



What the FDA Changed & How to Best Prepare

New FDA Pre-Market Submission Guidelines for Cybersecurity in Medical Devices

March 2019

Abstract:

The FDA has changed the way it views Cybersecurity in regards to premarket submissions for medical devices. Christopher Gates, Principal System Security Architect at Velentium, will break down what has changed in the new guidance, as well as provide instruction as to how to best comply with the requirements. Christopher has spent over 30 years developing medical devices and has participated in regulatory and standards bodies pertaining to Cybersecurity in order to better define tools, techniques, and processes to enable the creation of secure products. He has worked with several industry-leading device manufacturers, as well as the National Telecommunications and Information Administration (NTIA), the U.S. Department of Commerce, MITRE, Bluetooth special interest groups, and IEEE, in order to present, define, and codify techniques to control embedded cybersecurity.

Contents

Background	3
Device Tiering.....	4
Trustworthy Devices	5
Prevent Unauthorized Use.....	6
Code, Data, & Execution Integrity	7
Maintain Confidentiality of Data	7
Attack Detection	8
Device Configuration	8
Software Configuration Management	9
Incident Management.....	9
Cybersecurity Labeling.....	10
Cybersecurity Documentation Required for Submission.....	12

Key Takeaways

- 38 new design mitigations
- 14 new Cybersecurity Labeling & IFU procedures
- Description of handling of field software updates and patches
- Assignment and justification for a security tier assignment
- 20 new detailed security artifacts generated during the development lifecycle in the submission package

Background

In October 2014, the FDA released the first version of its Premarket Cybersecurity Guidance. This early version conveyed the FDA's heightened interest in ensuring secure medical devices do no patient harm, but was unclear about exactly what steps needed to be taken in order to meet the FDA's expectations.

Following that release, Velentium has reviewed dozens of clients' pre-submission meeting minutes, which included 483s, warning letters, and pre-market rejections that detailed the FDA's evolving expectations for security in new medical devices and their supporting systems. This cumulative insight allows us to tailor our client's development activities and generated artifacts to meet the FDA's threshold.

**GUIDELINES = FDA
REQUIREMENT**

Most of the activities and deliverables in the new 2018 guidance have been the FDA's expectations for the past 12 to 18 months, but

were not publicly documented and available to the industry.

Make no mistake, the content and approach described in this guidance is required. A recent Office of the Inspector General's report made three distinct recommendations to the FDA regarding medical device cybersecurity:

- Promote the use of pre-submission meetings to address cybersecurity and related questions
- Include cybersecurity documentation as a criterion in the FDA's "Refuse to Accept" checklist
- Include cybersecurity as an element in the smart template

In response, the FDA stated that *"it welcomes the OIG report as a means for strengthening the agency's already robust premarket review of networked medical devices. The FDA has already taken steps to implement these recommendations, and plans to update OIG as these items are completed."*

In the past, some manufacturers believed they would be able to justify **not following** the proper procedures needed to create a secure device by performing risk-benefit analysis on which they could base their case to the FDA reviewer in a premarket submission.

However, with the changes brought about by the report from the OIG and the FDA issuing its new guidance in response, any submission that fails to include all mandated security artifacts stipulated in the "Refuse to Accept" checklist will be rejected by a clerk. The manufacturer's justifications for ignoring security will never be seen by a reviewer.

Although this seems strict, the OIG and FDA now encourage manufacturers to utilize pre-market submission meetings to address cybersecurity questions. This practice can significantly reduce delay caused by manufacturer's misunderstandings and subsequent rejection of a pre-market approval request.

Although these guidelines were released as a “draft”, the FDA regards them as their current expectations for all new submissions, superseding the October 2nd, 2014 guidance. The FDA has adopted a prescriptive approach to clarify their expectations while raising the expected cybersecurity level of all new devices.

Cybersecurity expectations extend throughout the entire lifecycle of the product, including the complete development cycle (not just a post-design afterthought), as well as post-market. Even after a product is retired, manufacturers are now required to keep track of devices for up to a 5-year period.

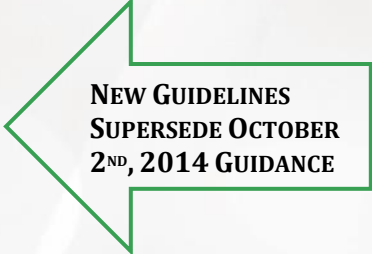
In addition, proprietary protocols are no longer deemed secure. Security researchers and hackers reverse-engineer and exploit these measures every day. No matter the level of expertise used, if you are not using strong and proven cryptographic primitives, then a proprietary protocol does not stand a chance of being secure. While utilizing cryptographic primitives does not ensure you have a secure device, it is a fundamental first step in the overall protection of a device.

In summary, the new expectations include:

- 14 new areas of coverage for security topics in labeling/Instructions for Use (Section 6)
- Handling description of field software updates and patches
- 38 new design-based mitigations (Section 5)
- Assignment and justification for a security tier designation
- 20 new detailed security artifacts generated during the development lifecycle included in the submission package (Section 7)

- Software & Hardware Bill of Materials utilized in the product – cross-referenced against known vulnerabilities
- Requirement stating protection mechanisms should prevent all unauthorized use through all interfaces
- Traceability matrix linking all security artifacts into requirements and hazard analysis

This new guidance is a welcome step forward, as it accurately communicates to designers and manufacturers the FDA’s expectations for securing medical devices.



**NEW GUIDELINES
SUPERSEDE OCTOBER
2ND, 2014 GUIDANCE**

Device Tiering

Our first topic concerns classification of medical devices according to the new FDA guidance. Two separate tiers of classification are described:

- Tier 1 – Higher Cybersecurity Risk
 - The device is capable of connecting (wired or wireless) to another medical or non-medical product, to a network, or to the internet **AND**
 - A cybersecurity incident affecting the device could directly result in patient harm to multiple patients, e.g. scaled attacks, which means an attack on one device that

exposes information in order to attack a larger set of devices

- Tier 2 – Standard Cybersecurity Risk
 - A medical device for which criteria for a tier 1 device are not met, AKA “everything else”

It should be noted that these cybersecurity tiers have nothing to do with any other type of classification of safety or risk as it pertains to a medical device or software. When evaluating whether a device is Tier 1 or 2, the first question should be “does the device have any form of connectivity over any medium to any other device, medical or otherwise?” Most products made in the last 20 years do. Yes, smartphones count! So do custom devices. Although the device itself may not

LARGE MAJORITY OF ALL MEDICAL DEVICES WILL BE DESIGNATED AS TIER 1

directly connect to the internet, if it communicates with any device, it does count toward Tier 1 classification.

The next Tier 1 identifying question is, “can the device be manipulated to harm a patient? It should be noted that this is not described on some scale of risk, where there is a documented difference between the creation of a blister on a finger and fatalities. Both in this case are considered harm. Very few medical devices are able to say they cannot cause any harm, but for the purposes of being considered a Tier 1, the potential for causing harm to multiple patients must be present. If all of these criterial are not met, then it is a Tier 2 device.

As an example, consider a patient-worn device that communicates via Bluetooth Low Energy to a smartphone, where an app controls device temperature. In order for the device to accept the app’s commands, the app must authenticate itself to the device via credentials common to all instances of the same device model. Under normal operating conditions, the device’s temperature can range between 100- and 170-degrees Fahrenheit.

When classifying this device according to the new guidance:

- ✓ This device could cause harm via thermal burns
- ✓ This device communicates with another device via Bluetooth Low Energy
- ✓ Multiple devices could be attacked if the common credentials were exposed

Therefore, it is a Tier 1 device.

Trustworthy Devices

Designing and developing trustworthy medical devices, according to the new FDA guidance on cybersecurity, will be our focus for the next two posts. [The guidance’s](#) stance is as follows:

“In particular, devices and systems should be designed to protect assets and functionality in order to reduce the risk of multi-patient harm due to the loss of:

- authenticity
- availability

- integrity
- confidentiality

Specifically, protection mechanisms should prevent all unauthorized use (through all interfaces), ensure code, data and execution integrity, and as appropriate, protect confidentiality of data (insofar as its release could be leveraged to effect multi-patient harm).

As a part of premarket submissions, manufacturers should submit documentation demonstrating how these design expectations are met.”

There’s a lot to unpack in the above statement and the changes it introduces. Let’s begin with a hypothetical example:

Assume there is a medical device which contains a serial port that is exclusively used for developers to interact with the device itself, and is not part of its normal operating communications system. The port’s only function is to transmit non-patient-related, non-patient-identifiable device operating information.

Under previous standards, this would have been an interface a manufacturer could have justified leaving unprotected, due to lack of therapy-modifying functionality, zero risk for exposure of patient data, and minimal potential for disruption of device operation. Under the new guidance, however, there is no leeway for interpretation. Any port that potentially could be used to infiltrate the device must be secured.

Moving on to consider individual requirements, we come to:

Prevent Unauthorized Use

One of the first requirements named in the section excerpted above includes both people and devices. If two devices or systems connect to one another, each must be authenticated to the other in a cryptographically strong way.

To clarify, we must understand the difference between authentication and authorization:

the former being complex and difficult, and the latter relatively simple in comparison.



Authentication is a process by which two systems communicate with each other to determine that each party is in fact who they say they are. Authorization is granting access to exactly the data or actions allowable to a given authenticated user.

For example, a junior engineer can request access to the source files of a specific project and can be authenticated within the system, but their authorization would only allow them to view and edit a subset of those files. The lead engineer, once authenticated, would be authorized for unrestricted access to all

files. In security terms, this is referred to as “user privilege”.

Code, Data, & Execution Integrity

Assuring code and data integrity requires processes that most medical device manufacturers already follow; those who don't, will be soon. These processes entail a signature or hash of the code, protected in a cryptographically strong way. (A cyclic redundancy check (CRC) does not qualify, as it is not cryptographically strong. Manufacturers should instead look to crypto hashes such as SHA256, HMAC, CMAC, etc.)

While code and data integrity were covered in the previous FDA guidance, execution

**EXECUTION
INTEGRITY =
NEW GUIDELINE**

integrity is a new guideline that will be difficult to implement. In simplified terms, execution integrity means defining how

the device can verify that intended software execution is being executed in the way the device was designed.

As an example, execution integrity can be thought of as an aggressive watch dog timer. It would monitor that the correct functionality was being executed within the processor and during the proper time frames. This functionality can be found and provided by ARM or TrustZone-enabled microcontrollers.

This can be particularly challenging in small, resource-constrained microcontrollers. While difficult, engineers and developers should remember that maintaining the intended execution integrity within the device (aka “essential clinical performance”) is what's most important. This can be partly linked to code integrity in that once the device boots up, it can check to see if the right program exists and confirm it is executing the intended one.

Maintain Confidentiality of Data

The new guidance does not restrict data confidentiality to patient health information (PHI) or patient identifiable information (PII) only. It is intentionally all-inclusive, encompassing all data transmitted through these devices. It also defines data into two separate classes - “at rest” and “in motion”.

“At rest” data has been or is being saved to memory within a cellphone, hard drive, flash memory, etc. Its defining attribute is that it is **not** being actively transmitted and is “at rest” in a data cache. This data must be encrypted, as its encryption is the only acceptable solution for confidentiality.

“In motion” data is the exact opposite of its “at rest” counterpart. This is data that is being transmitted or communicated across devices. Examples include a Bluetooth Low Energy transmission to a cell phone, a phone sending data over Wi-Fi to a network access, or between two PCs over ethernet. Once again, encryption is the only security solution that is accepted by the FDA for this process.

The difficulty with encryption is that encryption/decryption activities require “keys”. If keys are compromised, this could potentially expose a larger attack surface, where if one device is infiltrated, it can lead to the penetration of all similar devices.

To mitigate the risk of exposure, do not utilize a “common key” for any operation where the same key value is applicable to all devices of the same type. Each device should require a unique key value. This is completely different from the discussion between utilizing symmetric vs. asymmetric encryption, as in either case unique keys or key pairs need to be utilized.

Attack Detection

The next criteria the FDA has laid out for manufacturers involves the ability to detect cybersecurity attacks in a timely fashion. Conveniently, the definition of a “timely fashion” has been given some flexibility to vary depending on the medical device in question. This could be defined in microseconds, milliseconds, hours, or even days, depending on the application. An appropriate definition must be included in your security documentation.

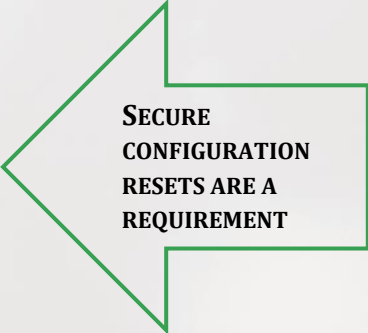
To comply, manufacturers must first have a system in place to log security events as they occur. Currently, few medical devices or systems log or store data appertaining to security events. An example would be attempting to authenticate to a device, or performing a data integrity check, which results in a failure. Such events should be placed into a security log.

These security events can occur by improbable releases of static electricity or a coordinated attack. The FDA is not looking for a deep level of forensic detail, but does want to make certain that these events are logged and accounted for as they take place -- even if there is little chance the log will ever be transmitted out of the device.

The guidance also makes note of forensic evidence capture (AKA “non-repudiation”). First, forensic evidence capture is used to make certain that the log’s integrity is maintained. As records are stored in this log, the integrity of the contents of each record and the integrity of the order of each record relative to the other records is continuously verified. This ensures the event logs cannot be tampered with while avoiding detection. In the event of patient harm or fatality, these logs could be used as forensic evidence in legal proceedings.

Device Configuration

If a device has been compromised, it is imperative that it can be reset to a safe & secure configuration. An example would be if a pacemaker loses its ability to communicate via BLE because of a “denial of service” attack, the device must continue to provide stimulation to the patient’s heart, ensuring its intended functionality remains uncompromised. The pacemaker must be designed to limit the impact from any vulnerability within its base mode of



**SECURE
CONFIGURATION
RESETS ARE A
REQUIREMENT**

operation, while retaining the ability to disable features to prevent the attack from progressing. Many medical devices do not have this capability, but it is now an FDA requirement.

Software Configuration Management

Manufacturers are now required to create an interrogation process which reveals device configuration, e.g. device and firmware versions, to authorized users. In addition,

DEVICE CONFIGURATIONS MUST BE ACCESSIBLE TO AUTHORIZED USERS

interrogation must reveal a software bill of materials that itemizes all third-party software components utilized in the device and its system, including the name, origin, and version of

each software component. This is important because if a vulnerability is discovered in one of these third-party software components and disclosed, hospitals must have a method to quickly determine whether a given device is now known to be susceptible to attack and whether patients are at risk. This information will determine what action the hospital should take, including updating or replacing the device as necessary.

Incident Management

As previously touched upon, a new task set out for manufacturers revolves around ensuring each device should have the ability to make a Software Bill of Materials (SBoM) available in a machine-readable format. This should be independent of whatever communications systems are used. A user must be able to interface with the device in order to determine which operating system, libraries, and software components were used in development. Like the guidelines surrounding the configuration of a device, if a hospital or other authorized user wants to know if a device is vulnerable to an attack due to 3rd party software, they should be able to quickly obtain this information. This could take the form of something as simple as a comma separated file or as complex as some existing standards, such as SWID or SPDX.

The device also needs to be able to respond to and contain the impact of a potential cybersecurity incident. This may include resetting itself back into a secure state, and it should notify the user if this occurs.

There also exists a requirement for the rapid deployment of software patches. Should a device need to be updated for any reason, specifically to patch a known vulnerability, it is necessary that it be able to have both its software and firmware updated while in the field.

Finally, the device needs to be designed to recover capabilities or services that were impaired due to a cybersecurity incident. Below are a few examples of these types of incidents:

- Ransomware – an unauthorized user penetrating a device to encrypt part of the system to prevent access or function of aspects of the device
- Secure Configuration –engineers must ensure manipulating a device, e.g. increasing the dosage or changing therapy settings, can only be done by authorized users of appropriate privilege levels
- Autonomous Functionality – if a device uses feedback from the body to fine-tune its therapy, and the component reading the body malfunctions and can no longer provide this information, the device should revert to a safe mode of operation and continue functioning at a baseline level, regardless of which features are temporarily unavailable.
- Tolerance of Quality of Service – If a device is expecting to receive a packet every 100ms, but a delay occurs, how will the device handle this delay? Designs needs to be sufficiently resilient in the event that communications do not occur at the expected intervals.

The good news is that most manufacturers have already considered these scenarios as part of normal operation and expected interference, not necessarily with security in mind. Where appropriate, this type of functionality is may already be in place.

Cybersecurity Labeling

Labeling is any written text that identifies any part of a medical device. This includes a label directly on the device, user manuals,

instructions for use (IFU), labels present on the outside of packaging, or anything related

14 NEW AREAS OF COVERAGE FOR LABELING AND IFUS

to ad copy, such as claims pertaining to the device functions in an ad campaign. This post will break down the fourteen (14)

new requirements the [FDA has set forth](#) for Cybersecurity labeling procedures.

- 1) **“Device instructions and product specifications related to recommended cybersecurity controls appropriate for the intended use environment (e.g. anti-virus software, use of a firewall).”** This instruction has much more applicability towards a PC embedded in a medical device than any other system. If there is some sort of cybersecurity control, a manufacturer should designate specific sections for their documentation, which would include areas such as wireless and wired authentication and pairing, front door security with the use of a password or ID swipe, and ensuring there is no default password.
- 2) **“A description of the device features that protect critical functionality, even when the device’s cybersecurity has been compromised.”** It is important for the manufacturer to convey in a manner that the average user can understand how they intend to protect the execution of the code, present data, confidentiality of patient information, and secure updates.

- 3) **“A description of backup and restore features and procedures to regain configurations.”** Once again, a very PC centric guideline that doesn't really extend to custom hardware-based medical devices. In most cases, a simple power cycle will satisfy this requirement.
- 4) **“Specific guidance to users regarding supporting infrastructure requirements so that the device can operate as intended.”** This will be needed for devices running a separate segment, such as using ethernet, Wi-Fi, or BLE. Examples would include how an IP address is assigned, how the device handles connections to Wi-Fi and BLE, and if any specific ports that need to be open.
- 5) **“A description of how the device is or can be hardened using secure configuration. Secure configurations may include end point protections such as anti-malware, firewall/firewall rules, whitelisting, security event parameters, logging parameters, physical security detection.”** While similar to #2, here is where a list of security features needs to be itemized.
- 6) **“A list of network ports and other interfaces that are expected to receive and/or send data, and a description of port functionality and whether the ports are incoming or outgoing (note that unused ports should be disabled).”** Similar to the items in #4, listing out the network ports and interfaces that are sending and receiving data will be sufficient.
- 7) **“A description of systematic procedures for authorized users to download version-identifiable software and firmware from the manufacturer.”** All that is needed for this guideline is a description of how to perform an upgrade of firmware/software from the manufacturer for the user.
- 8) **“A description of how the design enables the device to announce when anomalous conditions are detected, (i.e., security events). Security event types could be configuration changes, network anomalies, login attempts, anomalous traffic (e.g., send request to unknown entities).”** Once again, this requirement is similar to #2. The system will need to go through and list out the events it is looking for and which it will announce to a user.
- 9) **“A description of how forensic evidence is captured, including but not limited to any log files kept for a security event. Log files descriptions should include how and where the log file is located, stored, recycled, archived, and how it could be consumed by automated analysis software (e.g. Intrusion Detection System, IDS).”** This guideline could not be more associated with PC oriented systems, but manufacturers will still have to cover forensic logs that are being maintained in embedded products. Both of these have to be covered and there are standards for interfacing analysis software if connected to a PC. The logs can and should be provided as protected forensic evidence after an attack has occurred.

10) **“A description of the methods for retention and recovery of device configuration by an authenticated privileged user.”** If these methods do not exist for a device, it needs to be noted in documentation. If present, they need to be able to store the configuration of the device over the communications protocol. A vast majority of devices will not need to have this functionality so it will not be a consistent requirement.

11) **“Sufficiently detailed system diagrams for end-users.”** Our second-least-favorite item on this list. A manufacturer will need to show devices in the system and the interconnections between them. This can be indicated via diagrams where security mitigations are implemented.

12) **“A CBoM including, but not limited to, a list of commercial, open source, and off-the-shelf software and hardware components to enable device users (including patients, providers, and healthcare delivery organizations (HDOs)) to effectively manage their assets, to understand the potential impact of identified vulnerabilities to the device (and the connected system), and to deploy countermeasures to maintain the device’s essential performance.”** Depending on the system, the software BoM will include all 3rd party software components (libraries, operating systems), and will also need to be machine readable. At this point, it appears hardware will not be included in the CBoM.

13) **“Where appropriate, technical instructions to permit secure network (connected) deployment and servicing, and instructions for**

users on how to respond upon detection of a cybersecurity vulnerability or incident.” This requirement covers user instructions for performing a secure update and returning a device to its secure configuration (e.g., power cycling).

14) **“Information, if known, concerning device cybersecurity end of support. At the end of support, a manufacturer may no longer be able to reasonably provide security patches or software updates. If the device remains in service following the end of support, the cybersecurity risks for end-users can be expected to increase over time.”** This phrase should be copied word-for-word into the device’s IFU, as it is reasonable to assume risks to end users will increase over time. There should be a concerted effort by the manufacturer to remove devices from the market and from use once support is cut off.

Cybersecurity Documentation Required for Submission

The first item needed is written documentation that addresses each of the design requirements in section 5 of the Guidance, if the device has been classified as “Tier 1”.

(If its classification is “Tier 2”, then a manufacturer can leverage risk-benefit assessments, such as those described in ISO

14971, in place of descriptions of the implemented mitigations for each of the standards found in section 5 as would normally be required if the device is a Tier 1.)

This documentation must be sufficiently detailed to permit understanding of how the section 5 elements are incorporated into the design. Manufacturers will be required to add more clarification than is requested for labeling, such as how are keys exchanged, what cryptographic encryption methods are being used (AES 128), what integrity controls are being used on the data (hashing, signed certificates, etc.) and so on.

Submissions must also include description of how the software updated on the device. This should also include specifics on how the data is encrypted, how integrity is ensured, how the device verifies that the incoming software update package is an upgrade and not a downgrade. We recommend referencing the Department of Commerce's NTIA guidance document on how to implement secure updates.

To avoid confusion, it's best we define a few terms:


- **Vulnerabilities** are weaknesses within a device and its code
- **Threats** are exploitations of a specific vulnerability
- **Exploitability** is a metric of how easy it is to take advantage of a vulnerability. Vulnerabilities can exist, but are they hidden beneath multiple layers of security?
- **Severity** is the scale at which we measure the result of a vulnerability being attacked. Can a patient be killed? Can the device's safety, efficacy, or performance be reduced?

This new guidance mandates that the metrics of exploitability and severity be included with each vulnerability. We recommend that manufacturers take advantage of the industry-standard CVSS (Common Vulnerability Scoring System) rubric, which includes these attributes plus a few more that deepen insight into a vulnerability.

We consider the AAMI TIR57 security risk management report to be the most actionable and applicable standard for risk management. When following this report as a guideline, a submission will include a security management report and a threat model, which will include supply chain and deployment. For supply chain, this would cover policies at the manufacturer level that prohibit components purchased from the so-called "grey market," limiting sources to only certified vendors.

In terms of deployment, how are the product artifacts, e.g. software and hardware, supplied to the manufacturing facility? Could they have been corrupted during the delivering process or during manufacturing? If you are using an internal manufacturing process or a contract manufacturer, how is your company ensuring the integrity of these components?

A list of all cybersecurity risks and a detailed methodology, such as the STRIDE decomposition methodology, including justifications and traceability to requirements, must be included as well in order to show how risks were considered.



**ALL POTENTIAL RISKS
MUST BE CONSIDERED
AND DOCUMENTED**

Furthermore, a list of all cybersecurity controls and their justifications and traceability to requirements need to be documented. These traceability requirements must have a list of vulnerabilities that are associated with the mitigations created to solve them, a written justification for why the mitigation is effective, and the traceability into the requirements documentation to show that the mitigation was actually a requirement and was implemented into the system. Moreover, both a code (SAST/DAST) and boundary analysis, as well as penetration testing, must be performed in order ensure the mitigations put in place are working as intended.

In order to tie back into all the cybersecurity controls pertaining to traceability, a traceability matrix will be required to connect all the document structures together. This should be an integral part of the artifacts that are created and traced through the development lifecycle.

In other words: Bring security into the mainstream of product development documentation.

Finally, a software SBoM needs to be cross-referenced to a vulnerability database, such as the NVD, as a manufacturer will want to show that currently there are no known vulnerabilities. Failing to include this risks a pre-market submission rejection.

Conclusion

In the end, the new Cybersecurity guidance put out by the FDA has changed the focus on secure device development in medical devices for the foreseeable future. With all these new requirements, embedded cybersecurity will now need to be a top priority within your product development cycle. As technology continues to adapt and change, it is reasonable to assume that guidelines like these will only continue to become stricter and increase in number.