The WHO Cybersecurity Risk Assessment

Quick Start

- Send an email to cs_red_team@who.int. A member of the team will reach out to you.
- Please have a good understanding of your business requirements. A rough draft is acceptable – they do not need to be in their final form.
- You should already be in contact with the IMT Business Relations Management (BRM) team. If not, please use their form to request a consultation.
- It is helpful, but not required, to have an IMT application code.

Motivation

Cybersecurity is the responsibility of everybody in an organization. We are all, collectively and individually, responsible for ensuring we use technology in the safest way possible. We are also responsible for ensuring 3rd parties we contract with use technology in the safest way possible.

When, at WHO, we engage in a digital product regardless of the size, scope, or importance, we take up this responsibility. We owe the Organization a duty of care and caution.

This is true when we create IT systems which are meant to process vital information. It is also true when we create systems that are meant to disseminate information to the public. It is especially true when we ask the public to trust us with *their* precious confidential and personal information.

Origin

The Office of Internal Oversight (IOS) conducted its first audit of IT systems and IT systems security in 2016-2017. One of the recommendations which IOS made following that audit was for the Cybersecurity Team (CST) within IMT (Information Management and Technology) to create this risk assessment.

In response to this recommendation the CST through 2018 and 2019, collected a list of WHO policies and staff rules pertaining to the recommendation and used them to construct the risk assessment procedure.

Objective

The objective of the cybersecurity risk assessment process is to protect information assets against unauthorized disclosure (confidentiality), modification (integrity), and destruction, and to help prevent attackers from denying access to critical systems (accessibility).

These goals merit different controls, depending on the risk-level of the application, its hosting model, and other criteria to provide the best level of security we can to the organization, our partners, and to the public. The cybersecurity risk assessment is meant to provide the sponsors and principals of digital products with a way to rapidly convert these goals into a concrete checklist of controls to implement to carry out the initiative as securely as possible.

A secondary, but also important, goal of the assessment is to prepare the team behind a digital product for audit. The CST, collectively have hundreds of hours of experience working with auditors,

from the IOS, and from entities like the EU (European Union) and other member states. The assessment gives us a platform to help application sponsors and their teams to successfully navigate all such audits.

Applicability

The IOS recommendation is – in the fullness of time – to perform the risk assessment for all WHO applications. However, since resources are limited, the CST help application sponsors perform it in the following situations:

- At the beginning of a new procurement process for a new digital product.
- When the business requirements or hosting model of a digital product change significantly.
- When a vulnerability with an application is observed in the system and reported to the CST.

Compliance

The role of CST is to help you adhere to these policies and document that implementation for an eventual audit. The controls themselves are based on Staff Rules, and on policies in the e-Manual, specifically paragraphs 42 and 43 of the e-Manual Cybersecurity policy.

Responsibilities

Business Product Owner:

Responsible for identifying the business processes executed by the technical system, determining its sensitivity and criticality, and all types of data processed, stored, or transferred by the system and. Responsible for the identification of a System Cybersecurity Risk Level.

PM (Project Manager):

Responsible for ensuring the assessment has taken place and for the implementation of security controls in the system. The assessment must be performed in Phase 1 (Initiating project phase based on the WHO Project Management Framework at link – IT variant).

Cybersecurity Team:

Primarily responsible for helping and guidance with risk assessments and penetration testing execution.

Process

The assessment has five stages:

1. Risk Classification
2. Application of the Cybersecurity Controls
3. Implementation
4. Documentation
5. Validation

Risk Classification

In the risk classification stage, the Business Product Owner of a digital product responds to seven questions concerning confidentiality, integrity, and availability of the digital product which they are sponsoring based on the business rules they propose.

The CST help establish a risk level for the digital product based on an analysis of the answers to these seven questions. The cybersecurity team may make suggestions at this stage about how the business requirements of the product might be adjusted to reduce the overall risk. These recommendations are non-binding but can often accelerate the realisation of a project or initiative.

Application of the Security Controls

Based on the risk level calculated in the classification stage, between 23 and about 45 controls are applied to a digital product. The controls are organized in sections drawn from ISO 27001:

1. Physical and Environmental Security
2. Asset Management
3. Access Control
4. Operational Security
5. Communications Security
6. System Acquisition, Development, and Maintenance
7. Supplier Relationships
8. Incident Management
9. Business Continuity
10. Compliance

The Cybersecurity team must ensure that the team executing a project understand each of the requirements which apply to their specific project, how to implement the requirement, and how to demonstrate to an auditor that the implementation is done. To help achieve this, a member of the cybersecurity team will meet with developers and other implementors in the project team to explain each control, its implementation, and its documentation requirements.

The controls are meant to be reasonable, and there is some room for interpretation. The Cybersecurity team member assigned to your risk assessment will guide you based on their experience with attacks and audit and will help prepare you for both.

Implementation

Once the team behind the digital product understand the controls, it is their task to implement any of them which has not been implemented in the original plan.

Documentation

After implementing the controls, the project team must produce documents to demonstrate to an auditor that each control has been implemented withing the limits of due diligence.

Validation

The project team must then meet with the CST and with the IT Architect to validate its documentation. This is intended to help avoid unhappy surprises when the project rolls out to production. Following this meeting, the project is free to deploy to production. However, any announcement must be held until the CST conduct an automated scan of low-risk projects to reveal any vulnerabilities which might be present. For medium and higher risk projects the team will conduct manual penetration testing.

If severe or significant vulnerabilities are found during automated scanning the Cybersecurity team will conduct manual penetration testing, per the prerogative of the Chief Information Security officer.

Vulnerabilities which emerge from this scanning or testing must be remediated through bug-fixing or through a revision of the business requirements of the project.

Once all severe and significant vulnerabilities are closed the product can be released.

The stages of the assessment may seem daunting, but they are intended to be conducted rapidly. The estimated time for each part of the assessment is:

1. Classification – 10 minutes
2. Controls read-through – 30 to 45 minutes
3. Implementation – should normally already be done
4. Document Production – no more than 2 hours
5. Validation meeting – 30 minutes

The burden of scanning falls upon the CST, however remediation remains the responsibility of the project team.

CST, may, as a prerogative of the Chief Information Security Officer (CISO) assist with vulnerability remediation and any attendant mitigation, however this remains the responsibility of the project sponsor. The time required for this is wildly variable depending on the quality of the project code and interests and aptitudes of its developers.

RASCI chart

R – Responsible
A – Accountable (to the organization)
S – Supports
C – Consults
I – informed

Task	Project owner	Project manager	Developer or Integrator	Cybersecurity (CST)	Architecture (AIS)	Audit (IOS)
Arrange meetings		A				
Classification	A	R		S		I
Controls Read-through	A	R	R	S		
Implementation	A	R	R	C	C	
Documentation	A	R	R			
Validation	A	R	R	S	S	I

Timeline

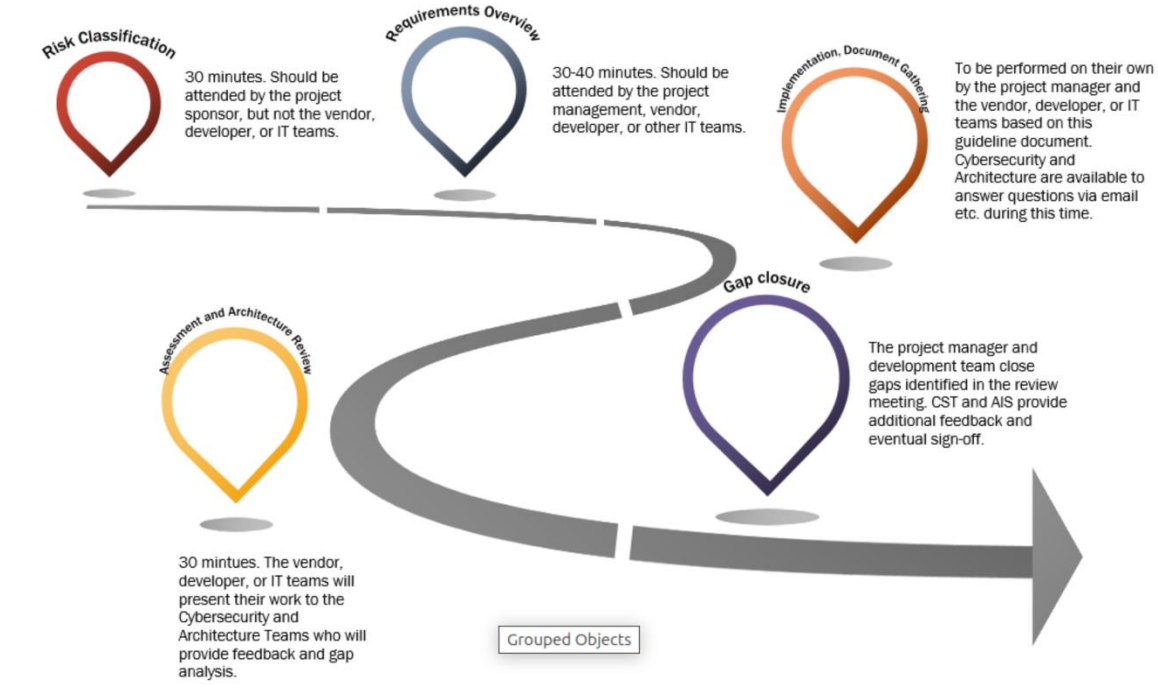

Annex I – Classification Questions

What best describes the most sensitive information handled, processed, or stored by the system(s)?

- Confidential (e.g., financial records, medical records, passwords, personally identifiable information, phone numbers, addresses, passport numbers, etc)
- Information for internal use only
- Public documents or open information

These three levels of classification are defined by the organization outside of the Cybersecurity risk assessment and apply to all information which the organization handles in digital or print form.

What is the worst outcome if that information is leaked?

- High Impact (e.g., loss of lives, Involvement of DG and Member State commissions, etc.)
- Medium Impact (e.g., damage to the credibility and/or Involvement of DCO and/or LEG)
- Low Impact (e.g., extra effort or higher cost, disruption in operations)
- No Impact (e.g., no extra effort, no disruption in operations)

Keep in mind that information which can be used to accurately geolocate a person can be extremely dangerous if the person is involved in some way in a real conflict.

What is the worst outcome if an attacker tampers with that information?

- High Impact (e.g., loss of lives, involvement of DG and Member States commissions, etc.)
- Medium Impact (e.g., damage to the credibility and/or involvement of DCO and/or LEG)
- Low Impact (e.g., extra effort or higher cost, disruption in operations)
- No Impact (e.g., no extra effort, no disruption in operations)

This is usually the hardest question, since we must think hard about what an attacker might have to gain from tampering with the data in the system, and what effect that might have on the organization. The CST have observed several variations on this kind of attack, ranging from defacement, through using WHO systems as an advertising platform, all the way through falsifying of data with the apparent aim of moving markets and creating illicit outsider/insider quick gain opportunities.

Who will use the system?

- A restricted group of WHO staff
- WHO staff and others under contract with WHO
- WHO staff, contractors, and a restricted group of an external registered third party/users
- The public

A system with sensitive information which is available to the public has a different security profile from one which is not available to the public, for example. So, we need to take the audience into account.

Is the system business-critical?

- Yes, it effects the whole of WHO
- Yes, it effects a specific office, team, unit or division
- No, it is not critical

For the purposes of this question, we need to carefully consider criticality. A system is critical if the organization or some part of it rely on the system to do their day-to-day work. Otherwise, it is not critical no matter how important.

If the system must be taken offline after an attack, for how long can it be down?

- It is critical and must be available all (100%) of the time
- Up to 4 hours of downtime/unavailability can be tolerated
- Up to 24 hours of downtime/unavailability can be tolerated
- Up to 1 week of downtime/unavailability can be tolerated
- More than 1 week of downtime/unavailability can be tolerated

During a cyberattack often the simplest way to deny access to an attacker is to simply take the system offline. The Security Operations Centre analysts collaborate with the product team to decide if this is possible. The answer to this question should be the one you imagine making during that meeting. It is not an SLA.

If the system must be restored from backup, what is the oldest acceptable backup to restore?

- The system is critical and the loss of new data or any recent changes could not be tolerated
- The loss of new or changed data from the last 24 hours could be tolerated
- The loss of new or changed data from the last 7 days could be tolerated
- The loss of new or changed data from the last week or more could be tolerated

If we believe that an attacker has tampered with the data in the system, the most prudent course of action might be to restore the system from backup. Assuming perfect backups we are interested in knowing how far back in time we can safely go. This should be based primarily on the volume of transactions in the system.

Annex II – Hosting Models

Software as a Service (SaaS)

Software as a Service is a model in which all responsibility for security and maintenance is assumed by the provider. This model is the easiest and fastest to implement, and many of our controls do not apply, however we still need documentation that applicable security controls have been implemented at the provider.

Platform as a Service (PaaS)

Platform as a Service is a model in which the cloud provider provides services like a database service like SQL Server or MySQL, a framework service, like .NET or PHP without need for the project team or the WHO hosting team to install these services on a machine, virtual or not. The cloud provider takes responsibility for some security measures like operating systems updates and patching, so those controls do not apply.

Infrastructure as a Service (IaaS)

Infrastructure as a Service is a model in which the project team and their suppliers run virtual machines in the cloud environment. In this model the project team is responsible for every security control except for the remote access control.

On-Premises Hosting

When hosting on the WHO premises all controls apply, however the hosting team will ensure that most of them are satisfied.

Client-Side Only (flat HTML + JavaScript)

Flat-file websites are naturally immune to many of the kinds of attacks which we worry about, but not all of them. Specifically, a flat HTML and JavaScript website can be vulnerable to cross-site-scripting (XSS) and document object model (DOM) attacks.

It is entirely possible to integrate such a site with SSO should authentication be required, but it must be rigorously checked, since any resulting secrets must be stored on the client.

Mobile-Only

Most mobile applications rely upon and have some connection with an online data source. Many even write data back to that data source, making it a data sink in cybersecurity terms.

However, some mobile applications do not rely on online data, or if they do on a data source which already exists and has already been vetted through this process. For those mobile applications many of the controls in this risk assessment process do not apply directly, especially if users do not authenticate to the application in any way

Annex III – Requirements.

In the second stage of the risk assessment, we apply requirements based on the Staff Rules, policies in the WHO e-Manual, and various audit recommendations to each project. The project owner is responsible for the application of these requirements, and the Office of Internal Oversight are charged with policing them.

The role of CST in this process is to explain the reasoning behind the requirements, and their goals, and to advocate for the best application of the requirements to help a project team to be successful in meeting their goals as securely as possible, as well as being compliant with all known audit requirements.

This will save time for both the project team and auditors in the event of an audit.

We have organized the requirements according to sections drawn from ISO 27001:

11. Physical and Environmental Security
12. Asset Management
13. Access Control
14. Operational Security
15. Communications Security
16. System Acquisition, Development, and Maintenance
17. Supplier Relationships
18. Incident Management
19. Business Continuity
20. Compliance

For each requirement in this annex, we explain the rationale behind the requirement, and explain how project teams can implement the requirement. We also explain how to demonstrate to auditors that the requirement has been implemented, based on our experience, and advice from the Office of Internal Oversight.

Each requirement has a basis either in the Staff Rules, in the e-Manual, a specific audit recommendation, or industry best practice.

The Cybersecurity Risk Assessment tool provides a rapid guide to applicability of these requirement, and a basic document for tracking your project. Once non-applicable requirements are filtered out you can use it as a to-do list for security and compliance.

11.1 Data Centre

Requirement	Applications or systems must be hosted either in the HQ data centre, one of the regional office data centres, or in a third-party datacentre that has ISO/IEC 27001 certification or the equivalent.
Section	Physical and Environmental Security
Risk Level	All
Hosting Models	All
Policy	WHO e-Manual XIV.3.1 Cybersecurity Policy

Rationale:

Cybersecurity starts with physical security, ranging from the physical locks on data centre doors to HR (Human Resources) policies through things like environmental controls like air conditioning and even neatness policy. Standards like ISO 27001, SOC 2, and EN 50600 provide a way for data centre owners to demonstrate that they have solid policies and that they follow them.

Implementation:

Most cloud hosting providers have one of these certifications. You should check before signing a hosting agreement.

Documentation:

The auditor should be able to easily find the hosting provider's certification. For WHO, Azure, AWS (Amazon Web Services) and Google Cloud Services the auditor already has these on record. If you are using a different hosting provider from these you should provide the certification documents, or links to them in your risk assessment's documents folder.

Applicability:

This control is always applicable.

12.1 Labelling

Requirement	The application or system owners must ensure that the application supports labelling for information classification. This is to help protect WHO in case confidential or internal-use information is printed or appears in a screenshot.
Section	Asset Management
Risk Level	All
Hosting Models	All
Policy	WHO XIV.2.3 Information Classification Policy

Rationale:

Carefully applied permissions and user/role schemes *within* an application often do a decent job of protecting sensitive information and are always implemented. However, data can escape from your application, and will escape when users print or take screenshots.

Implementation:

This is an extremely easy requirement to implement. If the application contains data which is either confidential or for internal use only the UI should contain some text which will *hopefully* appear in a screenshot or especially in a printout.

Documentation:

The auditor should be able to easily find a section in one of your documents describing how this has been implemented, with evidence of the implementation.

Applicability:

Obviously if the application handles only public data, or if it has no UI, as in the case of an API (Applications and programming interfaces), then this is not applicable directly. API developers should, however, consider the needs of downstream applications which might have a UI.

12.2 Removable Media

Requirement	Removable media must be protected with cryptographic techniques, preventing data degradation when transferring to fresh media and making content unrecoverable when it is no longer needed.
Section	Asset Management
Risk Level	All
Hosting Models	Physical devices, Projects relying on removable media
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Removable media, like USB keys and even internal hard-drives and SSDs present a serious confidentiality risk, in that it is easy for them to be misplaced or stolen, and then to fall into the wrong hands. We consider internal hard drives to be removable because they usually are, if an attacker has a few minutes alone with your device, and the correct tools.

Encrypting all such media makes it much more difficult for an attacker to use or copy the contents.

Implementation:

WHO-issued hardware, like Synergy desktops, and laptops, and iPhone mobile devices have this feature enabled for internal drives, so it only need be done for removable drives. For every removable drive used in a project encryption must be enabled, and a key management scheme must be implemented.

Documentation:

Describe your key management plan, and device encryption scheme, and provide examples of the result as proof. Screenshots showing the use of encryption tools like BitLocker and Cryptomator can demonstrate this for the auditor.

Applicability:

Clearly this only applies directly to projects involving such hardware.

12.3 Decommissioning

Requirement	The application or system owners must ensure that assets are destroyed or erased in a secure manner when they are to be decommissioned. Domain names must be safely parked to avoid abuse.
Section	Asset Management
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Teams rarely consider the possibility of a project reaching end-of-life while they are still in the planning stages, yet this is exactly what needs to be done to make sure that systems are retired rather than being abandoned. Abandoned projects create several distinct kinds of vulnerabilities, from a failure to patch upstream sources when vulnerabilities are discovered, to the potential loss or degradation of permission structures, as the organization structure changes over time.

Implementation:

Plan for end-of-life of the product. Try to define the kinds of events which should trigger this event, and who needs to be informed and consulted. Specify which data must be handled carefully, and how to delete it securely – consider your backups, keys, etc.

Documentation:

Write the plan down in a clear and concise set of events and decisions. A flowchart can sometimes be helpful.

Applicability:

This is applicable everywhere, but especially when sensitive information is involved.

12.4 Documentation

Requirement	The application or system owners must ensure that all technical documentation including architecture diagrams are updated and available. The information must include details of physical/virtual servers, operating systems, web servers, web frameworks, applications, databases, and network infrastructure.
Section	Asset Management
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Having good documentation is important at every stage of a project, and so normally these documents will be produced at an early stage in the project and kept up to date for the internal use of the project team. We ask for these documents for three audiences and three distinct reasons.

1. For audit. An auditor will look for these documents, and so having them is crucial at audit time.
2. For compatibility with WHO hosting and systems. PaaS, IaaS, and internally hosted projects must follow the guidelines of the infrastructure and hosting team.
3. For the Security Operations Centre during and after a cybersecurity incident. An agent must be able to quickly find and read these documents to quickly understand the layout of the system, and how an attacker might understand it.

This kind of document will also usually enumerate the version numbers of upstream dependencies, etc. This will help us make sure that we start with a sound and secure system architecture.

Implementation:

We ask that the documents provide all relevant information, but that they be terse, and not overly wordy. We also require a clear, and accurate infrastructure architecture diagram, with references to all the kinds of tools mentioned in the requirement text, and anything else of interest.

Documentation:

You can upload copies of the documents to the “Documents” folder within your risk assessment folder, or you can create hyperlinks to documents on a vendor site.

Applicability:

This varies a bit with hosting models. For On-Premises, IaaS, and Paas, we ask for the most complete documentation possible.

For SaaS solutions the second and third reasons for this documentation are not applicable, but we still need to demonstrate that this documentation exists for audit purposes. Happily, most SaaS providers have ample documentation on their websites. Please provide links, especially for any document which satisfies some other requirement in this list.

12.5 CMDB

Requirement	The Application or system owners must ensure that all related IT assets are documented and wherever possible available within a Configuration Management Database (CMDB).
Section	Asset Management
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Having a CMDB is helpful in a lot of ways, but for cybersecurity we are interested in the ability of a CMDB to show if a change in the application configuration has happened without your knowledge. This protects the application from changes made by an attacker as well as hasty production-only changes which skip change management.

Implementation:

Unfortunately, WHO has no CMDB capability at the time of this writing, so we cannot offer you this facility. However, cloud hosting models have several features which satisfy the same requirement. Namely infrastructure-as-code in the form of a templating system like Terraform or the templating systems native to your cloud environment.

You can also satisfy this with a development pipeline like CI/CD or Azure Dev-Ops.

Documentation:

A mention of the method used to capture a last known good configuration should appear in your documents with screenshots demonstrating the use of that method.

Applicability:

This is always applicable.

For SaaS applications it might be difficult to demonstrate, but it might be covered in certifications. You may find reaching out to our internal audit helpful to understand what they will accept.

13.1 SSO

Requirement	The application or system owners must ensure that authentication in integrated with the WHO Single Sign-On (SSO) system in Azure.
Section	Access Control
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Registering and authenticating users locally in an application exposes the organization, the users, and the application owner to a bevy of risks, and must be avoided. For this reason, WHO, like most organizations provides a central single-sign-on system.

In the absence of single-sign-on application owners are left managing users locally. This means that it is *possible* during the registration and authentication processes for an attacker who has gained access to the application to obtain the users' credentials. Sometimes even a good faith actor, like a developer debugging an application can obtain users credentials.

Having even the *possibility of* exposure to user credentials is a very compromising situation for an application owner or developer if not least because a significant percentage of users *reuse their password* in multiple systems. Having user passwords means that such a user might believe that WHO has mishandled their password if some other account of theirs becomes compromised in around the same period as their registration.

SSO prevents this kind of scenario, as well as real attacks against credentials, and allows us to apply policies which protect your users.

Implementation:

Integration can be accomplished with SAML2 or OpenID-Connect, both of which are industry standards. There are no peculiarities.

We do also help as required up to and including help with your code. Our SSO system also generates sample code for bespoke applications.

Documentation:

The architectural diagram should mention Azure AD (Entra ID) integration. Additionally, your SOPs (Standard Operating Procedure), and administrator's manual should include information about external partner registration, if required.

Applicability:

This is applicable to any application which authenticates users or requires any form of user management. It is an *especially* important and hard requirement for WHO workforce users.

The CISO has granted a few exceptions for *external* users in cases more than a million or so external users are expected, for example with WHO Academy.

13.2 MFA

Requirement	The application or system must enforce multi-factor authentication (MFA) for all users.
Section	Access Control
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Passwords are vulnerable to attacks like phishing. Worse, as noted above a certain number of users tend to use very weak passwords, or the same password everywhere. Even though we must try to guide users to use the best passwords possible it is never enough.

Using multifactor authentication, with TOTP or better methods like FIDO2 helps protect users from themselves and attackers.

Implementation:

Integration with the WHO single-sign-on system achieves this goal with no further effort, so if your users are WHO staff, or even a small number of external users (fewer than one million) then single-sign-on is the best way to implement this.

Otherwise, we recommend using FIDO2 keys with either a portable phone authentication app, or physical keys like Yubikey.

Documentation:

If you have implemented SSO there is no need for further documentation. However, in the exceptional situation of having many external users then you need to document on-boarding policies etc, for your external users.

Applicability:

While this requirement is always applicable it becomes the responsibility of WHO when using SSO (see 3.1 SSO)

13.3 Secondary Accounts

Requirement	Application or system owners must exclusively use secondary accounts for situations requiring elevating privileges. Secondary accounts must be removed when no longer required.
Section	Access Control
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Day to day accounts used for regular productivity activities like email are especially vulnerable to phishing and other confidence-based attacks.

Using a secondary account for activities like service administration and architecture adds an additional layer of complexity for an attacker, and can help at least mislead the attacker, preventing some kinds of attacks.

This can also help defend against lateral movement by an attacker who has gained a foothold by compromising a regular user account.

Implementation:

File a request with the global service desk for either on-premises “privileged accounts” or cloud-only Azure admin accounts as required. Grant elevated privileges to those accounts only.

Documentation:

Create a plan to stick to the use of secondary accounts for privileged access, and document it in your Admin Guide document. Include screenshots to demonstrate that you are following your procedure.

Applicability:

Applicable in any situation in which elevated privileges will be used to manage servers and other resources.

13.4 Delegation

Requirement	Application or system owners must ensure that login credentials (usernames/passwords) are not shared with anyone. If security delegation is required then this must be done by a traceable, auditable mechanism.
Section	Access Control
Risk Level	MEDIUM+
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Password sharing nearly guarantees password leakage, and more importantly makes auditability of actions impossible. For these reasons WHO staff are forbidden from sharing passwords in our Staff Rules.

However, when working with applications which have approval workflows units are placed in a position in which the temptation to share passwords is strong, since an approval might be with a person who is away on leave or duty travel, or otherwise unable to complete the approval.

This can be solved by functionality which allows a user or her supervisor to delegate that user's approval processes to some other user for a period.

Implementation:

This should be implemented through an explicit delegation feature. Role-based delegation can be acceptable.

Documentation:

The Admin Guide should contain a strong admonishment to not share passwords, and an explanation of how to use the delegation feature.

Applicability:

Any system with workflow and approvals must implement this. It is not applicable in the absence of workflow which can block in the absence of a given user.

13.5 Periodic Review of Access

Requirement	Application or system owners must ensure that all access permissions are reviewed periodically to determine that only the correct people have access to given roles and resources to avoid a scenario of toxic combinations.
Section	Access Control
Risk Level	All
Hosting Models	MEDIUM+
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Roles and role assignments are not static. Their memberships change, and sometimes the meanings of roles change. Roles are added and removed. New people join the team and others leave the team.

So, from time-to-time role assignments must be checked to make sure that only the right people have a given role. Our auditor recommends that this be done at least twice per year for most applications, but more frequently for applications with business-critical roles. For financial applications this should be done at least monthly.

Implementation:

Role review can be implemented within an application by adding a certification function. In certification the system creates a workflow from time-to-time for each role assignment, which must then be re-approved. If the workflow is rejected or left incomplete for a certain time the role is removed from the user concerned.

Another way to implement role review is offline, with a report created from application data, and then the reports used to manually search for incorrect role assignments. Roles would then be unassigned by manual procedures, as per the application.

Documentation:

Automatic certification should be auditable. The auditor should be provided with a method for generating a report on certification in the application. Documentation of this method should suffice to prove that the method exists.

In case of manual role review you should upload the artifacts from the review each time it takes place, or at least provide a link to the location where those artifacts are stored.

Applicability:

All medium risk and higher applications must implement role review.

13.6 Central Authorization Control

Requirement	Access rights within the application must be governed centrally. This can be fulfilled in several different ways, including the use of Entra ID application roles or security groups, on-premises AD security groups, or role-based provisioning from Oracle Identity Manager or Service Now.
Section	Access Control
Risk Level	HIGH+
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Just knowing who a person is (Authentication) might be enough to protect access to an application which has only one role, and those exist, but most applications have more than one role.

Knowing what role an authenticated person has in an application helps us make sure that actions are taken according to the rules and the procedures of the organization.

Centralized assignment of the roles and responsibilities allows systems to assign them automatically more easily, which leads to fewer exceptions. It also makes it easier for auditors to extract reports. It also helps us make sure that application roles are mapped as well as possible to real world roles.

Implementation:

Role and role membership date can be consumed from central systems using several mechanisms:

1. Direct provisioning from the Identity Governance System
2. Ahead-of-time provisioning from the authentication system (Azure AD)
3. Just-in-time provisioning from the authentication system during authentication

Since any of these options require special integration which depends on the application, you will need to work with the Cybersecurity team and our colleagues in Workplace and Communications to perform the integration.

Documentation:

The auditors will want to understand which security groups and application roles control access to various functionality in the application, so you should document that connection. You should also document how they can produce a report on demand.

Applicability:

This applies to all High-risk and VERY High-risk applications and is especially important where financial transactions are involved in some way.

13.7 The principle of least privilege

Requirement	Application or system owners must ensure that access permissions are applied based on the principle of least privilege. Users must be provided the minimum level of access or permission that is required to perform their role.
Section	Access Control
Risk Level	HIGH+
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Users should be given, at any moment, the permissions they need to do the tasks required of them and exactly no more permissions. Higher privileges should be granted when a user needs them, and then taken away again when the user no longer needs them.

This protects the organization, the system, and the user against a situation in which an attacker gains at least some level of control over the user's account. If the user has, at that moment, access to limited functionality then the attacker also will have access to only that functionality.

This also protects users from inadvertent access to functions upon which they are not trained, or likewise to data which they might naively mishandle.

Implementation:

Define processes for assigning privileges, and ways for a user to gain access to more restricted privileges when they need them, and then to give them up when they are no longer needed.

Documentation:

Simply share the documents produced in the implementation.

Applicability:

All high-risk and very-high-risk applications must implement this control, but it should be implemented everywhere.

13.8 PAM

Requirement	Administrator level access must be protected in such a way that each new access event requires the knowledge and an approval of a superior officer and is logged and audited. Technologies supporting this type of access are known as Privileged Access Management (PAM).
Section	Access Control
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Privileged Access Management (PAM) systems automate the process of temporarily granting advanced privileges in a system. They therefore reduce the possibility for human error, especially in terms of lingering privileges.

PAM systems can make administration of easier for the user/administrator as well, since the workflow to temporarily gain access can be automated and distributed, including peer-level grants.

Finally, PAM systems normally feature robust logging of actions taken during the time advanced privileges are being used.

Implementation:

There are many choices on the market, with proprietary, Open Source, and Open Core business models. The application team are advised to choose a tool which matches their needs and budget. The cybersecurity team can provide advice.

Documentation:

Documentation of your PAM implementation should be available to the auditor, including instructions about how to extract reports.

Applicability:

This is applicable to all high-risk and very-high-risk projects.

14.1 Operating Systems

Requirement	Application or system owners must ensure that the approved WHO operating system images for servers and virtual machines are used.
Section	Operations Security
Risk Level	All
Hosting Models	On-Premises, IaaS
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

The Architecture and Infrastructure (AIS) team prepare operating system images for those operating systems which they support, and they keep those images up to date.

Virtual and physical machines which are deployed from those images will automatically implement many of the controls in this section, and the following section, and will have other characteristics which make them easier for members of the AIS team to support. Use of these operating system images will cut down on surprises all around.

Implementation:

In most instances a member of AIS or the regional will provision a system for you to work with. Otherwise, AIS can provide installation media for you, or recommend a cloud installation image.

Documentation:

Please include the operating system major and minor versions in your architecture diagram, normally in a bullet point on basic infrastructure.

Applicability:

Since SaaS and PaaS hosting models do not expose the operating system to the project team the requirement obviously cannot be applied. Otherwise, it is always applicable.

However, AIS and the auditors have shown some flexibility in situations where an installation can be considered something like an *appliance* which we can think of as a kind of SaaS-on-premises hybrid. It is important to make sure that such a situation is fully supported by your vendor and that contracts cover all operating systems maintenance issues, and patching, as in SaaS, but with the additional burden of accessing an environment provided by or through WHO.

14.2 Outbound Proxy

Requirement	The system owner must ensure that the WHO Proxy is enforced on the underlying infrastructure. Direct access to the internet must not be permitted.
Section	Operations Security
Risk Level	All
Hosting Models	On-premises, IaaS, PaaS
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

When an attacker investigates what they believe to be a vulnerability, one of the most useful techniques at their disposal is to try to get the system under attack to try to contact an attacker-controlled system over the public Internet.

Successful outbound contact of this type can prove to the attacker that a vulnerability exists and is worth investigating.

At worst, an attacker can create a type of connection called a *reverse shell,* which will give the attacker full control over the target system.

Implementation:

Block all outbound traffic from the system via Firewall and configure the WHO outbound proxy server for those outbound connections which are required. Ask your WHO Hosting team about the outbound proxy settings.

Documentation:

Demonstrate that all outbound traffic is routed via the provided HTTPS proxy using PCAP examples or screenshots.

Applicability:

It is difficult to imagine applying or proving the application of this control in a SaaS situation. However, a reasonable attempt should be made to demonstrate due diligence.

14.3 Logging

Requirement	The application must generate log files about all activities happening within application (e.g., logon/logoff events, access right modification, permissions modification, business critical transactions)
Section	Operations Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Having a trace of who has logged in to an application, and what they have done there is indispensable in diagnosing things that have gone wrong, whether the access is malicious. In the case of malicious access then these logs become part of the key forensic evidence which will help our analysts understand the attacker's behaviour and goals.

Implementation:

If you have implemented SSO (control 3.1) then the only thing left is to be sure to log out important transactions. If there are no important transactions, then the requirement is fulfilled by SSO.

However, if the normal operation of the application includes important transactions, then they must be logged, financial transactions are an obvious example. Other types of events which should be logged include privilege escalation and role assignment.

Documentation:

Demonstrate that the application is integrated with the WHO SSO system, and if possible, catalogue those transactions which might be critical.

Applicability:

This is always applicable.

14.4 SIEM

Requirement	The application owner should ensure that all critical application event logs, administrator, and system operator activities, recording user activities, exceptions, faults, and information security events are monitored and logged in tamper-proof mode. Critical logs must be replicated to a central WHO logging solution.
Section	Operations Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Logging covered in item 4.3 is necessary, but not sufficient to provide cybersecurity forensic services for an application. The logs must also be stored in an unalterable way. This is sometimes referred to as log *shipping.*

Attackers will often attempt to cover their tracks by deleting or altering log files. If the logs are shipped to a central log server, this improves the chances of the attacker not being able to alter them.

One log shipping target is our Security information and event management (SIEM) system. SIEM goes beyond simple centralized logging to provide search and event hunting capabilities to our Security Operations Centre.

Implementation:

Our SIEM is LogRythm. Integration can be done with it directly or via a GNU Syslog server, and GreyLog.

Documentation:

Your architecture documents should provide proof of SIEM integration for the auditor, with screenshots if possible.

Applicability:

This is applicable to any application with critical transactions or events.

14.5 Product Manuals

Requirement	The application or system owner must ensure that installation and operating procedures are documented and available to all users that require them. (Administration manual and SOP (Standard Operating Procedure) for users)
Section	Operations Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Well document installation procedures help prevent a situation in which the original project team move on to some other work, and the new team need to know how to create an environment from scratch. This might be for creating a new test instance, or it might be for rebuilding a production instance in an emergency.

Complete operating procedures can help identify the correct behaviour of an application, so that if its behaviour changes due to an attack, we have a better chance of detecting the change. It can also help prevent the accumulation of bad data by improper use of the application.

Implementation:

Please create and test these documents and procedures.

Documentation:

Simply upload a copy of the completed documentation to the Documents folder in your risk assessment folder, for the auditor.

Applicability:

This is always applicable.

14.6 Vulnerability Scans

Requirement	The application or system owner must ensure that the application and its underlying infrastructure undergo regular vulnerability scanning.
Section	Operations Security
Risk Level	All
Hosting Models	HIGH+
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

New vulnerabilities can be introduced in the codebase of an application in several ways: the development of new features, regressions caused by fixing other bugs, or even in an application which itself has not change, by vulnerabilities being discovered in libraries or other dependencies.

For this reason, it is useful to perform a black box, dynamic application security test (DAST) from time to time.

Implementation:

Identify a DAST tool either within your existing development environment or ask for help from the WHO Cybersecurity team to identity a tool you can use. Build a schedule for running the tool as appropriate to your project, and make sure it is followed. The Cybersecurity team can also provide expert advice on remediation of any vulnerabilities you might find during these tests.

Documentation:

The auditor will want to see documentation of your testing strategy and schedule, so those should be included, in your architecture documents. As tests are run, in production, the test results should also be stored in the risk assessment Documents folder, as well as documentation of the remediation of any vulnerabilities you find during these tests.

Applicability:

Applicable to all high-risk or very-high-risk applications.

14.7 Patching

Requirement	The application or system must be compliant with the WHO patch management process to address security vulnerabilities.
Section	Operations Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

New vulnerabilities are discovered daily in operating systems, in libraries, frameworks, and all other kinds of upstream dependencies. Presuming your upstream dependencies are within the vendor's support model, the vendor will release a patch.

It is easy to find unpatched versions of these dependencies with a simple web scan, or by using a specialised search engine, or sometimes even with Google. Failing to patch would not only leave you vulnerable, but it also actually makes you a target.

This concept extends to upgrades as well, since at some point vendor support is dropped for older major versions of these dependencies, and they stop receiving patches, even when new vulnerabilities are discovered.

Implementation:

Devise a patching and release strategy for your project and stick to it.

Documentation:

Document your project's patching and release strategy, and then provide logs as you perform patching cycles so the auditor can be sure that you are following your strategy.

Applicability:

This is always applicable. For systems with on-premises hosting the responsibility may be divided between the project owner and a WHO hosting team.

14.8 EDR

Requirement	All systems must be equipped with the global WHO Antimalware/Endpoint Detection Response solution, where applicable. If not, detection, prevention, and recovery controls to protect against cyberattacks must be implemented.
Section	Operations Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Endpoint detection and response (EDR) consists of a light agent which runs in the background of your operating system. This agent observes system activity for telltale signs of intruder-initiated events, which is to say that somebody other than you is running scripts or programs in the background.

The agent reports this information to the Security Operations Center. If an analyst believes that your system is in fact under attack, they will notify you, and the Cybersecurity team, and with your consent take appropriate action.

Implementation:

Obtain an EDR installation package from the Cybersecurity team (if it has not already been installed).

Documentation:

Note in your document the EDR agent version.

Applicability:

On premises and Infrastructure as a Service (IaaS) systems must run EDR. This is not applicable to Software as a Service (SaaS) or Platform as a Service (PaaS)

14.9 Scanning of Uploaded Files

Requirement	All applications that require file upload functionality must scan all uploaded files for viruses and malware (See OWASP recommendations).
Section	Operations Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Many file types, including Excel and PDF files allow for executable scripts to be embedded within them. Other file types, even images, can be coerced into containing scripts and programs. For example, with careful crafting by an attacker a valid JPEG image can simultaneously be a valid Java or PHP library, or even a .NET installer.

When these files are uploaded by users from outside of the organization, they can be used to attack your system, or the desktop and laptop computers of your users.

https://owasp.org/www-community/vulnerabilities/Unrestricted_File_Upload

Implementation:

In the best-case files should be scanned before they are placed in storage, otherwise there will be a race between the scanner and any malicious code in the uploaded file. The project team is free to choose an implementing method. One option is to use a solution provided by WHO, the Cybersecurity team will provide guidance.

Documentation:

Whatever scanning method you implement, you should document it in your architecture document.

Applicability:

This is applicable to any application which allows *end-users to* upload files from endpoints outside of the control of the organization.

14.10 Limiting File Types to be Uploaded

Requirement	The file types allowed to be uploaded should be restricted to only those that are necessary for business functionality.
Section	Operations Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Given the dangers of accepting the upload of untrusted files, it is best to limit file types to the absolute minimum. This reduces the attack surface. If you allow only image files, then you limit your exposure to only the most difficult class of attacks.

If on the other hand you allow all file types, then you and your users will have to deal with executable files, installers, and other risky file types.

https://owasp.org/www-community/vulnerabilities/Unrestricted_File_Upload#

Implementation:

Configure your application to prevent the upload of files outside of a narrowly defined set of filetypes.

Documentation:

Please document the list of filetypes which your application allows and explain why each type is needed.

Applicability:

This is applicable to any application which allows *end-users* to upload files from endpoints outside of the control of the organization.

15.1 Insecure protocols

Requirement	The application or system must not use insecure TCP/UDP protocols such as (but not limited to) HTTP, FTP, Telnet or TFTP to transfer data over all wired and wireless networks. Only encrypted traffic is allowed.
Section	Communications Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

The protocols we use in most applications we use today provide encryption during their authentication phase, or offload authentication to a Single-Sign-On provider. Older protocols like FTP and Telnet, however, send plaintext credentials over the network during their authentication phase.

Such protocols must never be used and should be considered obsolete.

Implementation:

Replace Telnet with SSH (Secure Shell), and FTP with SFTP (Secure FTP) (at least) or consider replacing both with HTTPS.

Documentation:

Please make sure that your architecture diagram and document specifies which network communication protocols are used everywhere in your architecture, and that it is complete and exhaustive.

Applicability:

This is always applicable

15.2 Infrastructure SIEM

Requirement	All network infrastructure such as (Firewalls, Routers, Switches, Virtual Servers) must be configured to send security events to the Cybersecurity security operations centre log (SIEM) solution.
Section	Communications Security
Risk Level	HIGH+
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Infrastructure components produce logs which can be valuable for an analyst hunting for attacker behaviours.

They can also be a source of a large amount of noise. For this reason, we only ask that these logs be sent to the SIEM solution for high-risk and very-high-risk applications.

Implementation:

The cybersecurity team will provide guidance on SIEM integration for your infrastructure components.

Documentation:

Please include SIEM integration in your architecture document and diagrams.

Applicability:

High-risk and very-high-risk applications which run on premises, in Infrastructure as a Service, or Platform as a Service. Software as a Service vendors must have their own SIEM.

15.3 Infrastructure authentication

Requirement	Administrative access to all network infrastructure (Firewalls, Routers, Switches, Virtual Servers) must be integrated with the WHO global identity solution and must not use any local identity services. (Local administrative access must only be used in the event of emergency).
Section	Communications Security
Risk Level	All
Hosting Models	On-Premises, IaaS, and PaaS solutions outside of the Azure Cloud
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Integration with the WHO global identity solution provides user-account lifecycle management. This means that when a user leaves the organization or takes a different job in the organization their access to your infrastructure will be automatically removed.

When access to infrastructure is provided with local accounts it is common for old, stale accounts belonging to users who have left the organization to exist, unmaintained, and with important privileged access.

This is clearly a serious risk.

Implementation:

For PaaS solutions in the Azure cloud this implementation will be enforced by the hosting team. For other situations you must be sure that your infrastructure is integrated with the WHO identity solution. This includes cloud systems which are not under the direct control of WHO.

There is clearly, however, a need for break-glass accounts. These must be limited to one per infrastructure item, and their credentials must be guarded as the most precious kind of secret.

Documentation:

Please include this in your architecture document.

Applicability:

This is applicable to all but SaaS solutions, which must have their own security infrastructure.

15.4 Network Security

Requirement	The application owner must ensure that appropriate services for network security as defined in the Global Cybersecurity Policy including related documents, and the WHO Architecture Guidelines are implemented.
Section	Communications Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

The way network infrastructure is designed (architected) and configured can mitigate many types of attacks. Firewalls, proxies, and the new Web Gateways or Web Application Firewalls all play a part.

The Cybersecurity team maintains a set of policies and guidelines for the implementation of such architecture, as does the Architecture and Infrastructure team.

Implementation:

The Architecture and Infrastructure team will guide you and have produced a set of guidance documentations in addition to the official Cybersecurity guidelines.

Documentation:

Please make sure that your architecture diagram is complete.

Applicability:

This is always applicable.

15.5 Final Architecture Review and Sign-Off

Requirement	The application or system owner must review the service architecture with the WHO Global Architecture and Infrastructure Team (AIS) and the WHO Cybersecurity Team (CST)
Section	Communications Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Guidelines are open to interpretation and misinterpretation. So, to be sure as sure as we can that your implementation complies with the network security requirement (5.4) and with our other requirements we conduct a review before go-live.

This meeting is conducted by the team lead of the Architecture and Infrastructure team in collaboration with the Cybersecurity team.

Implementation:

When you feel that your documents are ready for audit, and that your application is ready for go live please schedule a meeting with the Team Lead for HQ/AIS and your Cybersecurity focal point. The Cybersecurity team will provide additional guidance as required.

Documentation:

Meeting minutes will be sufficient, or a recording.

Applicability:

This is applicable to all projects.

15.6 Encryption at Rest

Requirement	Extremely sensitive information must be encrypted at rest per the certified WHO encryption method according to the e-Manual Encryption rule.
Section	Communications Security
Risk Level	VERY HIGH
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Regardless of the cybersecurity controls we apply elsewhere there remains a chance that an attacker can gain access to a system. If that should happen the attacker could gain access to extremely sensitive secret data.

Encrypting such data at rest (in storage) can prevent the attacker from accessing the data, or at least make doing so much more difficult.

Implementation:

Follow the OWASP cryptographic storage cheat sheet to find an encryption algorithm which is both recommended and supported by your project's platform. Implement this with proper key handling and careful key management.

Documentation:

Document your encryption algorithm, and your key handling practices in your architecture document.

Applicability:

This is applicable to situations in which a project stores extremely sensitive information. Normally we apply it to very high-risk applications.

15.7 Network Isolation

Requirement	The systems must be deployed in logically isolated networks protected by a security system and deployed with the appropriate security level or zone.
Section	Communications Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

If an attacker gains access to some part of your system their next step is to try to understand what nearby systems to which they might also be able to gain access, and to repeat from there, building up access, and privileged access if possible.

The answer to this is to compartmentalize, to use traditional firewalls and network routing to make it more difficult for access to one system to lead to access to another.

Implementation:

Implementation in this case depends on the hosting model and is most crucial in the case of on-premises hosting. An on-premises application must not expose other systems to an attacker who might gain access to it.

In a Platform or Infrastructure as a Service situation this must be part of your architecture consideration.

For Software as a Service providers this is also important, but we may not have visibility, and so must rely on their certifications.

Documentation:

For SaaS, please provide the vendor's certification documents. Please include this in your architecture documents in all other cases.

Applicability:

This is always applicable, but as we have seen the implementation is specific to the hosting model.

15.8 Remote Access to WHO Systems

Requirement	The application/system must use the current approved secure methods per the Remote Access rule within the e-Manual.
Section	Communications Security
Risk Level	All
Hosting Models	On-premises
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Sometimes remote administration must be done on systems which reside on WHO premises. To facilitate this work WHO provides a VPN (Virtual Private Networks) which can be used by WHO staff and external users, including contractors and consultants from vendors. This VPN solution is integrated with our central identity system, and thus with user-account-lifecycle, so that users leaving the organization are not accidentally left with access.

It is also secured with MFA, monitored, and logged.

Sometimes there is a temptation, either because of provisioning issues, or because people do not like MFA, or for some other reason, to create a backdoor using commercial remote access software.

This control forbids that behaviour.

Implementation:

Do not create a backdoor.

Documentation:

Document the enrolment of each named user required for your project in our identity system and the VPN system (MFA).

Applicability:

This is applicable to on-premises systems.

15.9 Private Key Handling

Requirement	Secret (private) encryption keys must be kept safe.
Section	Communications Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Because so much other sensitive and confidential data, and in some situations the very identity of WHO depend on them, private or secret encryption keys must be handled with the utmost care. Private keys should be stored in such a way that only the application or resource which needs them has access. If possible, no users should have access, and if that's not possible then only privileged users, and only for the shortest time possible.

Private keys must be re-encrypted if they are to be sent over a network. Ideally a private key for encrypting extremely sensitive content should transported on encrypted hardware in a diplomatic pouch.

Implementation:

Learn about the secure key handling facilities of your hosting environment. Make sure someone on your team understands proper private key handling and storage. You may of course seek the guidance of the Cybersecurity Team and the Architecture and Infrastructure team.

Documentation:

Your architecture document should refer to any secure key storage facility which your project uses.

Applicability:

Applicable anywhere encryption of any kind is used (including TLS).

15.10 Anonymisation and de-Identification

Requirement	The application or system must anonymize or encrypt personal identifiable information (PII) and sensitive data.
Section	Communications Security
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

We owe a special burden of care to our users, our colleagues, and anybody else who's personal information we handle or store. An attacker who has assembled other information about people could correlate that with information leaked from WHO and do actual harm to the person in question.

So, as a rule, if we have sensitive personal information, especially, but not limited to medical records, and other medical data, then that data should be either anonymised, or encrypted.

Implementation:

If the data cannot be anonymised then it must be treated as extremely sensitive, as per 5.9 above. Otherwise, the most recent best practices in data anonymization must be used, for example this paper from the journal Nature:

https://www.nature.com/articles/s41597-023-02256-2

Documentation:

Document your encryption techniques or your anonymization process.

Applicability:

Applicable wherever Personal Identifying Information is handled or stored.

16.1 Non-Repudiation Techniques

Requirement	The application owner must ensure that vendor contracts include non-repudiation methods and fraud prevention whenever financial transactions or value are involved. The methods themselves should be described by WHO (or suppliers) in an appendix of the contract, along with fraud protection.
Section	System Acquisition Development and Maintenance
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

When financial transactions are conducted electronically a method, or several methods, as appropriate must be implemented to guarantee non-repudiation, which is to say that there must be some mechanism in place to prevent anyone from denying the validity of a correctly executed transaction.

Since there are many ways to perform non-repudiation, the requirement is that this be written into the vendor's contract or attached as an annex.

Implementation:

Make sure that the vendor's contract includes a non-repudiation and fraud prevention clause in the main contract or in an annex.

Documentation:

The contract itself serves as its own documentation.

Applicability:

This is only applicable in applications which support financial transactions.

16.2 Transport Layer Encryption

Requirement	The application owner must ensure that communication paths between participating parties are end-to-end encrypted.
Section	System Acquisition Development and Maintenance
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Attackers can monitor, eavesdrop on, and even falsify unencrypted network traffic at various points across networks, both within WHO and on the public Internet.

For this reason, all communications between all parties to an application must be end-to-end encrypted when sent over the network (transport layer) using Transport Layer Security (TLS).

TLS provides the ability to encrypt traffic, but also to verify the identity of a host (presuming that they implement careful private key handling) and for a client to verify WHO's identity (given the same presumption on our part). None of these things is possible with unencrypted communication.

Implementation:

Following the private key guidelines above and in the e-Manual obtain a TLS certificate from an appropriate certificate authority and use the key-certificate pair to establish encrypted connections with your clients and servers.

Establish a mechanism to keep certificates up-to-date, preferably an automated one, since the current TLS certificate lifespan is only 90 days.

The Architecture and Infrastructure team can help you with this if you are hosting in Azure or on premises.

Documentation:

Document your decisions about certificate authorities and certificate renewal in your architecture document.

Applicability:

This is always applicable.

16.3 OWASP Top-Ten for Web Applications

Requirement	The application owner must ensure that developers follow secure coding best practices as defined in the OWASP Top Ten Proactive Controls.
Section	System Acquisition Development and Maintenance
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

The OWASP (Open Web Application Security Project) project is an industry-wide standards organization committed to creating checklists and cheat-sheets to simplify the application of security best practices by developers and systems administrators.

The OWASP Top-Ten Proactive Controls is a checklist which developers can follow during development and code review to avoid the most common kinds of vulnerabilities in web applications.

https://owasp.org/www-project-proactive-controls/

Implementation:

Look up the most recent version of the OWASP Top Ten, and then use it as a checklist during development and code review to help secure your code. As always, the Cybersecurity Team are available to provide guidance in the specifics.

Documentation:

Anonymised screenshots or evidence of code review conversations should be sufficient to document this process for an auditor. If your application is medium risk or above the Cybersecurity Team will also conduct penetration testing which will further exercise the kinds of vulnerabilities which the OWASP Top Ten is meant to prevent.

Applicability:

Web applications

16.4 OWASP Secure Headers

Requirement	Developers must enforce HTTP (Hypertext Transfer Protocol) secure response headers per the OWASP secure headers recommendations.
Section	System Acquisition Development and Maintenance
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

The OWASP project is an industry-wide standards organization committed to creating checklists and cheat-sheets to simplify the application of security best practices by developers and systems administrators.

The OWASP Secure Headers recommendations is a cheat sheet for web hosting, providing administrators and architects with a quick and straightforward way to assure that their web hosting environments comply with industry best practices in the use (and non-use) of HTTP headers.

https://owasp.org/www-project-secure-headers/

Implementation:

Follow the checklist to configure the headers of your Web Access Gateway, CDN (Cybersecurity for Developing Nations), Load-Balancer or other front-end system. Avoid introducing non-compliant headers on your backend systems, or mask them out at the reverse proxy layer.

Documentation:

Include a sample of HTTP responses in your architecture document, demonstrating that the secure headers recommendations have been applied.

Applicability:

Web applications

16.5 Change Management

Requirement	The application owner must ensure that any changes within the application or underlying infrastructure follow the WHO change management process, or that some other change management process is agreed upon between the vendors and project owners.
Section	System Acquisition Development and Maintenance
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Change management is the discipline within IT application management which aims to assure that all changes are *deliberate* and thoughtfully planned and that communication about the change takes place before and after the change. Having a change management process prevents at least some misunderstandings and can prevent nasty surprises.

It can also help in Cybersecurity forensics: when an attacker introduces a change, we can be sure that they will not follow the process, so change management documentation can help identify attacker supplied code and similar changes.

Implementation:

If your project is a WHO corporate project, then it must follow the corporate change management process which is documented elsewhere and is owned by the Operations and User Support Team.

If your project is a Public Health IT project or initiative you are free to define your change management process as you see fit, but you must have one.

Documentation:

Please document your change management process. This can be in your architecture document, or in a document of its own, as befits the size and complexity of the process.

Applicability:

This is always applicable

16.6 No Prod Data in Lower Instances

Requirement	The application owner must ensure that confidential or sensitive data is stored in the production environment only. Copies of this data must not be stored in a non-production environment. Any environment containing production data is production.
Section	System Acquisition Development and Maintenance
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Non-production instances of web applications and APIs (Applications and programming interfaces) may not have the resilience of the production deployment. Very frequently dev and test instances lack stringent configuration and controls. Sometimes things like WAFs (Web Application Firewall) (Web Application Firewall) and frontend balancers are not even configured for these instances. Often these instances have either no identity management solution or use ad hoc identity management with weak password hygiene.

Production data must *never* under any circumstances be copied into such an environment.

If you cannot for whatever reason produce appropriate test data, based on test cases, then at a minimum data must be totally anonymised and scrubbed to falsify addresses, monetary values and the names of users, vendors, pets, etc.

Implementation:

Just do not do it. Do not use production data in a non-production instance.

Or, in the worst case apply every cybersecurity control that is applicable to production to those other instances. Effectively this makes them production instances since they need to work with production Identity Management, to take advantage of user-account-lifecycle.

Documentation:

You can use screenshots and reports to demonstrate to the auditor that your non-production instances contain crafted test data, or randomly generated test data.

Applicability:

This is always applicable.

17.1 Relationship Management

Requirement	There must be a supplier relationship maintained for the entire lifetime of the project, and this relationship must be documented for audit.
Section	Supplier Relationships (3rd Parties)
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

In control 4.7 Patching we observed that new vulnerabilities are frequently found in upstream dependencies of a project which we enumerated to include things like operating systems, libraries, frameworks, and others.

This is also true of custom code, configuration, and architecture delivered by a vendor. Therefore, all projects must have vendor support or some other kind of support over their entire lifetime.

If a vendor cannot continue with the project for some reason, then a replacement must be found. Or at worst the project must budget for and hire in-house support staff.

Implementation:

Well in advance of the expiry of a vendor support contract a new contract must be negotiated either with the same vendor or with a new vendor.

Documentation:

The auditor would like to be able to cross-search and check between your risk assessment folder and GSM procurement. Therefore, we will add a document to your risk assessment called “PO (Purchase Order) Numbers.docx.” Place all PO Numbers in that document.

Applicability:

This is always applicable

17.2 Certification of Third Parties

Requirement	The application or system owner must ensure that all third parties in the chain are ISO 27001 certified or equivalent.
Section	Supplier Relationships (3rd Parties)
Risk Level	HIGH+
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Standards-based certifications can demonstrate a vendor's commitment to security and compliance and do ensure that the vendor has certain internal policies and that they follow them. Knowing that a vendor has these kinds of certifications helps assure us that they follow secure best practices, especially when we have limited visibility into their practical day-to-day procedures.

Implementation:

Make sure that any vendor you work with is ISO 27001 Certified.

Documentation:

Provide proof of certification.

Applicability:

This is applicable to high-risk and very high-risk applications, especially in a Software as a Service situation.

17.3 Supplier Commitment

Requirement	The application or system owner must ensure of the third parties' commitment to WHO's Cybersecurity requirements (Included in the contract).
Section	Supplier Relationships (3rd Parties)
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

A supplier's commitment to WHO Cybersecurity processes and requirements is the cornerstone of application security. Since we must trust the supplier, if they are not committed to our requirements then nothing else, we do have any real meaning, because it can all be worked around.

Implementation:

Suppliers should participate in good faith and demonstrate a commitment to our requirements and procedures.

Documentation:

The cybersecurity team will provide a sworn statement of commitment as a short email in this form:

"Vendor XYZ has demonstrated a commitment to WHO's Cybersecurity procedures and requirements by fully participating in meetings and document gathering in good faith."

Applicability:

This is always applicable.

18.1 Incident Management

Requirement	The application or system owner must ensure that the WHO Cybersecurity incident procedure is followed when a security incident is identified.
Section	Supplier Relationships (3rd Parties)
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

When a cybersecurity incident is detected, we need to start our response and forensic operations as soon as possible, especially if there is a possibility of either stopping the incident in real time, or observing an active attacker to learn their methods, and thus how better to recover when we finally eject them.

Therefore, an annex of the Cybersecurity policy in the e-Manual (section XIV 2.3) provides a procedure for contacting various parties at WHO in response to an incident.

Implementation:

Your architecture document or the equivalent should already have a procedure for responding to incidents, like downtime or other system failures. You should be sure that there is a decision point in this procedure for notification to WHO when you decide that an outage or other odd behaviour is due to an attack.

Documentation:

Please document your incident management procedure.

Applicability:

This is always applicable.

18.2 Formal Incident Reporting

Requirement	The application or system owner must provide bi-annual reporting on information security incidents related to the services delivered by third parties.
Section	Supplier Relationships (3rd Parties)
Risk Level	HIGH+
Hosting Models	SaaS
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

In contemporary development and hosting models we might not be aware of the entire chain of upstream suppliers. Therefore, we rely on the primary supplier for information about cybersecurity events concerning their own contracted suppliers. This kind of transparency builds trust, especially when there is an incident.

Implementation:

Make an annual report, detailing any incidents which have occurred with any third-party supplier, and how they were resolved.

Documentation:

Upload the reports to the Risk Assessment Documents directory.

Applicability:

This is applicable to high-risk and very-high-risk projects.

18.3 SIEM Use Cases

Requirement	The application or system owner, in cooperation with the Cybersecurity team, must develop use cases to be implemented in the WHO Security Information Event Management (SIEM) solution.
Section	Supplier Relationships (3rd Parties)
Risk Level	HIGH+
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

When an application handles many transactions, and therefore logs a lot of potentially valuable information there is a risk that the Security Operations Centre (SOC) analysts will miss actionable or interesting events in the general flow of log data. Worse the application might produce many false-positives, or events which look like a cybersecurity incident but in fact are not.

Implementation:

Produce examples of transactions or events which might be indicative of a cybersecurity incident. Schedule a meeting with our SOC analysists (via the Cybersecurity team) to discuss these cases and brainstorm others.

Documentation:

Include a document listing your use-cases in the documents folder.

Applicability:

This applies to high-risk and very-high-risk applications which produce send logs to the SIEM.

19.1 Backups

Requirement	The application or system owner must ensure that backup methods are identified and deployed based on the required recovery time and point objectives. (e.g., hot/cold backup, tape backup).
Section	Business Continuity
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Backups are absolutely required for system availability, and not always provided by cloud solution providers. The product team is therefore responsible for the creation of a backup procedure.

Implementation:

Create a backup procedure for your application. Specify frequency and retention requirements. Provide a test plan for a full restore from backup.

Documentation:

Provide a copy of the backup procedure either as part of your architecture document or in a separate document.

Applicability:

This is always applicable.

19.1 Disaster Recovery

Requirement	The application or system owner must ensure that the application or system is included in a Disaster Recovery Plan (DRP) and Disaster Recovery testing is performed annually in cooperation with the Architecture and Infrastructure team and the Corporate Applications team.
Section	Business Continuity
Risk Level	HIGH+
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

Critical systems are subject to WHO's Disaster Recovery Policy. This policy requires that for any such system there be a plan to restore the application to service quickly in case of the total loss of its normal hosting environment in a natural or human-caused disaster.

Implementation:

Create an appropriate disaster recovery plan and test it at least annually.

Documentation:

Include the plan and testing reports, or links to them, in the documents folder of your risk assessment.

Applicability:

Applicable to high-risk and very-high-risk applications, especially when they are critical to the entire organization.

20.1 Legal Compliance

Requirement	The application or system *sponsor* must make a written statement that the project will comply with WHO published policies in the Staff Rules and e-Manual
Section	Compliance
Risk Level	All
Hosting Models	All
Policy	WHO XIV.3.1 Cybersecurity Policy

Rationale:

While every WHO staff member has already signed a compliance document covering both the Staff Rules and the e-Manual, our suppliers and some contractors and consultants may not have signed such a document. The IOS auditor has asked that we therefor collect a statement from the project owner of compliance, on behalf of these third parties.

Implementation:

The project sponsor should be familiar with the e-Manual, but in this case should refresh their knowledge by re-reading section XIV 2.4 Cybersecurity Policy. Then project sponsor should send an email to cs_red_team@who.int with the following text, verbatim: “We and all third parties involved in project <your project name> will comply will all Staff Rules and the e-Manual.”

Documentation:

A Cybersecurity Team member will upload a copy of the project sponsor’s email to the documents folder.

Applicability:

This is always applicable.

Annex IV – Document Production Guidelines

Specifying the exact list of documents to be delivered to pass successfully through some future audit is tricky, but we can specify a few minimal documents to be delivered

- An architecture document describing your network layers, and network security tools, key storage, frameworks, operating systems, database types and versions, etc.
- An administration manual, and if necessary, a users' manual.
- A file containing the PO Numbers for all suppliers.
- Any certifications and audit reports from the various suppliers.
- The email pledge for control 9.1.

Apart from the pledge email, these documents should already exist, and will be produced in the normal course of a healthy project. For this reason, at least we do not *require* that they follow any specific form. We do provide some samples, however, in the next annex.

For a project with multiple software-as-a-service components the documents list can become rather long:

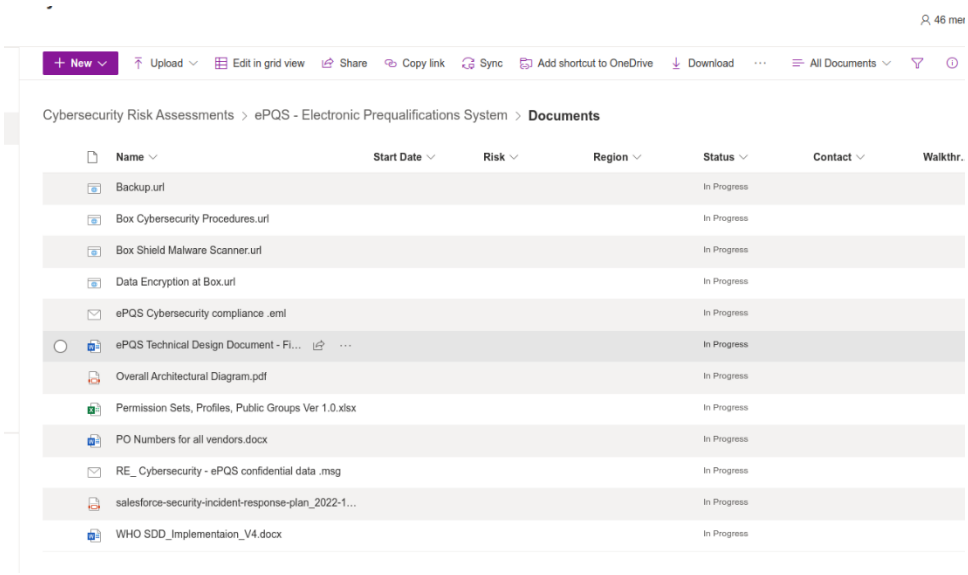

Again, this should not be intimidating. These documents should already exist in a healthy project. However, the requirements in Annex III do specify certain documentation to demonstrate that the control has been implemented. Most of these should be added to the architecture document.

Annex V – Sample

These are sample pages from an architecture document. **Do not feel bound** to the use exactly the same architecture. This is an example of how to document *your* architecture decisions.

In this example the developer used C# and .NET. But you could use something else. Just put that in the document instead.

Likewise, this project used PostgreSQL. You are not bound to use PostgreSQL, but please tell us which database you have chosen.

Development environment

- Visual Studio 2022
- Visual Studio Code
- JetBrains Rider
- .NET 6.0 LTS
- PostgreSQL >=11
- Node >=v16
- Angular 13

GIT repository structure

- Build directory contains all required files for CI\CD.
- Docs directory contains this project documentation.
- Sources directory contains source files.

Application infrastructure

The aim of the UHC CMS project is to manage content of the existing UHC SPDI tool (https://uhcc.who.int/); therefore, the UHC CMS project based on the infrastructure of the UHC SPDI tool and logically extends it when needed.

We already have three application gateways on the UHC SPDI tool for our 3 environments: dev, test and prod. We will add 3 new static web applications allowing managing the content for these three environments to the existing UHC SPDI gateways. Each of these new gateways will use the following path routing:

{environment.domain}/**packages** - UHCC SPDI static angular web application
{environment.domain}/**cms** - UHCC CMS static angular web application
{environment.domain}/**api** - .NET API 6 web service

To depict these changes in the architecture diagram of the UHC CMS project, we have extended the architecture diagram of the UHC SPDI tool as following:

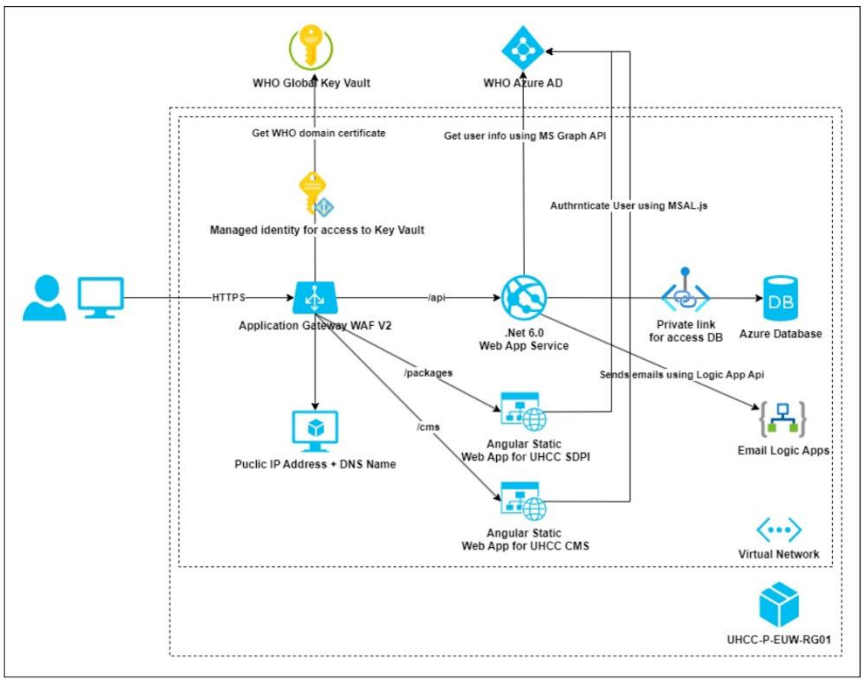

Back-end architecture

This section was added for UHC CMS Solution Architecture document for consistency only. It completely coincides with the Back-end Architecture of the UHC SPDI tool.

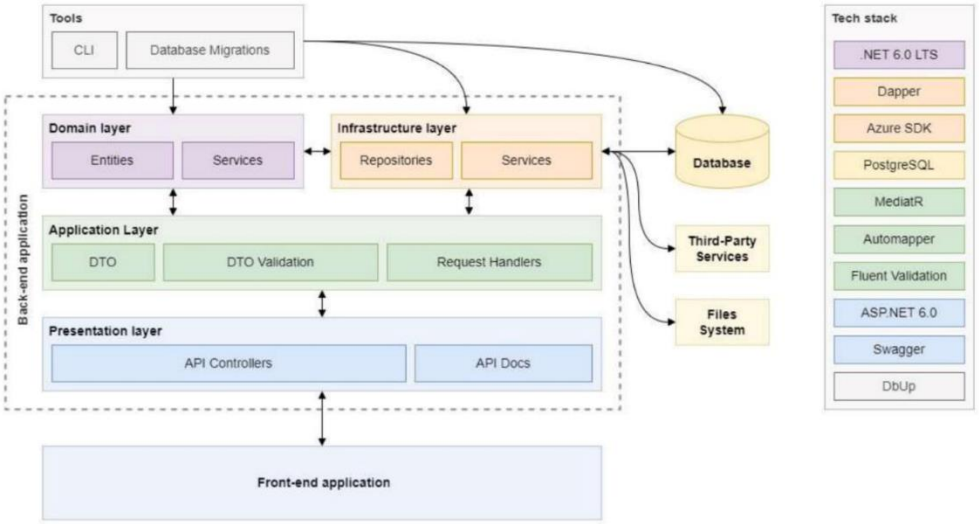

Key libraries and frameworks

- .NET 6.0 LTS is a cross-platform version of .NET for building websites, services, and console apps.
- ASP.NET 6.0 is a web framework built upon .NET Core
- MediatR simple mediator pattern implementation in .NET https://github.com/jbogard/MediatR
- Dapper is a simple object mapper for .NET and owns the title of King of Micro ORM in terms of speed and is virtually as fast as using a raw ADO.NET data reader. An ORM is an Object Relational Mapper, which is responsible for mapping between database and programming language.
- Fluent Validation is a popular .NET library for building strongly-typed validation rules. https://fluentvalidation.net
- Automapper is a simple little library built to solve a deceptively complex problem - getting rid of code that mapped one object to another. https://automapper.org
- DbUp .NET library that helps you to deploy changes to SQL Server databases. It tracks which SQL scripts have been run already, and runs the change scripts that are needed to get your database up to date. https://dbup.readthedocs.io

.NET solution structure

- Tools folder contains utility console applications o Who.Uhcc.Cli assembly - command line interface that can be used by Azure Web Jobs
- Who.Uhcc.Migration assembly - database migration tool
- Who.Uhcc assembly implements Domain, Infrastructure, and Application layers
- Who.Uhcc.Server assembly implements Presentation layer

Front-end architecture

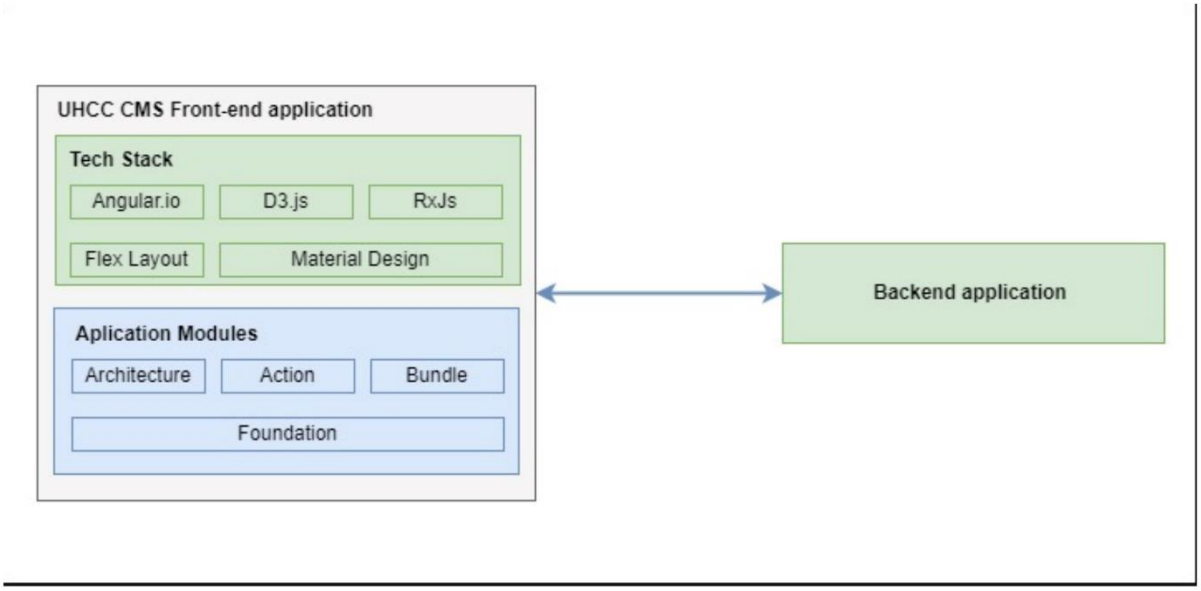

Files structure

- Sources directory - root of Angular project template
- Foundation directory - group of core modules
  - core module contains very high level and common utils
  - ui module implements all UI components according to the design
  - uhcc-api module implements integration with UHCC backend
- Features directory - group for feature modules
  - common module contains shared components between feature modules
  - package architecture module contains components that related to the package architecture: groups, subgroups, categories, interventions etc.
  - action module contains components that related to the action editor feature
  - bundle module contains components that related to the bundle editor feature

Docs folder

Documents located in ~\Docs directory of the root of the repositories (frontend and backend).

CI / CD processes

CI and CD are configured in Azure (https://dev.azure.com/ and https://portal.azure.com/). After each commit in the repository, build and deploy automatically starts.

develop branch -> DEV environment

test branch -> TEST environment

main branch -> PROD environment

Management process

Do a change in a DEVELOP branch => Deploy to DEV => Approve by QA =>
Merge DEV to TEST branch => deploy to TEST => Approve by QA and business => Deploy to PROD after business approval.

Azure DevOps pipeline

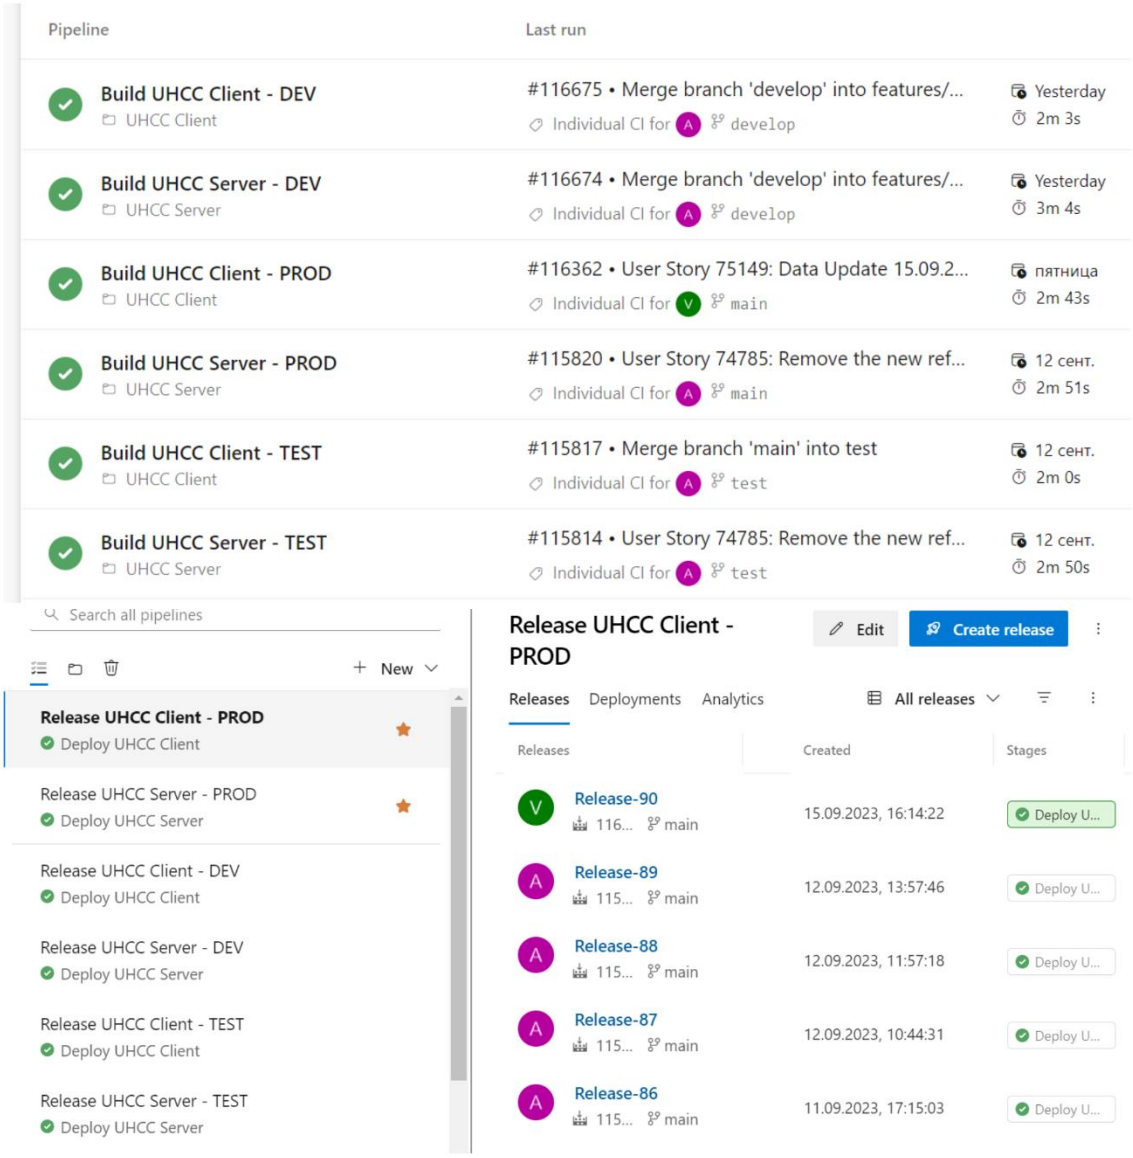

We are using Azure DevOps pipelines to build and to deploy static web applications and api web services for each environment. We will add additional 3 build and release pipelines for CMS application.

Back up

Database automatically backed up by Azure every 24h.

Environment monitoring

Logs are configured for the backend in Azure.