ALE | Where Everything Connects

Campus cybersecurity in the age of IoT and GDPR

White Paper

Campus cybersecurity in the age of IoT and GDPR

Alcatel·Lucent Enterprise

Table of contents

Introduction

Campus Cybersecurity is a perennial 'top of mind' topic for Higher Education. For the second time in three years, information security topped Educause's annual Higher Ed CIO Top 10 IT Issues survey[1]. This should come as no surprise according to the latest Verizon Data Breach Investigations Report (DBIR)[2]. The report, which shows cybercriminal activity trending upwards, identifies the three most targeted industries as Financial and Insurance, Healthcare, and Education.

Information security in education has always been a competition between ease-of-use and total security; sometimes security wins, but most often ease-of-use is the priority. Throw in the new European privacy legislation known as General Data Protection Regulation[3, 4], (GDPR) the introduction of a broad variety of IoT devices, and it is not surprising that Information Security, and protection of the institution's reputation, is foremost on the minds of Higher Education CIO's.

This focus on Information Security is timely. Never before has information technology played such a large role in education. Whether researching a topic for class or submitting classwork, information technology is at the heart of pedagogy, and the strategies implemented to identify and assist student academic success.

The 2018 Verizon Data Breach Investigations Report provides some very interesting insight on the attack vectors and motivations for malicious cyber activity.

Table 1. 2017 Education Industry Data Breach summary

Frequency	292 incidents, 101 with confirmed data disclosure
Top 3 patterns	Everything Else, Web Application Attacks and Miscellaneous Errors represent 76% of breaches
Threat actors	External (81%), Internal (19%), Partner (2%), Multiple parties (2%) (breaches)
Actor motives	Financial (70%), Espionage (20%), Fun (11%)
Data compromised	72% personal, 14% secrets and 11% medical

Figure 1. Patterns seen in education breaches

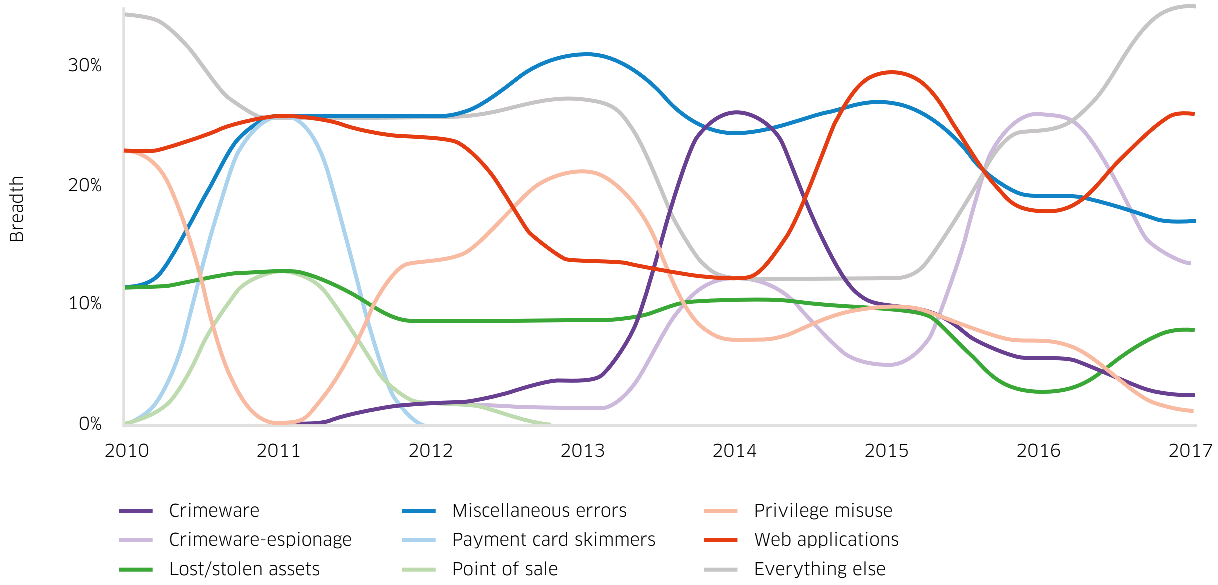

[1] EDUCAUSE Center for Analysis and Research (ECAR), "Top 10 IT Issues 2018," EDUCAUSE Research Snapshot, *EDUCAUSE Review*, January 29, 2018.

[2] Verizon Enterprise, "Verizon Data Breach Investigations Report (DBIR) 2018", April 2018, pages 29-30

[3] https://eugdpr.org

[4] https://gdpr-info.eu

This whitepaper explores the implementation of risk-based security strategies, such as the 'defense in depth' concept. This concept promotes protecting a computer network with a series of individual defensive mechanisms such that if an attack is launched, the attacker will need to defeat multiple, independent layers to succeed.

The "CIA triad" will be referred to throughout this paper. This construct identifies cyber targets that possess "Confidential" information, that impact the "Integrity" of the information, or deny "Availability" of the information. The CIA triad is useful when formulating campus-wide security policies and assigning risk values. For instance, a stolen or lost laptop that contains Personally Identifiable Information (PII) would be classified as a Confidential information attack; a hacker who successfully alters grades would be violating the Integrity of the Student Information System; and, BYOD printers in students dorms could be co-opted to participate in a Denial of Service (DoS) attack, which would be classified as an Availability attack, because the purpose of a DoS attack is to flood a web server with so many bogus requests that the server will stop responding to legitimate requests.

For the purposes of this paper, information security threats will be referred to as malicious cyber activity, which is defined by the U.S. Federal government as:

> "Activities, other than those authorized by or in accordance with U.S. law, that seek to compromise or impair the confidentiality, integrity, or availability of computers, information or communications systems, networks, physical or virtual infrastructure controlled by computers or information systems, or information resident thereon."

Effects of malicious cyber activity include the following:

> "The manipulation, disruption, denial, degradation, or destruction of computers, information or communications systems, networks, physical or virtual infrastructure controlled by computers or information systems, or information resident thereon."

It has been a watershed year in 2018 with the flood of Internet of Things (IoT) devices on campus, and the introduction of privacy regulations in cyberspace, specifically the European Union's General Data Protection Regulation. The GDPR has ramifications for universities world-wide and both topics will be discussed in depth.

General Data Protection Regulation

The GDPR was enacted May 25, 2018, by the European Union to give back citizens control of their cyber information. Essentially, the GDPR enables every EU resident the right to know and decide how their personal data is being used, stored, protected, transferred and deleted. Individuals also have the "right to be forgotten" by requesting deletion of all of their data. Unlike other countries' definition of personally identifiable information, GDPR also covers location data, including IP addresses – which could have a major impact on deploying Location Based Services (LBS) on campus or even network equipment logs.

As the most progressive privacy legislation in the world, in addition to the rights spelled out above, the GDPR also has power to enforce any non-compliance. GDPR violations could result in non-compliance penalties of 4% of annual world-wide revenue or $23 million USD whichever is greater.

Planning for GDPR would be included in the Confidentiality section of the CIA triad and will be specifically addressed in the Application Security section.

Internet of Things (IoT)

The Internet of Things is an interesting topic. Unlike Software-Defined Networking (SDN) which relied on implementation consistency from the networking community, IoT devices are already on the campus and their footprint will continue to grow. An IoT device is a connected device that has the ability to send and/or receive information without the need for human intervention to operate it. On a campus, this could include a consumer Wi-Fi printer, security cameras and IoT sensors, or even lecture hall/classroom projectors.

Under normal usage, these devices are essentially innocuous. However, since they are network devices and they have an operating system, they are susceptible to hacking and malware. – putting them at risk to be recruited and enslaved into a bot army. You may recall the October 2016 DDoS (distributed denial of service) attack[5] against Dyn, a domain name services provider based in the U.S. with clients in Europe and North America. Dyn was attacked three times in the same day. The resulting analysis confirmed that IoT devices (cameras, baby monitors, Wi-Fi routers, and printers) had been compromised with a malware variant based on the Mirai virus source code.

Planning for IoT mostly resides in the Availability section of the CIA triad and will be specifically addressed in the Network Access Security section.

Defense in depth

Defense in depth is the practice of defending a computer with layers of independent security devices or strategies. It was originally conceived by the U.S. National Security Agency (NSA) as a comprehensive approach to protecting information. There are three areas of focus in this construct:

1. People: Information assurance is the goal for IT leadership and includes the application of security services such as: Availability, Integrity, Authentication, Confidentiality, and Non-repudiation. The application of these services should be based on the Protect, Detect, and React paradigm. Also included here would be hiring/firing practices and training.
2. Technology: Includes hardware and software that prevent access to the contents of a system.
3. Operations: Focuses on all the activities required to sustain an organization's security posture on a day-to-day basis and can include system security assessments, recovery and reconstitution, change control and data handling.

The remainder of this whitepaper will explore these areas of focus through the exploration of the following layers: Network Edge Security; Application Security; Network LAN Security, Network Segmentation and Network Management.

Network edge security

The edge of the network is where the institution's internal network interfaces with another network, including a carrier, REN or public Internet. This is the first layer in a defense in depth architecture. This layer can also be viewed as protecting Information Assurance from the "Barbarians at the Gate" and their malicious activity.

The Network Edge has both incoming and outgoing traffic, which means that we will need to address each traffic type.

[5] https://en.wikipedia.org/wiki/2016_Dyn_cyberattack

Incoming traffic

Incoming traffic: Whether the traffic is from a WAN link, the public Internet, or a research network link, some or all of following technologies should apply:

1. Firewall: Normally residing in the DMZ (demilitarized zone), a firewall can accomplish several security tasks, including NAT (network address translation), VPN (virtual private network) and of course, applying security logic to traffic.
2. Intrusion Detection System (IDS)/Intrusion Detection and Prevention System (IDP): Are important in protecting the network from attack. There are two types:
   1) Signature based which recognizes patterns of known exploits
   2) Anomaly based which recognizes deviations from 'base' network activity
3. Virtual Private Network (VPN): Provides an encrypted connection to the local network. This type of connection is best for secure communications and usually employs Multifactor Authentication schemes (MFA).
4. SPAM filters: Considered by many as an email server firewall. SPAM filters have grown in sophistication and capabilities to identify malware infected attachments, spoofed addresses (phishing) and other attacks. According to the Verizon Data Breach report, phishing is one of the top attack vectors for Higher Ed. Implementing a SPAM filter can help reduce the number of attacks.
5. Web traffic filter: Some firewalls can accomplish this task, and to be sure, not all institutions implement a Web traffic filter. However, this technology does enable a university to permit or deny access to web sites from their network. Most often seen in K-12, this security device helps to keep users out of sites that may lead to infection or loss of control.
6. Network security monitoring: Applications like Bro[6] or Splunk[7] can provide useful data on network traffic and potential security anomalies.

Outgoing traffic

In addition to requesting traffic, a network will also send information outside of the network. While this type of traffic is not usually suspect, monitoring it can inform the institution if there is a security breach (such as IoT malware compromised devices participating in a DDoS attack).

Outgoing traffic must traverse the Incoming traffic defenses, but there is one technology that is useful in protecting outgoing traffic, that is encryption.

1. MACSec Encryption is also known by the IEEE designation 802.1AE. This technology supports encrypted transmissions across a link. For instance, if you have a Disaster Recover (DR) site for your data center, you can populate that site via a MACSec encrypted link, ensuring data confidentiality, data integrity and data origin authentication.

Application security

The next layer in addressing today's cybersecurity threats is Application Security. This layer encompasses both end user computing devices, as well as network-based applications. Application Security is an important tool in a university's security arsenal. It addresses the Confidential and Integrity domains of the CIA triad, and is important to consider when implementing GDPR protections.

Technologies and tactics in this layer include:

1. Data encryption: Both major PC OSes have disk data encryption features. This is important as it protects the user's data from casual access. Encrypting network storage is gaining popularity and improves the integrity and confidentiality of the information.

[6] https://www.bro.org

[7] https://www.splunk.com

2. Multifactor Authentication (MFA): Multifactor authentication is a mechanism that requires two devices to work in concert in order to successfully access an application or resource. Usually, the flow is userid/password, then a string of numbers is input and texted to the user's smart phone. This provides another layer of confirmation that the person is who they say they are. Another mechanism to provide MFA is to use hardware tokens like those from RSA and Google instead of a smart phone text. This strategy has the advantage of being able to be leveraged in VPN and other secure communication strategies.
3. Micro-segmentation: The data center has evolved from physical servers to virtual servers, with dozens of virtual servers occupying the space of one or two physical servers. However, with this density, comes risk. Traditional stateful inspection firewalls do not have the capacity to analyze, at wirerate, a data center's traffic flows. VMWare[8], a major hypervisor provider, has implemented a technology called "micro-segmentation". This means each application can now have its own security perimeter without relying exclusively on VLANs.

Practices implemented in this layer should include:

4. End user security awareness training: This activity is usually conducted during Security Awareness Month (October). However, periodic reminders about phishing, spear phishing, and social engineering exploits (divulging credentials), should be issued throughout the year to help reduce this attack vector's surface area.
5. Security Patches: It is important to test security patches and application updates prior to applying them. However, releases from the manufacturer are important and should be prioritized. According the Verizon Data Breach report, 6% of successful attacks exploited security holes that a patch would have covered. In fact some of the most damaging attacks have come from delayed patch application.

Network access security

Local Area Networks (LAN) and Wireless Local Area Networks (WLAN) are entrance points for campus computing users. A LAN connection usually means physically connecting an Ethernet patch cord from the network device, to a RJ-45 wall port. A WLAN connection does not require physical connectivity, and instead uses the built-in Wi-Fi chip in the device to see and connect to the network.

Universities have always had to balance ease-of-use against total security. It is important to not make secure access too burdensome, otherwise students and others will find a work around, or increase "Shadow IT" devices. For instance, in some campuses, a guest will need a sponsor to be able to authenticate to the network; in others, it's OK to send the guest to a self-service portal where they input personal information about themselves, agree to acceptable usage language and can then sign in. Both provide the same protection to the network, but one is a bit more onerous than the other.

In a defense in depth architecture, Network Access Security is one of the most important investments to make. This layer is all about the 4 A's: Authenticate, Authorize, Audit and Administer. Successful authentication leads to authorization to the allowed network resources for a specific role. Audit is about monitoring network traffic and behavior, and if anomalous, Administer quarantine or network traffic rules.

Network Access Security technologies that support the 4 A's include:

1. 802.1X: An IEEE standard, port-based network access control mechanism that provides an authentication mechanism for both LAN and WLAN devices. 802.1X requires a supplicant (computer or device requesting access); authenticator (usually the Ethernet switch or WLAN AP or controller); and an authentication server (usually a RADIUS or EAP server). The supplicant supplies the credentials to the authentication server, if correct, the authenticator grants access to the network.

[8] https://www.vmware.com

2. Biometric: A new and potentially easier to use security authentication method is to utilize biometrics, or something unique to the user, such as a fingerprint, iris scan or voiceprint. While this won't help with IoT devices, it does provide an easy way to securely grant a human access to the network.
3. Encryption: While previously referenced as security mechanisms for data at rest and in transit from data center to data center, MACSec encryption is becoming more prevalent in the LAN – from core to edge. In addition, IEEE WLAN standards have addressed encryption with the 802.11i standard, which protects the Wi-Fi transmissions from access point to user from Man-in-the-Middle attacks (one of the main reasons to not use public, unencrypted Wi-Fi). As well, the Wi-Fi alliance recently announced WPA3 (Wi-Fi Protected Access 3). WPA3 adds new features to simplify Wi-Fi security, enable more robust authentication, deliver increased cryptographic strength for highly sensitive data markets, while maintaining the resiliency of mission critical networks.
4. Hardened OS: Many network devices are purchased with default administrator authentication profiles. Unfortunately, and most often in the case of IoT devices, these defaults are never changed, making the device an easy target for malware insertion or even code change. A hardened device OS is one that scrambles the code and memory location so that a successful breach of one device will not be successfully repeated in an automatic fashion. In addition to protecting the integrity of the device, a hardened OS can also possess the following features:
    1) DoS attack awareness and mitigation: Some network infrastructure devices can recognize that a DoS attack is in progress and immediately thwart the attack by dropping the offending traffic.
    2) IP based attack awareness: Some network infrastructure devices are able to integrate with SNORT or other IDS/IDP products and react to a positive IP attack signature and quarantine the offending traffic.
5. Unified Network Access Policies: After successful authentication, this structure authorizes network access based on parameters such as MAC address, time of day, user role, department (such as student, guest, faculty, staff, vendor, administration, admissions, athletics) or even the location from which they are authenticating. This structure supports both LAN and WLAN access. It eliminates security profile mis-matches and duplication, and provides consistent, secure network access.
6. IoT Devices: Unified Network Access Policies are a critical feature for IoT device enablement. Additionally, DHCP Device Fingerprinting can quickly identify IoT devices from multiple manufacturers. This feature leverages DHCP options that provide vendor-specific information about the device hardware or operating system. The exchange is done using DHCP options as defined by RFC 2132.[9] Utilizing DHCP options provides vendor, device, and OS information, which combined constitute the device "fingerprint". For example, a recent CCTV RFP was awarded to single manufacturer. These new products will coexist with legacy CCTV cameras. With the implementation of UNAP, the security team has two ways to identify and apply network security rules: DHCP Fingerprinting, or MAC address masking. MAC address masking uses the first 24 bits of the MAC address, which contains the Organizationally Unique Identifier (OUI)[10]. Leveraging UNAP with a masked MAC address policy would allow the institution to quickly and securely roll out their new surveillance cameras. DHCP Fingerprinting would allow the campus to identify all CCTV cameras and apply consistent security policies.

DHCP Device Fingerprinting is usually supported by WLAN systems, however, it is most powerful when LAN devices are included as well, as not all IoT devices connect to the network via Wi-Fi.

[9] http://www.ietf.org/rfc/ rfc2132.txt

[10] http://standards-oui.ieee.org/oui/oui.txt

Figure 2. Auto recognition and classification

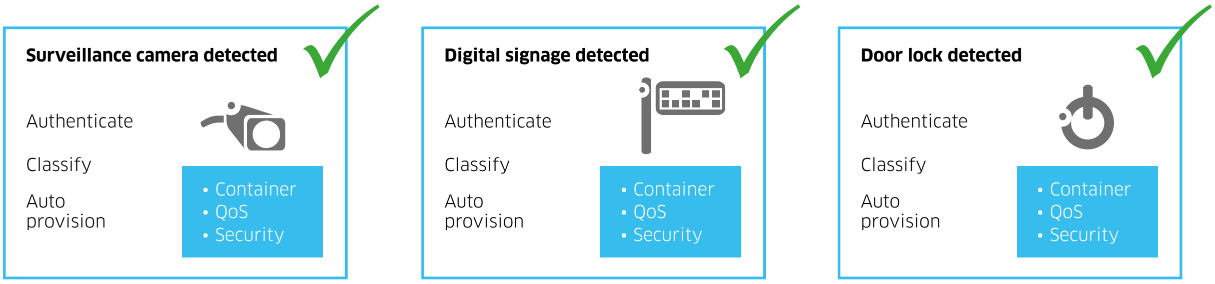

Network segmentation

With the introduction of LAN switching, network segmentation was expanded from physical segmentation to include virtual LANs (VLANs). This provides the ability to limit network services to users who are members of the VLAN, essentially securing applications and services. Most network equipment supports up to 4,096 VLANs, which should be more than enough, but in reality can be a challenge and must be architected properly.

The topic of VLAN exhaustion is real and has been addressed in several technologies used by campuses. MPLS (Multiprotocol Label Switching)[11] and its derivative, Shortest Path Bridging (SPB – IEEE 802.1aq)[12], both use the concept of ‘service interfaces’ to further segment network activity.

Figure 3. SPBM – Network Segmentation, No Spanning Tree

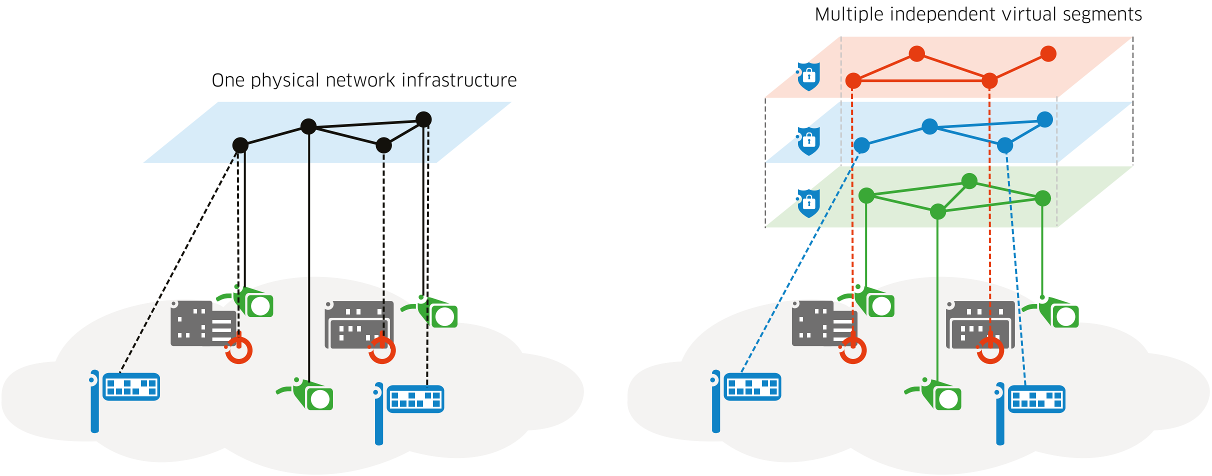
One physical network infrastructure
Multiple independent virtual segments

Network management

Network management is an often overlooked application from infrastructure equipment providers. There are many reasons for this. Chief among them is the prevalence of third party applications that manage multi-vendor environments. However, there are trade-offs when utilizing a third party Network Management system (NMS). Many vendor provided NMS platforms can integrate with support entitlements, and even suggest maintenance releases that address potential bugs in deployment. In addition, the OEM NMS can provide a wealth of statistics, analytics and trends that can allow proactive network management.

[11] https://en.wikipedia.org/wiki/Multiprotocol_Label_Switching

[12] https://en.wikipedia.org/wiki/IEEE_802.1aq

IT professionals in education are also starting to look at Deep Packet Inspection (DPI) capabilities for the LAN infrastructure. The ability to know which applications are consuming the most bandwidth can inform the entire IT department as to where investments need to be made. For example, if traffic to/from the LMS application is the top consumer of bandwidth, it would probably be a good to review the resources that are assigned to the LMS VM, or if it is cloud-based, to investigate the user experience and determine if further investment is needed to increase bandwidth, memory, processing or storage.

With regards to cybersecurity on campus, the NMS is the platform where it all begins. From enabling SSL/encrypted CLI access to the infrastructure equipment, to implementing unified network access policies, to monitoring traffic for anomalies, to quarantining misbehaving devices/users, the NMS is the platform that can accomplish all of this and much more.

In addition to a strong NMS, some third party programs should be considered:

1. perfSONAR[13]: While this is tool used by many research institutions, it has capabilities that can help you understand your complete network performance. Here's an excerpt about what it is meant to do:

   "Ensuring that things are operating well, on an end-to-end basis, is critical. Monitoring within a single domain is a common and accepted practice; cross-domain performance monitoring is difficult to do with traditional tools. perfSONAR is a widely-deployed test and measurement infrastructure that is used by science networks and facilities around the world to monitor and ensure network performance."
2. Tools within perfSONAR that provide additional troubleshooting:
   1) pScheduler: Throughput tests to remote locations
   2) OWAMP: Continuous checks for latency and packet loss

Summary

Education institutions are the third most targeted industry in the world. Cybersecurity in education is more than a 'nice to have'. It is a critical dimension in enterprise architecture and can contribute to a university's positive brand recognition.

Implementing a risk-based security plan allows the university to allocate budget according to need or risk. Following the defense in depth architecture ensures that an attacker must defeat different technologies throughout the exploit in order to succeed.

Employing the CIA triad of Confidentiality, Integrity and Availability classifications in your cyber assets helps to determine the risk to the institution, and weight to the topic.

Implementing the 4 A's (Authentication, Authorization, Audit and Administration) provides a unified structure to network access and behavior on both LAN and WLAN networks.

Segmenting the network with MPLS or SPB allows granular control of services and the devices/users accessing those services.

Lastly, training: Security technology, and architecture will take you a long way in protecting your assets. However, as identified in many studies, phishing and user error are the most prevalent breach methods. Training your students, faculty, staff and vendors about cybersecurity can help reduce your number one risk factor.

[13] https://www.perfsonar.net/about/what-is-perfsonar/

References and resources

National Institute of Standards and Technology: This link takes you to the landing page for NIST's cybersecurity framework: https://www.nist.gov/cyberframework

EDUCAUSE: A non-profit association that helps Higher Education elevate the impact of IT. The 2018 annual "Top 10 IT Issues" survey results and links to other resources can be found here: https://er.educause.edu/articles/2018/1/top-10-it-issues-2018-the-remaking-of-higher-education