Government Information Quarterly xxx (2016) xxx–xxx

Contents lists available at ScienceDirect

Government Information Quarterly

journal homepage: www.elsevier.com/locate/govinf

Beyond the Castle Model of cyber-risk and cyber-security

Christian Leuprecht [a,*], David B. Skillicorn [b], Victoria E. Tait [c]

[a] Department of Political Science, Royal Military College of Canada, P.O. Box 17,000, Station Forces, Kingston, Ontario K7K 7B4, Canada
[b] School of Computing, Queen's University, Kingston, Ontario K7L 3N6, Canada
[c] Department of Political Science, Carleton University, B640 Loeb Building, 1125 Colonel By Drive, Ottawa, Ontario K1S 5B6, Canada

ARTICLE INFO

Article history:
Received 13 July 2015
Received in revised form 27 January 2016
Accepted 29 January 2016
Available online xxxx

Keywords:
Cyberdefense
Security
Boundaries
Organizational boundaries
Millennials
Generational differences
Compromised environments

ABSTRACT

The predominant metaphor for secure computing today is modeled on ever higher, ever better layers of walls. This article explains why that approach is as outmoded for cyber security 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 cyberwalls 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 cyber security is needed: Organizations and individuals need to get used to operating in compromised environments. The article's conclusion hints at more nuanced forms of computation in environments that must be assumed to be potentially compromised.

1. Introduction

The Castle Model is a metaphor for cybersecurity, in which the presence of walls (boundaries), often in layers, create a space that is considered "inside" and, therefore, safe, in contrast to a conceptual "outside" that is considered potentially dangerous. The metaphor draws on conventional castles, with their emphasis on strong walls that are difficult and costly to breach, and gateways that allow traffic out and in, but only in a controlled way that keeps the inside safe.

Walls have a dubious history as tools of defense. From the Stone Age, humans surrounded their settlements by walls, but history is full of examples of 'impregnable' castles being penetrated. 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, great efforts notwithstanding. The Castle Model of cyber security is as alluring as these physical defenses but, as we shall show, creates an equally false sense of security.

Cyber security and cyber risk are conventionally addressed as technical problems with a small cultural component. We argue that solutions to the rapidly growing problems associated with cyber security require a more balanced understanding. The mindset associated with "defense as walls" risks creating a blind spot to some of the most substantial forces preventing progress in cyber security, forces that are not associated with malignity or laziness, but with the need to get useful and productive work done. Our focus is on the international, national, organizational, and personal forces that are responsible for the present parlous state of cyber security.

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 cyberspace. It is their nascent, emergent, tentative behavior, and the social processes that ensue, that generate cyber risk in the first place. Luhmann, Giddens, and Habermas observed how risk is related to decision-making: decisions often create largely unintended consequences for others (Leydesdorff, 2010). By virtue of its interconnectivity, unintended consequences can be multiplied several million-fold, and in extremely short timeframes.

The Castle Model for cyber security 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 cyber security industry is evidence to that

* Corresponding author.
E-mail addresses: christian.leuprecht@rmc.ca (C. Leuprecht), skill@cs.queensu.ca (D.B. Skillicorn), victoria.e.tait@gmail.com (V.E. Tait).

http://dx.doi.org/10.1016/j.giq.2016.01.012
0740-624X/Crown Copyright © 2016 Published by Elsevier Inc. All rights reserved.

effect (Zedner, 2009, chap. 5; Gill, 2006). In many countries, cyberdefense is regarded, at least partially, as a societal good, akin to public health or policing but it is not provisioned in the same way. Instead, the Castle Model directly reinforces the digital divide, and indirectly the digital divide's economic and social fault lines across individuals, households, businesses, geographic areas, class, race, ethnicity, and gender (Castells, 2001; Lu, 2001; National Telecommunications and Information Administration, 1995; Norris, 2001). It also blinds governments and organizations to cooperative opportunities for collective benefit.

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 cyber security policy by Harknett & 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" delineates the two sides. 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.

No organization can exist as an island. Boundaries must inevitably have gateways that permit resources and information to flow in and out. There is a natural and inevitable tension between walls, preventing access, and gateways, allowing it. This tension reflects the balance between security and usability.

This metaphor of "inside" and "outside" is called the Castle Model (Frincke & Bishop, 2004) because it replicates the medieval mindset: strong (often layered) walls preserving the integrity of the inside against attack from the outside – and the ability to impose strict controls over movement in and out (but often 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 issues that reveal themselves in, for example, the problems that national governments have in adapting taxation regimes to the modern world.

Just as physical castles were built to be imposing, as well as defensible, a great deal of cyberdefense infrastructure adds little to real defense but creates the impression that defenses are in place. This has been called "Security Theater" (a term which Bruce Schneier is credited with having coined). A common example is in the domain of password control. Many organizations insist that passwords contain both upper and lowercase characters as well as symbols. Using this larger character set increases the effective resistance of the password to brute-force cracking by the equivalent of approximately three lowercase characters, a tradeoff that any user of a tablet or phone would happily make. Similarly, many organizations require passwords to be changed regularly. Once an account has been infiltrated, a sophisticated intruder will install a keylogger to capture the new password as soon as it is changed. Furthermore, many users simply change their passwords enough times in succession that their organization's policy allows reuse of the original. Both aspects of password controls are, therefore, largely a form of theater. Although Security Theater is ineffective in increasing actual protection, it remains popular as a way of signaling concern to the wider public.

The problem is aggravated by technologies of protection that are expensive to build; consequently, most organizations buy them off the shelf. This results in defensive monocultures where many different organizations use exactly the same walls. Attackers' sunk costs are thus reduced and optimized as they can invest in one attack technology, knowing that it can be leveraged across many targets. A recent example is the Heartbleed vulnerability, an error that made apparently encrypted communication traffic vulnerable to access by an attacker in a straightforward way. The vulnerability affected, by some estimates, two-thirds of web sites and had serious knock-on effects by invalidating security certificates. Its cause was a programming error that had gone unnoticed for two years in open-source software.

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 domain 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; and
- 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 cyberwall technologies are the response to systems that are connected to the Internet, making their internal content 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).

Technology enables three new possibilities that did not exist for real-world castles. Defense in-depth historically meant more concentric layers of defenses, but defense in-depth today reflects the fact that defenses are no longer concentric. The first new possibility is that attack scenarios for complex systems can be computed, taking into account the individual vulnerabilities of walls and gateways, and how they can be chained together to create intrusion pathways. These scenarios are called attack graphs (Sawila & Skillicorn, 2012); they provide a roadmap that could be used by a potential attacker to find the easiest way in. However, they have limitations: they require detailed knowledge of the potential vulnerabilities of each component of the system under attack, detailed knowledge of the possible or plausible connections among these components, and a substantial amount of computing power to perform the required analysis. This makes their use by attackers problematic — attackers do not usually know enough to exploit them. Thus their use is typically restricted to defenders, who use them to assess overall vulnerability and to select cost-effective strategies to harden defenses.

A second novel possibility enabled by technology is that intrusions can, in principle, be detected as they occur by real-time monitoring tools that allow analysts to see where and how a system is under attack. This is typically ineffective for purely pragmatic reasons. Real-world systems are under attack constantly, and it is extremely difficult to differentiate the significant attacks from those that have no chance of succeeding. Imagine a real-world castle able to report every missile that struck its walls to a medieval control center. In any attack, the number of reports would be large, and would not provide much information about the actual ebb and flow of battle, let alone an invading party quietly slipping in through a remote entrance. Furthermore, such systems are vulnerable to attacker-generated false alarms. System monitoring is, therefore, a helpful addition to the defenders' arsenal, but it is often defeated by the rate of false positives. Such was the case with the JP Morgan Chase hack of 83 million accounts in 2014: having allegedly spent $250 million on cyber protection that year alone, the system ended up being overwhelmed by false positives, and, by all accounts, unable to parse high-risk from low-risk intrusions. The consumer-facing version of this problem can be seen in the warnings generated by antivirus, malware detection, malicious website detection, and security-certificate checking tools. Almost all users are unable to judge which, if any, of these warnings should be heeded and often the only alternative to ignoring them is to abandon the planned task.

Physical castles did not have the resources to maintain checks between different regions of the castle in an ongoing way. If an intruder entered stealthily, there was typically nothing to stop them going anywhere inside. In the analogous cyber setting, it is possible to impose further levels of checks when an individual accesses different regions of an online system using tools such as Role Based Access Control (Sandhu, Coyne, Feinstein, & Youman, 1996). In such models, the areas accessible to an individual depend on their role in the organization. However, as the ability of Snowden, Manning, and Delisle, in three major incidents, to exfiltrate and publicize vast amounts of classified data show, role based access control is either not used or ineffective in real-world systems, even security-minded systems where the stakes are high.

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.

2. Organizations are tearing down walls from the inside

The first driver of change is the opportunities that having weaker boundaries creates 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.

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). 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.

3. 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 robust, well-resourced border-security regime, 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 cyber security domain, the cyberwall technologies discussed above are all becoming increasingly porous (McDougal, 2009). Quigley and Roy found that these porous networks are allowing cyber security threats to flourish. Websense Security Labs found that, over the span of half a 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 & Roy, 2012). A Center for Strategic Studies report, in collaboration with McAfee (CSIS and McAfee, 2014) suggested the global losses from cybercrime cost between $US375 billion and $US575 billion in that year. These statistics show that there is a growing, threatening, and expensive attack ecosystem in the cyber sphere, and frequent news stories indicate how these figures are increasing with time.

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 defense 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.

Cyberwall 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 cyberwall 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 cyberwall 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 or cell phone and the organizational network is now a vulnerability, even if it is encrypted (as the Heartbleed vulnerability dramatically showed) (CVE, 2013); and the home computer and cell phone have become, in practice, a part of the organizational network, together with any virus and malware infections they 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. There have been attempts to build and deploy organizational apps or an organizational app store so that the code downloaded onto employee devices, even those not owned by the organization, can be controlled or monitored. Such apps attempt to create a space on each employee's device that behaves as if it were inside the organization. However, the security regime of popular mobile platforms (which can often be 'rooted') is not strong enough to make this a general-purpose solution. Despite the vulnerabilities created by the ability to connect remotely, many organizations feel compelled to allow telecommuting, and want their employees to be available 24 × 7 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 device to be incorporated into an organizational network 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 (Wall Street Journal, 2014). Organizations that implement strong security themselves can still 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, flow 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. The bar is constantly raised as more and more services are piggybacked 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 USB devices 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 of any organization is the ability of insiders to carry out attacks from within. 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 because the National Security Administration (NSA), for which he was a contractor, is among the most secure in the world.

The technologies that implement the cyberwall 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 accessed remotely. For example, many network switches, which see all of the traffic inside an organization and, therefore, represent a high-risk threat and may themselves implement some of the firewall rules, contain hidden user accounts to enable their software to be updated remotely. In some cases, the user names and passwords associated with these accounts are hard-coded, so that even an organization sophisticated enough to see the problem cannot fix it. Less sophisticated organizations may not even be aware that such loopholes are built-in to the devices they buy and deploy and, because they are built into the wall, other wall technologies may not notice them.

4. 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*

Please cite this article as: Leuprecht, C., et al., Beyond the Castle Model of cyber-risk and cyber-security, *Government Information Quarterly* (2016), http://dx.doi.org/10.1016/j.giq.2016.01.012

se. Many of their attitudes are also held by earlier generational cohorts, but with lower intensity. Previous generations 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 to the world differently than previous cohorts do, 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): 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 ("phubbing");
- 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; and
- 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 principles for 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 Millennials' 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 & Sadaghiani, 2010);
- they are members of a much wider number of interlocking communities than previous generational cohorts (Statistics Canada, 2006). 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);
- 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., 2009); and
- 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., 2007). In a fundamental way, communication is conceived as a multilogue, a conversation, rather than as a dialog. They have been called "ambient broadcasters" (Pew, 2010). Furthermore, the audience component of a communication is often not a coherent shared-interest group but something more *ad hoc* ("friends") (Jacobs & Diefenbach, 2012). This creates a plethora of problems:
  - o 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 particular recipient may not necessarily have enough context to understand their full meaning.
  - o Dissemination is not controllable by the original sender, and the technology makes it easy to pass communications on, far beyond their intended reach.
  - o There are no gatekeepers to control what does and does not get disseminated – individuals decide for themselves (Johnson & Kaye, 2010).

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). At the mundane level, they are accustomed to living their lives under the pervasive gaze of cultures of social surveillance (Bauman et al., 2014) 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., 2007). 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 (Drapeau & Wells, 2009; Hibbard, 2011).
- 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.

5. The way forward: computing 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 cyberwalls 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. Two out of three of these are social not technical; and all three are not problems to be solved, but forces that require a response.

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. 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 cyberwall 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".

The three forces of organizational change towards responsiveness, increasing weakness of technical defensive solutions, and changing attitudes in the populations of technological societies are working synergistically to destroy the existence of a "safe inside zone" where important work can get done.

A new kind of solution is needed (Karas, Moore, & Parrott, 2008). Although still in its infancy, the most hopeful direction is a strategy, or perhaps a metaphor, known as *computing in compromised environments*. It entails a paradigm shift from protecting a safe inside zone *per se* to protecting computations and their data by obfuscating and masking data both at rest and in flight, and by effecting computations in deliberately unique ways each time they execute. The strategy's 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 deterrents, and not as protection.

While there may be opportunities to mitigate, in a small way, changing organizational imperatives and the attitudes of a techno-population, solutions must be primarily technological. 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 lifespan of the virtual machine. Cloud services already make use of a version of this idea, and significant difficulties have already been encountered;
- 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) or selected from a pool of multiple versions of the same actions. 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. The number of combinations grows much faster than the number of versions of each piece of code. Investing in a few high-quality versions of each piece, therefore, provide a large return on the extra effort invested. Although this increases costs, these need to be compared with the costs of intrusions and other successful attacks;

- modeling 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, understanding and categorizing the transit in isolation. Once an attacker is “inside” there is often much less scrutiny. Behavior modeling 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; and
- 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.

  A password is a single secret that is shared by both ends of an authentication. If the secret is extracted from either end, the usefulness of the mechanism is completely destroyed. In contrast, extracting one part of a secret provides little or no information relevant to the other parts. Secret sharing can, therefore, 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. Most critical organizational knowledge is not about the existence of information, but the values of certain components (the strategic direction, the value of a tender). Suppose there is a document containing this kind of information that the organization does not want exfiltrated. The system creates multiple copies of the 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. For a tender bid, for example, the only part that needs to vary between the true and false versions is the actual amount.

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

  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.

Although many of these ideas are in the early stages of research, some sense of the evolution can be sketched by comparing a bread-and-butter example of an online activity as it is done today, contrasted with how it might be done in the future. Consider the purchase of an item from an online retail website. The purchaser interacts with the website using a web browser. The choice of web browser necessarily determines the encryption (or not) used for the transaction; the website owner determines the security level of the purchaser's authentication; the transaction probably already takes place within a virtual machine in the retailer's back-end system; validating the payment requires another interaction between the retailer and the credit card business. The main vulnerabilities are in the user's authentication (where a failure would allow someone else to order products for free), the back-end system (where a failure would allow products to be scheduled for delivery but not paid for), and the credit card transaction (where a failure would allow credit details to be captured for illicit use or sale).

Potential developments in the future center on better authentication and reducing the monoculture of interaction types. At the user's computer, the user's browser could run in a virtual machine created for the session, and the browser itself could be assembled from individual components so that no two sessions would be exactly alike. The interaction between browser and website could choose, at runtime, from a number of possible encryption systems and keys. The retailer's back-end order-completion system could run in a virtual machine that is assembled, on-the-fly, from individual components. The interaction of retailer and credit-card company could use the same kinds of refinements as the user-retailer interaction. If there are only 2 different ways to implement each step, then there are 2^k different structures for a *k* step interaction, so that even modest variability at the step level leads to much larger variability overall. An attacker must be prepared for the larger number of combinations, turning the conventional advantage possessed by attackers on its head.

These ideas are still in the early stages of development. However, they seem to hold more promise than trying to build higher and thicker walls, and to 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 social and organizational factors and their operational, conceptual and theoretical implications.

References

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

Bauman, Z., Bigo, D., Esteves, P., Guild, E., Jabri, V., Lyon, D., & Walker, R.B.J. (2014). After Snowden: Rethinking the impact of surveillance. *International Political Sociology*, *8*(2), 121–144.

Bloomberg Business (2011, June 27). 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

CSIS and McAfee (2014). Net losses: Estimating the global cost of cybercrime. *Washington, DC: Center for strategic and international studies*. Santa Clara, CA: Intel Security McAfee.

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).

Please cite this article as: Leuprecht, C., et al., Beyond the Castle Model of cyber-risk and cyber-security, *Government Information Quarterly* (2016), http://dx.doi.org/10.1016/j.giq.2016.01.012

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.

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

Fritzson, A., Howell, L.W., & Zakheim, D.S. (2007). Military of Millennials, Strategy + Business. *49* http://www.strategy-business.com/article/07401?pg=0.

Gill, M. (2006). *The handbook of security*. New York: Palgrave Macmillan.

Harknett, R.J., & Stever, J.A. (2011). The new policy world of cybersecurity. *Public Administration Review*, *71*(3), 455–460.

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

Hershatter, A., & Epstein, M. (2010). Millenials and the world of work: An organization and management perspective. *Journal of Business and 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., & 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.

Leydesdorff, L. (2010). The communication of meaning and the structuration of expectations: 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.

Pempek, H., Yermolayeva, Y., & Calvert, S. (2009). College Students Social Networking Experiences on Facebook. *Journal of Applied Developmental Psychology*, *30*(3), 227–238.

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.

Sandhu, R.S., Coyne, E.J., Feinstein, H.L., & Youman, C.E. (1996). Role-based access control models. *IEEE Computer*, *29*(2), 38–47.

Statistics Canada (2006). *Canada's Ethnocultural Mosaic*. Ottawa: Census, 2008, Minister of Industry Catalogue no. 97-562-X.

Sawila, R., & Skillicorn, D.B. (2012). *Course-of-action decision support using partial graph cuts, IEEE conference on technologies for homeland security, Nov 2012, 291–297*.

Shirky, C. (2008). *Here comes everybody: The power of organizing without organizations*. New York: Penguin Press.

Verdon, J. (2012 April). *The wealth of people: How social media re-frames the future of knowledge and work*. Brussels: North Atlantic Treaty Organization RTO-MP-HFM-201.

Wall Street Journal (2014 Feb 6th). Target breach began with contractor's electronic billing link. www.wsj.com/articles/SB10001424052702304450904579367391844060778 (accessed Sep 18th 2015)

Zedner, L. (2009). *Security*. Abingdon: Routledge.

Christian Leuprecht is Professor of Political Science at the Royal Military College of Canada and Senior Fellow at the Macdonald Laurier Institute. He holds a ministerial appointment to the governing Council of the Natural Sciences and Engineering Research Council of Canada, is president of the International Sociological Association's Research Committee 01: Armed Forces and Conflict Resolution, and a United Nations Security Structure Expert. He is cross-appointed to the Department of Political Studies and the School of Policy Studies at Queen's University where he is also a fellow of the Institute of Intergovernmental Relations and the Queen's Centre for International and Defence Policy. As a foremost expert on security and defense, political demography, and comparative federalism and multilevel governance, he is regularly called as an expert witness to testify before committees of parliament.

His award-winning publications have appeared in English, German, French, and Spanish. His over 100 publications include 9 books and nearly 40 articles, most of which are available as free downloads: , http://www.christianleuprecht.com/. Recent articles have appeared in *Armed Forces and Society* (2015), *Global Crime* (2015, 2013), the *Canadian Foreign Policy Journal* (2014, Maureen Molot Prize for Best Article), *Canadian Public Administration* (2014), the *Canadian Journal of Political Science* (2012, 2003), *Regional and Federal Studies* (2012), and *Terrorism and Political Violence* (2011). His editorials appear regularly across Canada's national newspapers and he is a frequent commentator in domestic and international media.

Leuprecht has been a visiting professor at the Université Pierre-Mendès France (2015), the University of Augsburg in Germany (2011), the Swedish National Defence College (recurring) and the European Academy (recurring), the Bicentennial Visiting Associate Professor in Canadian Studies at Yale University (2009–2010). He is a research affiliate at the National Consortium for the Study of Terrorism and Responses to Terrorism (since 2005), the Network for Terrorism, Security, and Society (since 2012), l'Université de Montréal's International Centre for Comparative Criminology (since 2014), the Centre interuniversitaire de recherche sur les relations internationales du Canada et du Québec (since 2015), l'Observatoire sur la radicalization et l'extrémisme violent (since 2015), the Austrian Institute for European and Security Policy (since 2010), the Solomon Asch Center for Study of Ethnopolitical Conflict at the University of Pennsylvania and Bryn Mawr College (2003), the World Population Program at the International Institute for Advanced Systems Analysis in Vienna, Austria (2002), and held doctoral (2001–2003) and postdoctoral (2003–2005) fellowships from the Social Sciences and Humanities Research Council of Canada. He holds a Ph.D. from Queen's University (2003), and graduate degrees in Political Science (1998) and French (1999) from the University of Toronto as well as the Institut d'Études Politiques at the Université Pierre-Mendès France in Grenoble (1997).

Since joining RMCC in 2005, he has served as Associate Dean of the Faculty of Arts and Deputy Head of the Department of Political Science and Economics. He is the recipient of the RMCC Commandant's Commendation for Excellent in Service. A long-time proponent of experiential learning, Leuprecht has also been nominated and short-listed repeatedly for RMCC's Teaching Excellence Award. He is a member of the editorial boards of *Commonwealth* & *Comparative Politics*, *Current Sociology*'s Manuscript Series, and the Springer book series in Advances in Science and Technologies for Security Applications. He has also been a member of the editorial board of the Queen's Policy Studies series published by McGill-Queen's University Press.

David Skillicorn is Professor in the School of Computing at Queen's University, and an adjunct Professor at the Royal Military College of Canada. His research interests are in knowledge disovery, particularly for counterterrorism, law enforcement, and fraud. He has authored more than a hundred papers, and several books. His Ph.D. is from the University of Manitoba, and his undergraduate degree from the University of Sydney.

Please cite this article as: Leuprecht, C., et al., Beyond the Castle Model of cyber-risk and cyber-security, *Government Information Quarterly* (2016), http://dx.doi.org/10.1016/j.giq.2016.01.012