, D. A., Bishop, M. (2004). "Guarding the castle keep: teaching with the fortress metaphor." IEEE Security & Privacy, 2(3): 69–72.

Gill, M. (2006). The Handbook of Security. New York: Palgrave Macmillan.

Goldsmith, A., Brewer, R. (2015). Digital drift and the criminal interaction order. Theoretical Criminology. Forthcoming.

Harris, Michael. (2014). The End of Absence: Reclaiming what we've lost in a world of constant connection. Toronto: Current.

Harknett, R. J., Stever, J.A. (2011). "The New Policy World of Cybersecurity." Public Administration Review, 71(3): 455-460.

Hershatter, A., Epstein, M. (2010) Millenials and the World of Work: An Organization and Management Perspective. Journal of Bus Psychology, 25 (2): 211-223.

Hibbard, L. (2011). Communicating with the Net Generation. Carlisle Barracks, PA: U.S. Army War College.

Jacobs, J., Diefenbach,V. (2012). The Use of Social Media in Public Affairs – A German Perspective. Brussels North Atlantic Treaty Organization RTO-MP-HFM-201.

Johnson, T.J., and Kaye, B.K. (2010). "Believing the blogs of war? How blog users compare on credibility and characteristics in 2003 and 2007." Media, War & Conflict, 3(3): 315-333.

Karas, T. H., Moore, J.H., Parrott, L.K. (2008). Metaphors for Cyber Security. SANDIA Report SAND2008-5381. Albuquerque, NM: Sandia National Laboratories.

Lee, C. K. C., Conroy, D.M. (2003). "Teenager's Consumption on the Internet. Australasian Marketing Journal." 13(1): 8-19.

Leydesdorff, L. (2010). "The Communication of Meaning and the Structuration of Exceptions: Giddens' 'structuration theory' and Luhmann's 'self-organization'." Journal of the American Society for Information Science and Technology, 61(10): 2138-2150.

Lu, M. (2001). "Digital Divide in Developing Countries." Journal of Global Information Technology Management, 4(3): 1-4.

McDougal, M. (2009). Castle Warrior: Redefining 21st century Network Defence. CSIIRW '09 Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies. Accessed 25 May 2015. http://www.cisr.ornl.gov/csiirw/09/CSIIRW09-Proceedings/Abstracts/McDougal-abstract.pdf

Myers, K. K., Sadaghiani, K. (2010). "Millennials in the Workplace: A Communication Perspective on Millennials' Organizational Relationships and Performance." Journal of Business and Psychology, 25(2): 225-238.

National Telecommunications and Information Administration. 1995. Falling through the net: A survey of the have nots in rural and urban America. Washington, DC: U.S. Department of Commerce.

Norris, P. (2001). Digital Divide: Civic Engagement, Information Poverty, and the Internet Worldwide. Cambridge: Cambridge University Press.

Pariser, E. (2012). The Filter Bubble: How the New Personalized Web is Changing What We Read and How We Think. New York: Penguin Books.

Pew Research Center. (2010). The Future of the Internet. Available at http://pewinternet.org.

Quigley, K., Roy, J. (2012). "Cyber-Security and Risk Management in an Interoperable World: An Examination of Governmental Action in North America." Social Science Computer Review 30(1): 83-94.

Resnyansky, L., Falzon, L., and Agostino, K. (2012). From transaction to meaning: Internet-mediated communication as an object of modeling. In: 8th International Conference on Social Science Methodology. Sydney, 9-13 July, Conference Proc. Vol II. Accessed 3 May 2015. http://itupl-ura1.ml.unisa.edu.au/R/?func=dbin-jump-full&object_id=116267.

Sassen, S. (2002). Towards a Sociology of Information Technology. Current Sociology, 50(3): 365-388.

Shirky, C. (2008). Here Comes Everybody: The Power of Organizing Without Organizations. New York: Penguin Press.

Tufekci, Z. (2008). "Can you see me now? Audience and disclosure regulation in online social network sites." Bulletin of Science Technology & Society, 28(1): 20-36.

Verdon, J. (2012). The Wealth of People: How Social Media Re-Frames the Future of Knowledge and Work. Brussels: North Atlantic Treaty Organization RTO-MP-HFM-201 (April).

Zedner, L. (2009). Security. Abingdon: Routledge, chapter 5.