Commissioned by:

Federal Office for Information Security

In Cooperation with:

secunet

Study

Open RAN Risk Analysis

5GRANR

Version: 1.2.1
Date: 21 February 2022

Authors: Stefan Köpsell (Barkhausen Institut)
Andrey Ruzhanskiy (Barkhausen Institut)
Andreas Hecker (Advancing Individual Networks GmbH)
Dirk Stachorra (Advancing Individual Networks GmbH)
Norman Franchi (Advancing Individual Networks GmbH)

Executive Summary

This study addresses the question of the **security risks** that arise from the **O-RAN** implementation of a 3GPP RAN as specified by the **O-RAN** Alliance. It starts with a functional description of a 3GPP RAN and the O-RAN architecture. The study then presents a risk analysis that was undertaken on this basis while factoring in the protection goals of **confidentiality, integrity, accountability, availability and privacy. Three stakeholders** are considered in assessing the risks associated with a violation of these security objectives: the **user** of a 5G network, the **operator** of a 5G network and the **state** as an entity representing the societal perspective.

Since the existing O-RAN specifications still lack specificity in many areas and provide few guidelines in the area of security in particular, two types of scenarios were considered in the risk analysis. The **worst-case scenario** assumes that none of the optional security measures have been taken, while the **best-case scenario** assumes that all the (optional) security safeguards have been implemented.

For the risk analysis, **various powerful attackers** have also been considered, including an **external attacker, a 5G user, an "insider", a cloud operator** and a **RAN operator**.

The risk analysis has revealed that **medium to high security risks** can be identified in **numerous** interfaces and components specified in the context of O-RAN. This is hardly surprising, as the current development of the O-RAN specifications does not adhere to the paradigm of **"security/privacy by design/default"**. It also **fails to take account** of the principles of **multilateral security** (i.e. assuming minimal trustworthiness of all stakeholders).

In the course of the risk analysis, some **potential improvements** were identified that could minimise risk. These can be found in the recommendations at the end of this study. It is important that security improvements are included in the specification **now** to **avoid** the **security weaknesses** that occurred in the development of the 3GPP standards.

Commissioned by:

Federal Office for Information Security

This risk analysis was commissioned and financed by the Federal Office for Information Security (BSI). The BSI had no influence on the results.

Disclaimer: This English version of the risk analysis has been translated by a translation agency, commissioned by BSI. While the authors proofread the English version to ensure best possible correctness of technical details, we need to refer to the German version in case of errors or ambiguity.

Table of contents

1 Introduction

Fifth generation (5G) mobile networks offer a variety of new use cases, particularly those relating to the networking of “things”. This gives 5G an increasingly central role in the area of basic communications infrastructures – especially as a basic communications structure for critical infrastructures such as those for power, water, logistics and transport. It is therefore essential to be aware of the risks involved in implementing 5G as a communications infrastructure in the context of IT security and data protection. This type of risk analysis is thus the main focus of this study. By way of limitation, this study does not address an overall 5G system, consisting essentially of a 5G radio access network (5G RAN) and the 5G core network (5G Core). Instead, the risk analysis solely considers the 5G RAN, with the focus on recommending a concrete 5G RAN implementation, which is specified by O-RAN Alliance[1]. This recommended implementation is referred to as O-RAN.

The main objective of this study is therefore to analyse the risks associated with O-RAN. It covers the threats and risks identified in the current specifications and can be used as a starting point for deciding which measures are necessary in the future to minimise these risks.

The study starts with an overview of the next generation radio access network (NG-RAN) from a technical and functional perspective. In particular, it examines the O-RAN architecture. This is followed by a presentation of the risk analysis methodology used in the study. In particular, it explains the protection goals and the underlying attacker model in detail. In chapter 4 the study briefly considers other studies that have analysed risks or threats related to NG-RAN or O-RAN. Of particular importance here are the analyses provided by the O-RAN Alliance itself. This is followed by the main part of this study, which analyses the actual risks of O-RAN. A related risk analysis with regard to NG-RAN in general can be found in Appendix A. Background knowledge of general NG-RAN risks is essential to understanding the O-RAN risk analysis. Readers who do not have this background knowledge are advised to read Appendix A before the O-RAN risk analysis in chapter 5. The risk analysis starts with individual components and interfaces of O-RAN in order to provide an evaluation of the security risks for O-RAN as a whole.

This study concludes with recommendations on safeguards for improving O-RAN security.

[1] https://www.O-RAN.org/

2 Next Generation Radio Access Network (NG-RAN)

NG-RAN, Open RAN and O-RAN are the concepts and current components of the RAN area under review. They appear throughout the entire document. The last two terms in particular are often used erroneously as synonyms. This is understandable, as they not only sound similar, but are also inextricably linked. For example, the “O-RAN Alliance” has been founded to realise the concept of “Open RAN”. The introductory chapter 2.1 explains the context of all three terms in relation to each other.

Chapter 2.2 delves deeper into the NG-RAN architecture, starting with the 3GPP standard and continuing to the O-RAN specifications. Chapter 2.3 explains the interfaces introduced by O-RAN that are essential for the security considerations in the later chapters. Chapter 2.4 deals with applications, focusing in particular on machine learning (ML) with O-RAN. Chapter 2.5 concludes with information on existing O-RAN software, which is open to anyone interested in developing their own solutions.

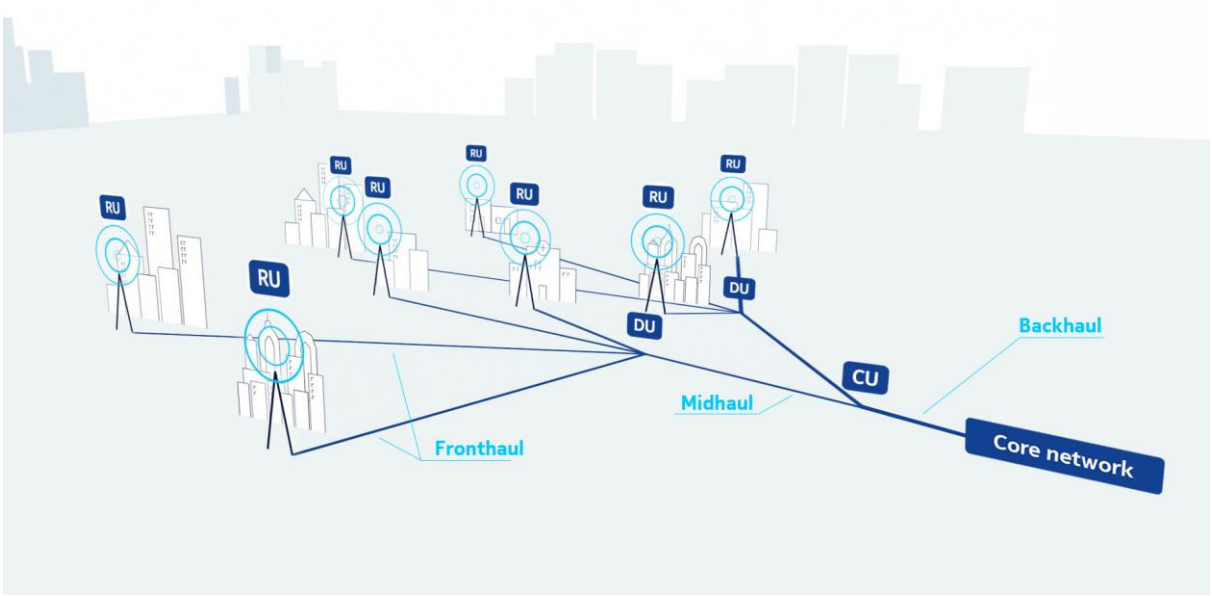

Figure 1: Simplified representation of an Open RAN architecture (source: Nokia)

2.1 Introduction and concepts

A radio access network (RAN) forms the link between end devices and the core network in a mobile radio network. For the fifth generation (5G) of mobile networks, 3GPP specified the New Radio (NR) mobile air interface and the next generation RAN (NG-RAN) in Release 15. NG-RAN offers access via both NR and E-UTRA (Evolved UMTS Terrestrial Radio Access), the air interface of 4G/LTE (Long Term Evolution). Mixed operation of LTE and NR base stations connected to a core network (4G or 5G) is called “non-standalone” mode (NSA mode); operation based solely on NR and a 5G core network is called “standalone” mode (SA mode) [1].

The first 3GPP studies on NG-RAN began in Release 14 and eventually led to the technical report [2]. In addition to the possibility of operating in both SA and NSA mode, Release 14 already includes another characteristic feature of NG-RAN: the option of separating the 5G base station into a centralised unit (CU) and one or more distributed units (DUs). This division of the base station results in additional possibilities and flexibility for deployment, as outlined in the example in Figure 1.

2.1.1 Open RAN[2]

In parallel to the NG-RAN work carried out by 3GPP, the concept of the "opening of the RAN" ("Open RAN") has been pursued for several years as an alternative to the traditional approach. In the traditional approach to the provisioning of a RAN, as found in most public mobile networks, a RAN is a monolithic solution. The internal interfaces within this solution are proprietary and largely undisclosed to third-party vendors. The related hardware is an independent solution that is not compatible with the equipment of other manufacturers. The functions and proprietary interfaces of a monolithic RAN solution are highly optimised, but inflexible as a result. Therefore, the goal of pursuing the Open RAN concept is to make the RAN more independent of proprietary technology by defining specifications for open interfaces and abstracting network elements and functions from the hardware.

There is a strong incentive for mobile network operators to realise the Open RAN concept because it is very expensive to equip their networks and there is very little competition in the market for RAN equipment. With each new generation of public mobile communications, the entire technology has to be replaced or built in parallel – at costs running into billions. The hardware grows more complex with every upgrade; more applications and features have to be supported, and thus prices increase. The market is dominated by three main vendors: Ericsson, Nokia and Huawei – and Huawei products are not allowed to be used by mobile network operators in some countries. The usually gradual entry of new vendors is not possible because the equipment for a RAN has to be supplied entirely by one vendor due to the interdependence of the product solutions [3]. It is therefore intended that the Open RAN will make it possible in the future to increase competition and reduce costs, enabling the RAN to be built and operated using interoperable components from different manufacturers. Furthermore, modularisation should open up the option of only replacing (software) components when extensions to a RAN are required rather than the complete technology.

By referring to the elements already standardised by 3GPP, it is possible to identify the next steps towards achieving the Open RAN concept. In addition to the already standardised elements CU and DU, the Open RAN environment specifies the radio unit (RU), which is integrated in or located near the antenna. This means that Open RAN defines the fronthaul as the connection between the RU and the DU. The midhaul between the DU and the CU and the backhaul for connecting the RAN with the core are already defined by the NG-RAN standardised by 3GPP (see Figure 1).

2.1.2 O-RAN

Based on the idea of Open RAN, the O-RAN Alliance was founded in 2018 with companies from the telecommunications sector, mainly including an international consortium of network operators. The goal of this alliance is to standardise an NG-RAN that is largely based on virtualised and interoperable components. The desired flexibility should enable tailor-made solutions for individual applications, which can be quickly reconfigured and efficiently optimised as required. These efforts have resulted in the development of the O-RAN architecture as an extension of the RAN architecture standardised by 3GPP. The O-RAN architecture should be seen generally as an enhancement to existing 3GPP standards. In the long term, the results of O-RAN may be incorporated back into the 3GPP standards. Until then, the O-RAN architecture will exist as an addition to the 3GPP architecture.

The O-RAN Alliance supports the O-RAN software community in a cooperation with the Linux Foundation. In summer 2020, the alliance announced its intention to launch the 5G SD-RAN (5G Software-Defined Radio Access Network) project together with the ONF (Open Network

[2] There are different spellings of this term (including with or without a hyphen or space) that are often used synonymously. In this study the form "Open RAN" is used unless the link to a particular reference specifies a different spelling.

Foundation) to promote and facilitate the creation of open-source software for mobile 4G and 5G RAN deployments.
The existence of the SD-RAN group does not mean that all components of the O-RAN architecture will be open-source, or even that they have to be based on the open-source code developed by the project group. However, the software community to be built around this open-source project will influence the development of applications for configuration and optimisation. One example is the involvement of RIA (RAN Intelligence and Automation), a sub-group of the TIP (see chapter 2.1.3) in the 5G SD-RAN programme. The goal of the RIA is to develop and implement AI/ML[3]-based applications for a variety of RAN use cases, including SON (self-organising network), RRM (radio resource management) and mMIMO (massive multiple input and multiple output) [4].

2.1.3 TIP

The Telecom Infra Project (TIP) was founded in 2016 as an engineering-oriented collaboration. The goal of the project is to develop and provide a global telecommunications network infrastructure that facilitates global access for all interested parties. TIP helps promote an ecosystem of hardware and software providers, initiates plugfests and develops blueprints. While it does not write specifications, it is involved in promotions, training and the implementation of Open RAN solutions around the world.
In 2020, the TIP “OpenRAN” project and the O-RAN Alliance announced a liaison agreement to ensure their alignment in the development of interoperable, disaggregated and open RAN solutions [5].

2.1.4 Additional groups

Alongside 3GPP and TIP, there are other organisations involved in Open RAN or the O-RAN specification and O-RAN development that also make important contributions. These include the Small Cell Forum, the Open Networking Foundation (ONF) mentioned above and the OpenAirInterface Software Alliance (OSA) [4].

2.2 NG-RAN architecture

This chapter explains the NG-RAN architecture in detail. It starts with the 3GPP specification (chapter 2.2.1) and an explanation of the division of user and control planes (chapter 2.2.2), which serve as a basis for the O-RAN architecture (chapter 2.2.3). Chapter 2.2.4 deals with the challenges posed by integrating RAN components from different manufacturers, followed by an explanation of the topic of RAN sharing (chapter 2.2.5) and a selection of options that the O-RAN Alliance has identified based on the current concepts.

2.2.1 Specification according to 3GPP

Figure 2 shows an NG-RAN consisting of two 5G base stations (next generation nodeBs, gNBs) and the interfaces for connecting the core and the end devices. The gNB is divided into two logical functions: the CU and DU (see introduction in chapter 2.1).

[3] Artificial intelligence (AI) refers to the ability of a computer to mimic human thinking; it is a generic term for describing the performance of a machine. Machine learning (ML), on the other hand, describes a range of methods used to gain insights from data for AI applications.

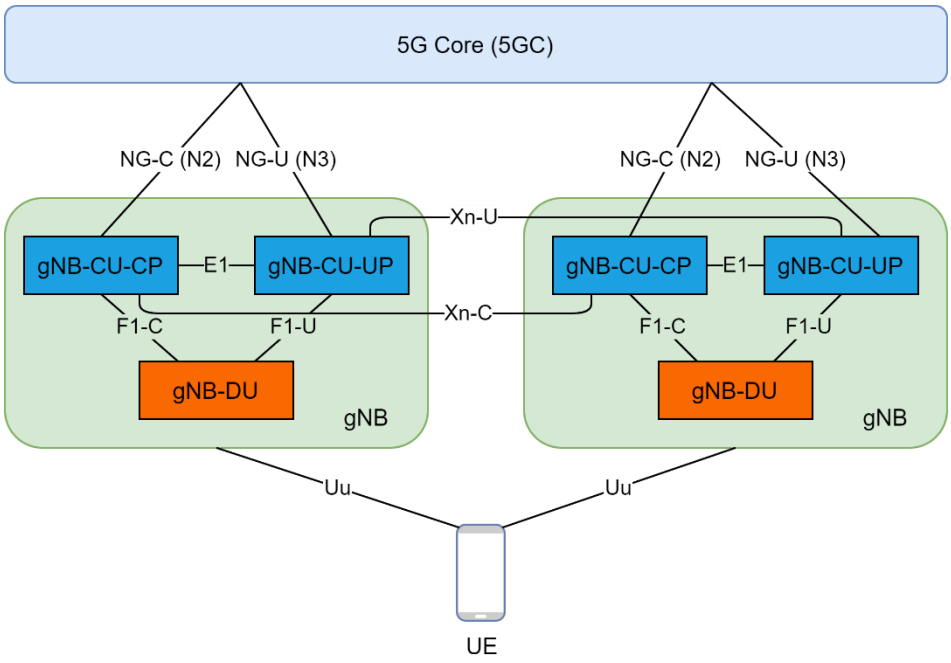

Figure 2: 3GPP architecture for 5G NR

The DU is the baseband unit in a 5G RAN. It handles the layer 1 and layer 2 processing and fulfils critical functions such as coding/decoding, scheduling, MIMO processing and beamforming. It makes sub-millisecond decisions about allocating radio resources within a cell based on factors such as policies, interference conditions or the type and distribution of user devices.
The CU is a new node in 5G that does not exist as a discrete function in 4G. In the 3GPP RAN architecture, the CU offers layer 3 functions such as connection and mobility management. It is divided into the CU-UP for user plane processing and the CU-CP for control plane processing (see chapter 2.2.2). The E1 interface defined in [6] serves to exchange configuration data and other information between the disaggregated units. The CU-CP and CU-UP are each connected to the DU via the F1-C and F1-U interfaces, which are defined in [7]. Connecting to the 5G core, there is the NG-C interface between the CU-CP and the AMF (Access and Mobility Management Function) and the NG-U interface between the CU-UP and the UPF (User Plane Function).
It should be noted that the 3GPP architecture does not specify the Remote Radio Unit (RRU), i.e. the implementation of the interface between PHY and RF layers is left to the vendors.

2.2.2 Control/User Plane Separation

Control/User Plane Separation (CUPS) refers to the complete separation of the control plane and the user plane in mobile radio networks. The control plane is responsible for functions like managing user connections, defining QoS guidelines or user authentication. The user plane is responsible for data traffic transport.
CUPS was first introduced for 4G in 3GPP Release 14 for the EPC (Evolved Packet Core) [8] and extended to 5G systems in Release 15 [9]. The main motivation for this separation is to be able to scale the user plane independently of the control plane, thus giving operators greater flexibility in the dimensioning of their networks. For instance, the user plane can be extended if the data traffic increases without changing the functions of the control plane.
CUPS allows for free configuration of the user plane to meet application-specific requirements for routing, data encapsulation, traffic control or other tasks received as specified from the control plane. This enables the implementation of different user data transmission solutions that coexist in the same user plane and can be dynamically selected according to the needs

of a specific traffic situation. Furthermore, a separate user plane enables shorter delays in transporting data traffic and thus makes it possible to meet higher latency requirements.
Figure 3 shows the protocols used in the user plane between the UE and the gNB on the Uu interface and their conversion to the protocols used on the NG-U (N3) interface between the gNB and the UPF in the 5G core. As usual, the physical layer L1 and the MAC layer L2 exist for network access. The RLC (Radio Link Control) is mapped to IP. PDCP (Packet Data Convergence Protocol) is converted in the gNB to UDP (User Datagram Protocol). UDP is a connectionless and unreliable transmission protocol that is also neither secured nor protected. SDAP (Service Data Adaptation Protocol) is converted to GTP-U (GPRS Tunneling Protocol – User). GTP-U tunnels are used to transfer encapsulated user data packages (transport packet data units, T-PDUs) and for signalling messages between a particular pair of GTP-U tunnel endpoints. The tunnel endpoint ID contained in the GTP header specifies the tunnel to which a particular T-PDU belongs [10].

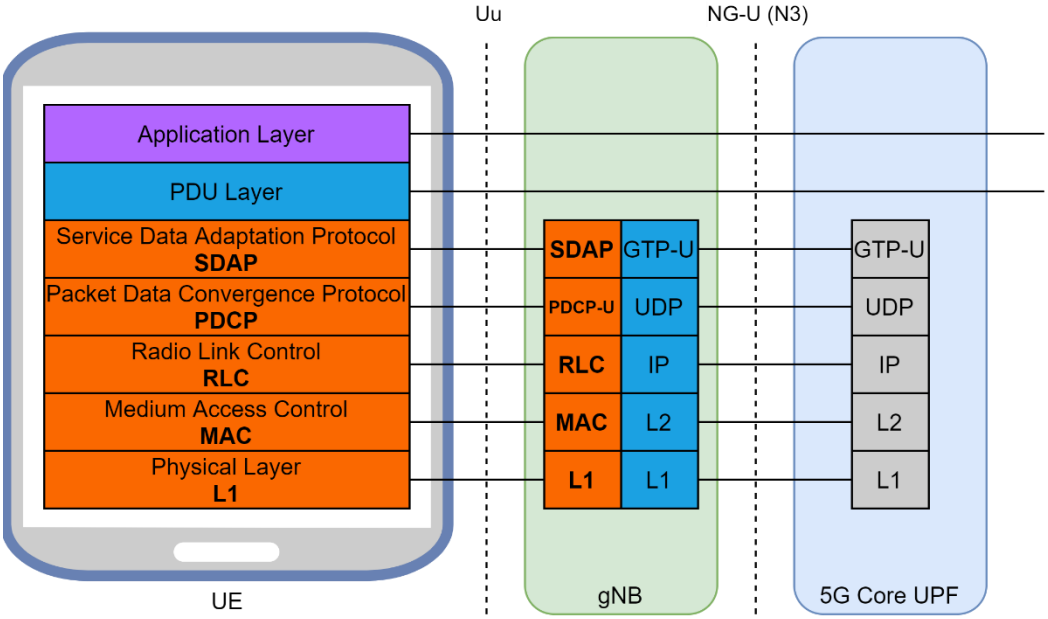

Figure 3: Protocol layers of the user plane

Figure 4 shows the protocols used in the control plane and the conversion to the protocols in the gNB that are used on the NG-C (N2) interface between the gNB and the AMF in the 5G core. The lowest three layers are analogous to the layers in the user plane (see above). For the control plane, PDCP is converted to SCTP (Stream Control Transmission Protocol). Unlike UDP, SCTP is a reliable, connection-oriented network protocol [11]. The Radio Resource Control (RRC) protocol is converted in the gNB to NG-AP (Next Generation Application Protocol), which is described in [12].

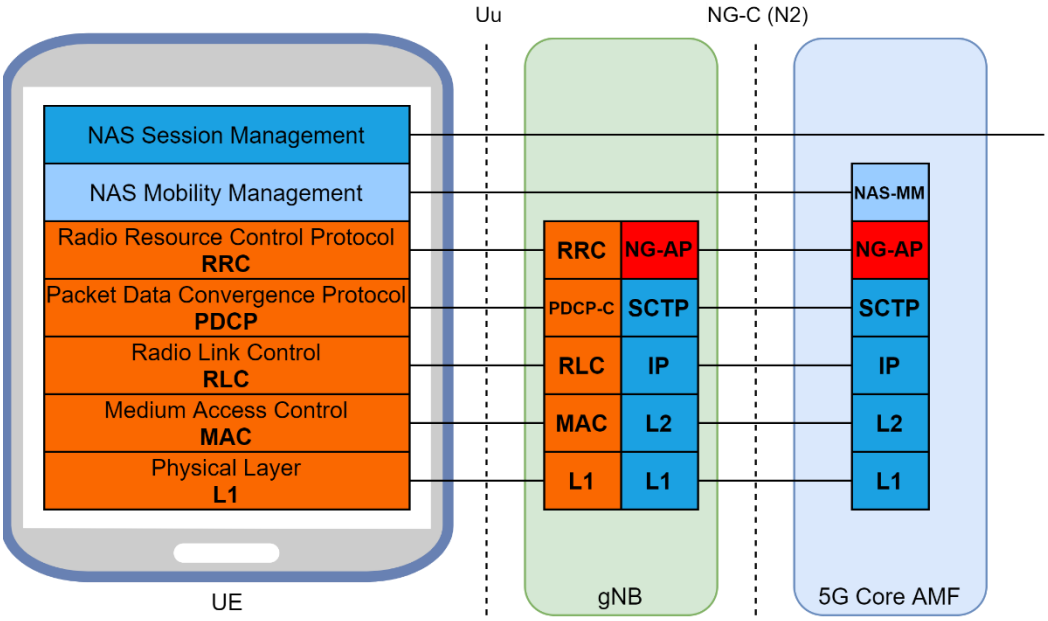

Figure 4: Protocol layers of the control plane

The NAS (non-access stratum) [13] describes a set of protocols that are used to transmit radio-independent signalling messages between the UE and certain functions of the core network. The two fundamental protocols are 5GS Mobility Management (5GMM) and 5GS Session Management (5GSM). The 5GMM protocol is used between the UE and the AMF to transport messages for UE registration, mobility and security. It also serves to transport the 5GSM protocol, which supports the management of PDU session connectivity and is applied between UE and SMF (Session Management Function) via the AMF.

2.2.3 Specification according to O-RAN

The O-RAN Alliance further divides the CU and DU network functions in accordance with the 3GPP definition. Figure 5 shows the architecture and highlights the division of functions and interfaces between 3GPP and O-RAN [14]. The elements defined in 3GPP, CU-CP, CU-UP, DU and eNB (evolved node B) are assigned the “O” prefix in the O-RAN specifications to clarify the difference. This abbreviation stands for O-RAN. For example, O-DU is the shortened form of “O-RAN Distributed Unit”.

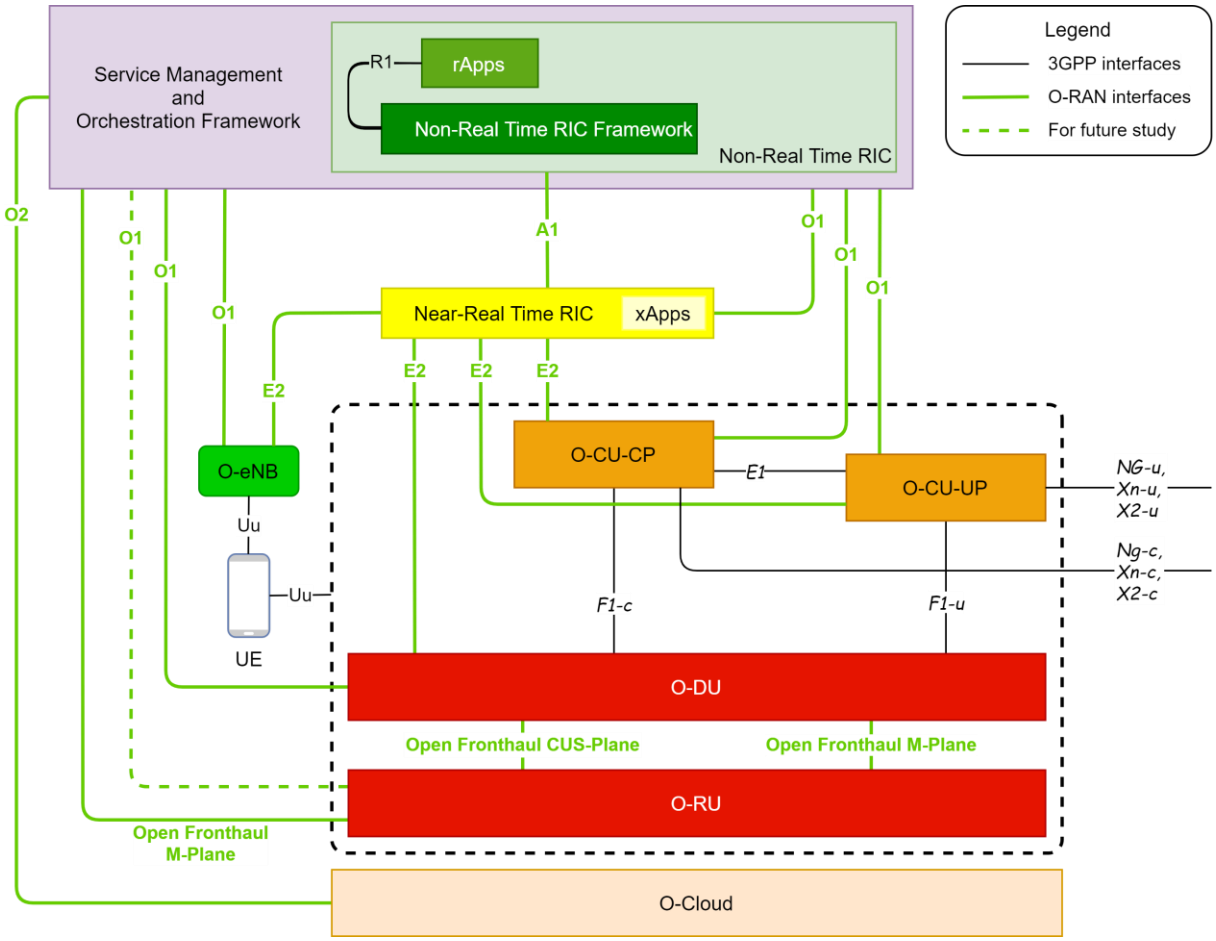

Figure 5: Logical O-RAN architecture, including Uu interface for O-RAN components and O-eNB [14]

The new functions defined by O-RAN are:

- The service management and orchestration (SMO) framework
- The RAN Intelligent Controllers (RICs) in their variants non-real-time and near-real-time, abbreviated hereafter as “non-RT RIC” and “near-RT RIC”
- The Remote Unit (O-RU)
- The O-Cloud

The relevant interfaces are defined in accordance with the defined functions:

- A1 between non-RT RIC and near-RT RIC
- E2 between the near-RT RIC and the E2 nodes O-CU-CP, O-CU-UP, O-DU and O-eNB
- O1 between the SMO and the managed units near-RT RIC, O-CU-CP, O-CU-UP, O-DU and O-eNB, as well as O-RU (which is still under investigation)
- O2 between SMO and O-Cloud
- Open Fronthaul between O-DU and O-RU, or between SMO and O-RU

In principle, each of these functions could be developed and provided by different suppliers using conventional hardware while ensuring the data exchange between the components via the open, specified interfaces. This would realise the goal of the Open RAN concept.
According to [15], by building its architecture on 3GPP's 5G NR architecture, the O-RAN Alliance will benefit from the enhanced 3GPP security features introduced for 5G and described in [16], including:

- Enhanced user identity privacy via Subscription Concealed Identifier (SUCI)
- Full protection of control/user plane traffic between UE and gNB (encryption and integrity protection) over the air interface

- Full protection of the gNB interfaces, including the E1 interface between CU-CP and CU-UP and the F1 interface between CU and DU
- Advanced home network control (authentication)
- Additional security for network slices based on service level agreements (SLAs)

2.2.3.1 Service management and orchestration (SMO) framework

In the O-RAN architecture, the SMO framework is responsible for managing the RAN domain[14]. This includes the following tasks:

- Supporting FCAPS via the O1 interface between SMO and the O-RAN network functions or, in the hybrid model, via the Open Fronthaul M-Plane interface between SMO and O-RU
- RAN optimisation via the A1 interface between the non-RT RIC in the SMO framework and the near-RT RIC
- O-Cloud management and orchestration for platform resource provisioning and workflow/workload management via the O2 interface between the SMO and the O-Cloud

A formal interface between SMO and non-RT RIC is not currently defined. The implementation of SMO and the boundary to the non-RT RIC framework is therefore a free design decision. Certain functionalities may be included or excluded in a non-RT RIC implementation. Essentially, interfaces A1 and R1 are inherent to the non-RT RIC, while interfaces O1 and O2 are not.

2.2.3.2 O-Cloud

The O-Cloud is a cloud-based computing platform that comprises a collection of physical infrastructure nodes and hosts the following components [14]:

- The relevant O-RAN functions near-RT RIC, O-CU-CP, O-CU-UP and O-DU
- The supporting software components, such as an operating system, virtual machine monitor, container runtime, etc.
- The corresponding management and orchestration functions

2.2.3.3 RAN Intelligent Controller (RIC)

The RIC is specified by the O-RAN Alliance as an integral part of the O-RAN architecture. As can be seen in Figure 5, the RIC appears in two forms, each of which is adapted to specific control loop and latency requirements. A detailed introduction to the topic of RICs can be found in [4].

The near-RT RIC has direct interfaces to the O-CU-CP, O-CU-UP and O-DU via the E2 interface. It enables their programmatic control in time cycles ranging from 10 ms to 1 second. Due to strict latency requirements with control loops of less than 10 ms, real-time functions such as the RRM remain on the DU [14]. In principle, the near-RT RIC can programmatically configure the O-DU to improve its operation; this is also conceivable in the future. For example, the near-RT RIC could be used to change the scheduler behaviour on the DU. In the O-RAN Alliance, Working Group 3 (WG3) is responsible for the specification of the near-RT RIC, including the E2 interface.

The non-RT RIC is specified for control loops of more than one second. It establishes policies for the higher network layers. They can be implemented in a RAN either via the near-RT RIC through the A1 interface or via the SMO connection to the RAN nodes through the O1 interface. Conversely, the non-RT RIC collects data generated in standardised formats from the RAN components via the O1 interface in order to serve traditional RAN optimisation functions. By being placed higher up in the architecture, the non-RT RIC can access larger RAN data records

generated over longer periods of time to gain deeper insights into performance and identify potential optimisations that are not visible in sub-second processing.
The non-RT RIC can also be connected to other network data sources (e.g. to jointly optimise radio, IP network and edge cloud performance). In addition, it can access data sets outside the network itself (those related to traffic, emergency services, weather, mass public events, etc.) so that the network can prepare for or respond to real-world events.
Working Group 2 (WG2), together with the associated A1 interface, is responsible for the non-RT RIC specification in the O-RAN Alliance. The O1 interface is part of Working Group 1 (WG1).
It should be noted that the control loops introduced by the RIC can potentially conflict with results of other procedures, such as those implemented by the mobile network operators to provide policy updates for the DU and CU functions.

2.2.3.4 O-CU and O-DU

The RAN nodes O-CU-CP, O-CU-UP and O-DU correspond to the nodes CU-CP, CU-UP and DU defined in 3GPP. As a result, the O-CU-CP and O-CU-UP serve the E1 interface and the F1-C/F1-U interfaces for the O-DU, as being specified by 3GPP. As specified in [17], the O-CU-CP implements the RRC and PDCP protocols and the O-CU-UP implements the PDCP and SDAP protocols for connecting to the UE, while the O-DU implements the RLC, MAC and high-PHY functions of the radio interface (see chapter 2.2.2).
The O-RAN architecture introduces the following changes [14]: The O-CU-CP, O-CU-UP and O-DU terminate the E2 interface to the near-RT RIC and the O1 interface to the SMO framework. The O-DU serves the Open Fronthaul interface, including the Open Fronthaul M-Plane interface for the O-RU to support O-RU management either in the hierarchical model or in the hybrid model (see [18] or chapter 2.3.7).

2.2.3.5 O-eNB

The O-RAN architecture incorporates 4G/LTE via the O-eNB. The O-eNB can be an eNB defined according to [17] or an ng-eNB (next generation eNB) defined according to [19]. The associated interfaces and protocols must be supported accordingly. The E2 and O1 interfaces must also be supported for O-RAN compatibility.

2.2.3.6 O-RU

Compared to 3GPP, the O-RAN architecture also includes the antenna unit, which is called the O-RAN Radio Unit (O-RU). It contains the antenna structure as well as the analogue radio frequency (RF) and power amplifier. The O-RU is connected to the O-DU via the Open Fronthaul (Open FH) interface. It maps low-PHY functions of the radio interface to the UE. Compared to the other RAN nodes, this is a physical node. The virtualisation of the O-RU is a topic for future studies [14].

2.2.4 Open RAN integration model

In a given RAN, the network operators have the largest share of costs in terms of CapEx (capital expenditure) and OpEx (operational expenditure). Therefore, efficient structures and processes for the deployment, integration and ongoing operation of a RAN are particularly relevant. With the introduction of the O-RAN architecture, integrating the various components of potentially different manufacturers becomes even more important – for successful commissioning and for secure operation and management of HW and SW upgrades. It is expected that in most cases, it will not be the network operator itself that will be responsible for the integration, but specialised system integrators (SI) or managed service providers (MSP).

The costs of different approaches are widely discussed, but the complexity and the relevance for security of the integration tasks are undisputed. New acceptance procedures and interoperability tests are necessary and may even have to be certified. In the operational concept, dedicated monitoring functions must be used that are exclusively directed towards the permanent verification of operational and data security of the distributed RAN components, especially if DevOps processes are used for automated software updates of the network functions.
To illustrate a possible scenario, Figure 6 presents a case of a CU provider and several DU and RU providers, each with different software versions that are to be integrated in an O-RAN, assuming that the software is installed on COTS (commercial off-the-shelf) hardware.

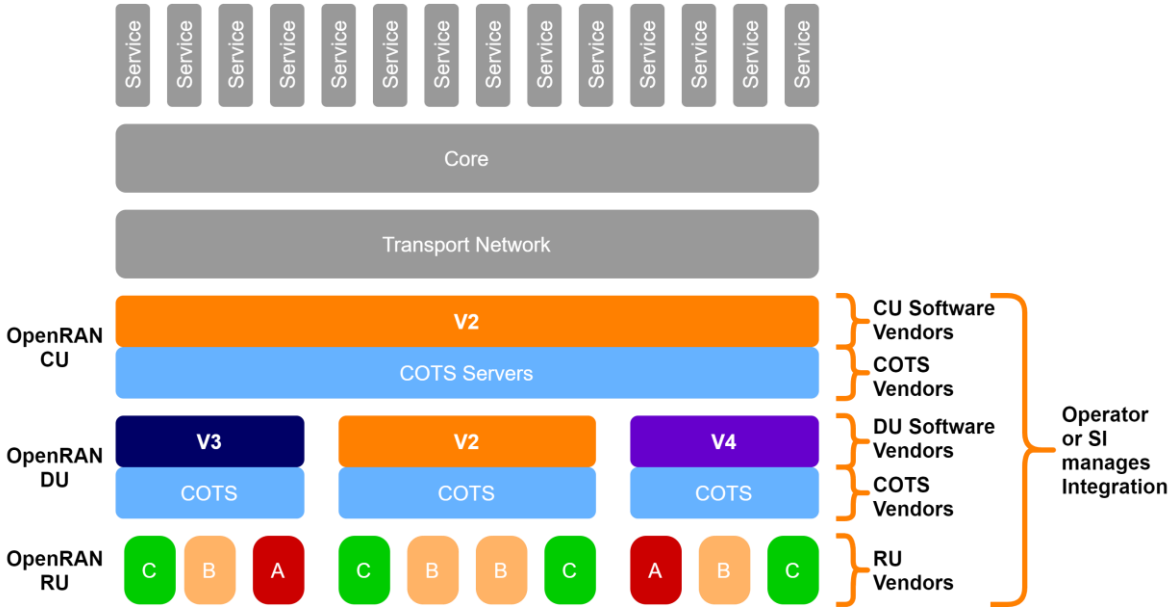

Figure 6: Integration models for an Open RAN managed by operators or system integrators ([source: Parallel Wireless], see [20])

It should be noted that in addition to the new NG-RAN functionality to be integrated, all the network components that are already in operation have to be considered in these models – including those of legacy technologies from 2G to 4G/LTE. Isolated campus networks are an exception, provided they do not include a legacy network.

2.2.5 RAN sharing concepts

One major distinction in RAN sharing concepts is between passive and active sharing. Passive sharing simply means the sharing of passive network elements. This includes the sharing of sites, masts, power supplies and cooling, as well as the connections between the sites and the respective concentration points all the way to the core network sites. It also includes site-related services and costs, such as security, fire prevention or property monitoring. The joint use of active network elements is called active sharing. Both of these concepts involve the shared use of electronic network elements.

2.2.5.1 MORAN and MOCN

Two RAN concepts [21] for active sharing are shown in Figure 7: Multi-Operator Radio Access Networks (MORAN) and Multi-Operator Core Networks (MOCN). In a MORAN, all components of the RAN (RAN system technology, antenna, mast, site, power supply) are shared by two or more operators. Each operator uses dedicated radio frequencies. In this approach, they can independently control the cell level (e.g. each operator can set its own optimisation parameters and transmit power to control cell range and interference).

With an MOCN, two or more core networks share the same RAN (i.e. the operators also share frequencies). The carriers are shared. Therefore, operators cannot control their networks at the cell level. The existing core networks can be kept separate. An MOCN is the most resource-efficient solution as it allows mobile operators to put their respective spectrum allocations into a common pool. The MOCN concept is specified by 3GPP for 5G in [9] and for previous mobile generations in [22]. It has been supported for UMTS since Release 6 and for LTE since Release 8; support for GERAN was added in Release 11.

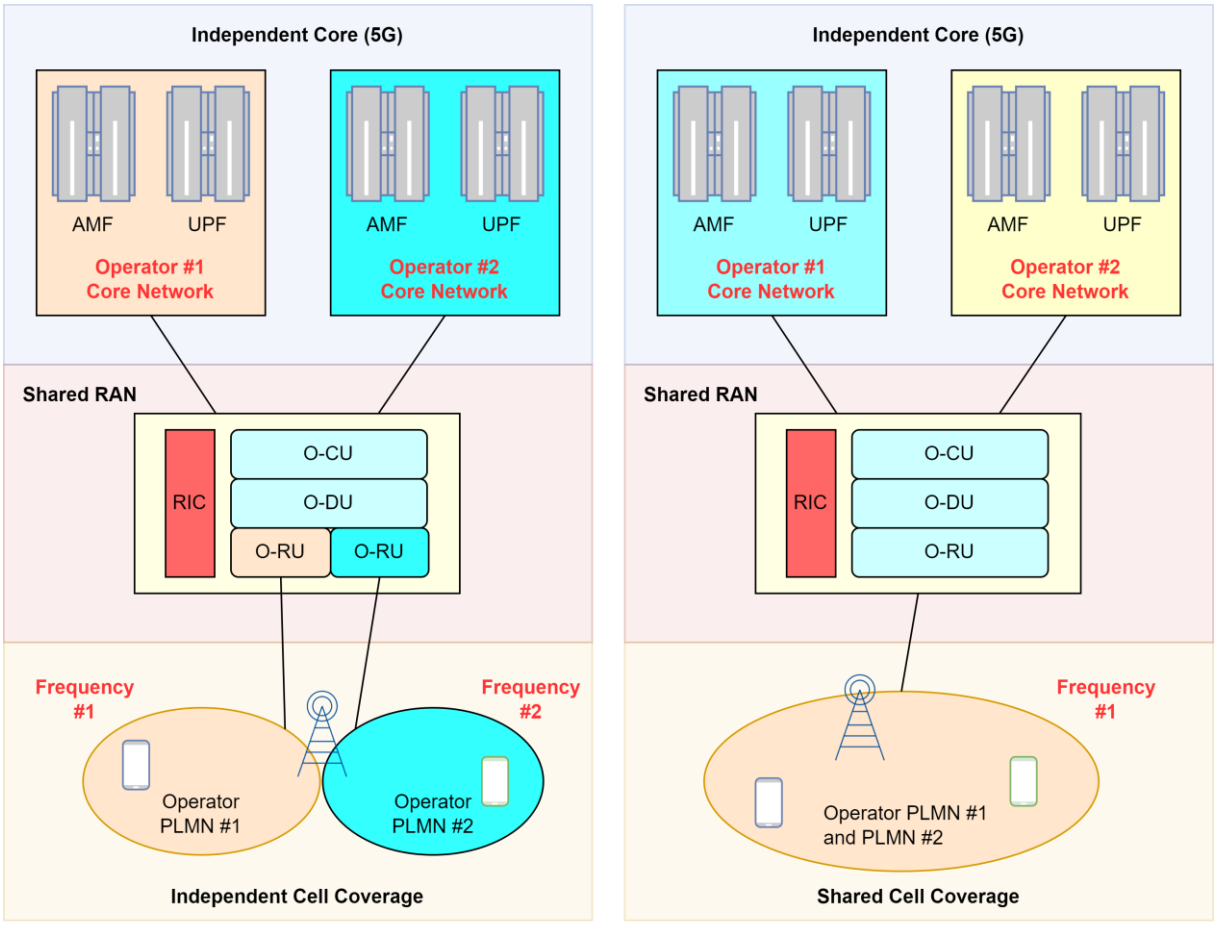

Figure 7: RAN sharing concepts MORAN and MOCN (taken from Techplayon)

In both approaches (MORAN and MOCN), mobile operators can choose to share transmission lines or to use them separately (in the same physical link).

2.2.5.2 Sharing concepts for O-RAN

There are many different designs for sharing the RAN components O-RU, O-DU and O-CU and the fronthaul, midhaul and backhaul transport layers. The decisive factor in each case is the placement of the individual elements in the RAN and the realisation in hardware, VMs or in the cloud (see Figure 8). The O-RU is always located at the antenna (radio) site. The server and software for the O-DU can be hosted at a dedicated site or in

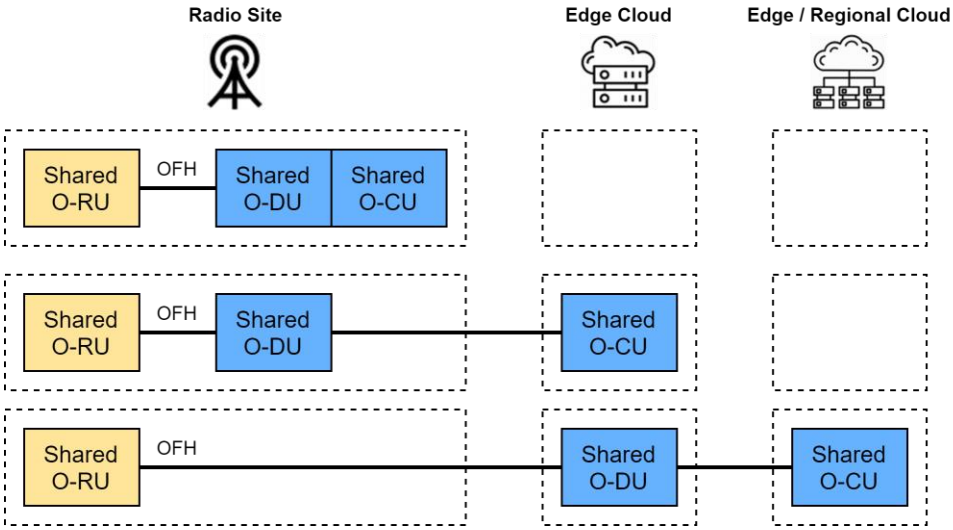

Figure 8: Main options for RAN sharing of O-RU, O-CU and O-DU

an edge cloud (at a regional data centre or the headquarters). The server and software for the O-CU can be co-hosted with the O-DU either at the radio site or in an edge cloud, or hosted separately at a regional cloud data centre.

A detailed compilation of possible scenarios is provided in the requirements document [23] prepared by Deutsche Telekom, Orange, Telefónica, TIM and Vodafone. Figure 9 shows scenarios 9, 10 and 17. In scenario 9, virtualised RANs (CU and DU) of different vendors are hosted on computing infrastructures provided at the radio sites (distributed vRAN on shared infrastructure). In scenario 10, the virtualised DUs are distributed at the radio sites, while the associated CUs are hosted on different edge clouds (distributed DUs, centralised CUs). Each edge-cloud may be a vendor-owned computing infrastructure. Scenario 17 represents a typical indoor scenario where, similar to scenario 9, the vRANs of different vendors are hosted at the radio site (or in an edge cloud). The O-RUs are connected to the DUs via a multiplexer.

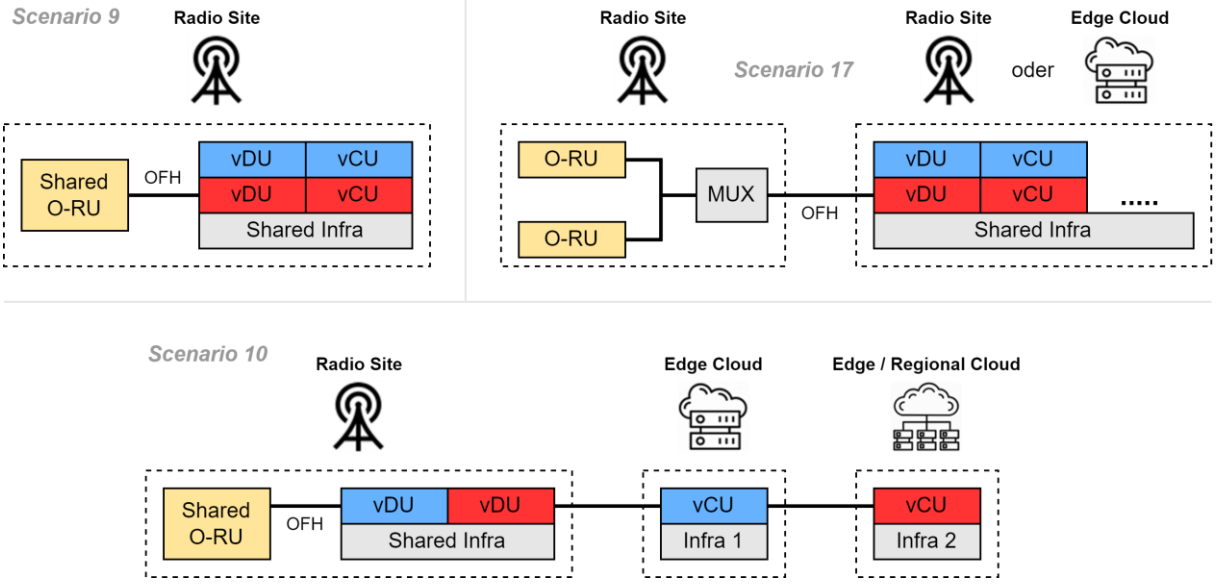

Figure 9: Selection of RAN sharing concepts: Distributed vRAN / Shared Infra (scenario 9), Distributed vDU, Centralised vCU (scenario 10) and Multi-Op Shared Infra (indoor scenario 17). [23]

Infrastructure sharing is preferred by network operators rather than the sharing of O-RAN units. Sharing the infrastructure for VMs or the cloud infrastructure is currently easier to coordinate than the joint management of O-RAN units. Configuration management for shared RANs has yet to be specified between operators.

For the user experience, RAN performance in terms of throughput, latency and reliability is usually key. Therefore, in terms of RAN sharing, RAN network slicing is a critical part of end-to-end network slices. The advantages are that slice-specific resource allocation, scheduling and admission control enable differentiated traffic handling and separation of subscribers or groups.

2.3 Description of O-RAN interfaces

This chapter explains the interfaces defined by O-RAN in more detail. Figure 10 shows the O-RAN architecture with all its interfaces, but without the interface to the UE (see Figure 5). In contrast, the fronthaul reference architecture for the O-RAN specification of

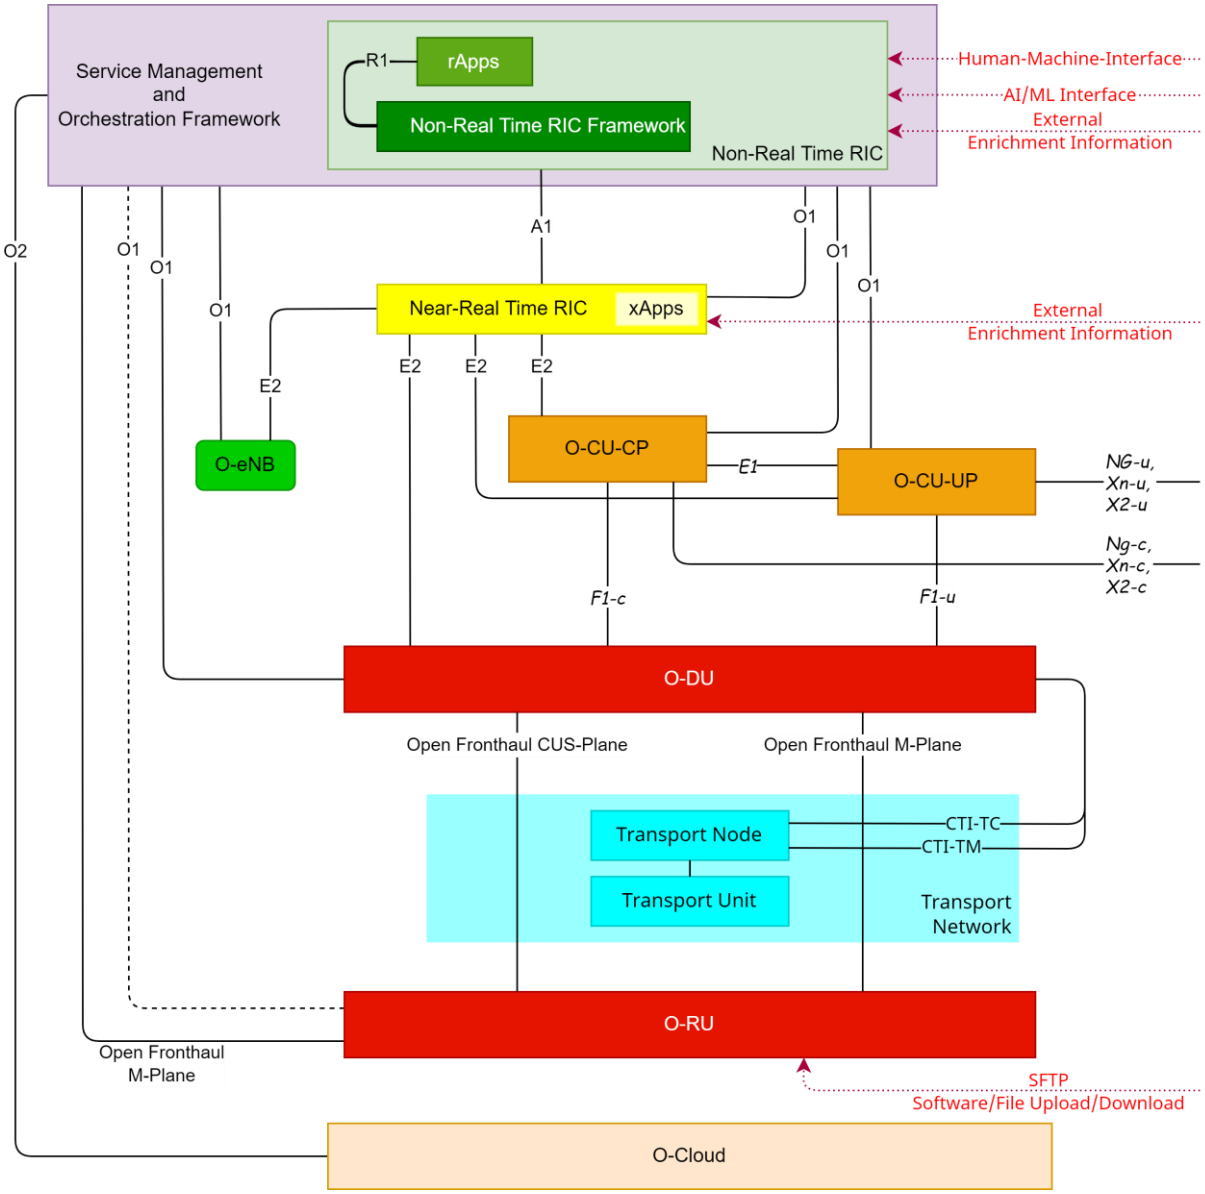

Figure 10: O-RAN interfaces

the CUS-Plane and M-Plane is extended by the Cooperative Transport Interface (CTI) and the transport network. The interfaces for human-machine interactions, AI/ML, external enrichment information and software updates are also added.

2.3.1 O1 interface

The O1 interface is the connection between all “O-RAN managed elements (MEs)” and the actual “management entities” of the SMO framework. The aim is to ensure the operation and management (e.g. FCAPS MGMT, software MGMT, file MGMT) of the O-RAN components via this interface. In other words, the O1 interface is used to enable the management of all O-RAN components that need to be orchestrated and the associated O-RAN network functions. The components managed via O1 include the near-RT RIC, the O-CU, the O-DU (in the case of 5G NR), and the O-eNB (in the case of O-RAN-compatible 4G/LTE networks). The O-CU corresponds to a predefined combination of O-CU-CP and O-CU-UP. Figure 11 shows an extract from the logical O-RAN architecture to illustrate the O1 interface and its influence on O-RAN MEs, including the ME O-eNB. The O-RU termination of the O1 interface towards the SMO is currently under investigation. For this reason, the interface connection is shown dashed.

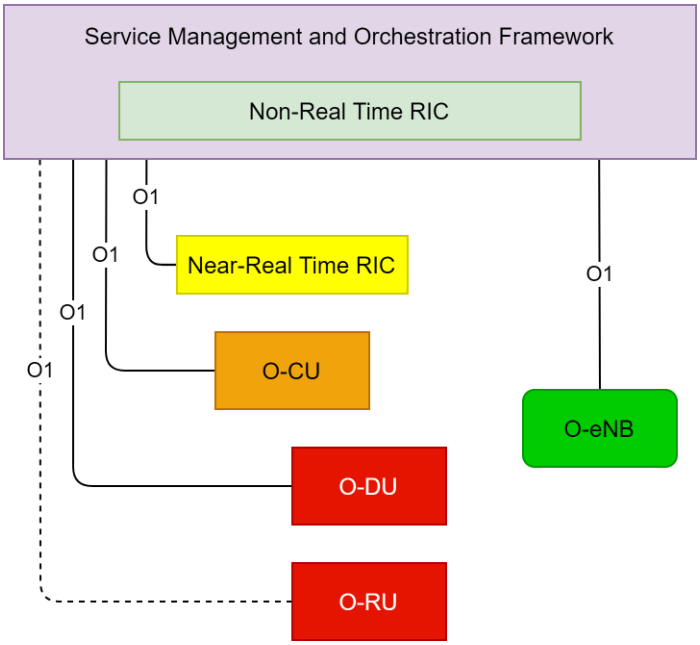

Figure 11: Network management and operation: O1 interface as connection between SMO and O-RAN managed elements (MEs), taking into account the ME O-eNB for O-RAN-compatible 4G/LTE RAN components (ref. [14])

O1 enables the SMO framework to access the O-RAN network functions. Network management is supported in line with the FCAPS model. FCAPS corresponds to the ISO model for network management, which describes and comprises the five functional areas of fault, configuration, accounting, performance and security management. Thus, the O1 interface plays a central role in the overall O-RAN architecture and network operation. O1 supports typical FCAPS and other management functions, including the following functions [24]:

- Discovery/registration
- Configuration of addressing
- Versioning
- Monitoring

In detail, the following management services (MnS) can be supported via the O1 interface (see also [18], [25]):

- Provisioning management services
  - General NETCONF requirements
  - Creating, modifying and deleting MOIs (managed object instances)

    - Reading MOI attributes
    - Notification of changes to MOI attribute values
    - Subscription control
- Fault supervision management services
    - Fault notification
    - Fault supervision control
- Performance assurance management services
    - File reporting and streaming of performance data
    - O-RAN-defined performance measurements and control of measurement jobs
- Trace management services
    - Call trace and streaming trace
    - Minimisation of Drive Testing (MDT)
    - Radio Link Failure (RLF) and RRC Connection Establishment Failure (RCEF)
    - Trace control
- File management services
    - File readiness notification
    - List available files and file download
    - Bi-directional transfer of files (e.g. configuration files for beamforming, certificates, ML files) between the client (File Management MnS Consumer) and the file server (File Management MnS Provider)
- Communication surveillance
- Heartbeat management services
    - Heartbeat notification
    - Heartbeat control
- Start-up and registration management services for Physical Network Functions (PNFs)
    - PNF plug-and-play
    - PNF registration
- Software management services for PNFs
    - Software package naming and content
    - Download, pre-check and activation of software
- Instantiation and termination of a virtualised network function (VNF)
- Scaling management services for VNFs

The management functionalities are realised by using standard protocols (e.g. SSH, TLS, NETCONF) and data models (e.g. YANG). For example, with the help of the provisioning management service, the SMO framework can receive information (updates) from the MEs via the O1 interface (e.g. on the current resource utilisation) and in return initiate an optimised configuration of the MEs.
In O-RAN-based mobile networks supporting AI/ML approaches, the O1 interface is used to collect (training) data that can be used for ML purposes from the MEs O-DU and O-CU. The (ML-based) non-RT RIC located in the SMO can offer policies to be considered in cell-level optimisation by providing (time-varying) optimal configuration sets for cell parameters via the O1 interface.
The O-RAN architecture allows the collection, access to and management of data records (history) relating to the traffic transferred over the RAN, the selected routing and the handover operations carried out. For this purpose, the data is transmitted via the O1 interface.
In the case of RAN sharing imagine the following scenario: an external operator (“guest operator”) gets access to the RAN infrastructure and computing resources of the actual network operator (home operator) by means of virtual RAN functions (VNF). Then, options for remote control and remote configuration of these VNFs must be made possible. In such a

scenario, “remote interfaces” (O1, O2) can also be introduced to enable the guest operator to transmit the desired configuration for the respective VNF at the home operator's location. The VNFs each represent a logical implementation of the O-CU and O-DU functionalities. The O-RAN architecture can provide a number of open interfaces as remote interfaces, including O1, to monitor the performance of remote users. This enables different optimisation strategies in terms of radio resource allocation and adaptation of QoS parameters.
There is currently a joint work item (JWI) in the O-RAN standardisation between the Working Groups WG1 and WG4 to define how exactly O-RU components can support management services in the RAN via the O1 interface. The decisions resulting from the JWI will be incorporated into future revisions and updates of the O1 interface specification.

2.3.2 O2 interface

The O2 interface is an open, logical interface within the O-RAN architecture and, like the O1 interface, is used as a vehicle for running open management and orchestration services. Its purpose is to ensure secure communication between the SMO framework and the O-Cloud platform. Depending on the deployment scenario selected within an O-Cloud instance, the O-Cloud platform can virtualise various network functions (NFs) and thus take over RAN functions within the overall architecture. The SMO framework provides the ability to manage a large number of O-Cloud instances in parallel. It is able to support the orchestration of available platform and application elements/resources, as well as workflow and workload management. The O2 interface is needed and used to implement these SMO tasks. It enables centralised management of the cloud infrastructure and cloud resource utilisation by the RAN, along with deployment life cycle management for the virtualised network functions (VNFs) running in the O-Cloud. The VNFs, the virtual machines (VMs) and container instances are to be managed via O2. However, the dependencies on cloud instantiation and the software applications (apps) to be executed in the could instances are not fully specified in the standard so far. The specification of the O2 interface, which is taking place in WG 6 (“The Cloudification and Orchestration Workgroup”), is only available in outline (see [26]).

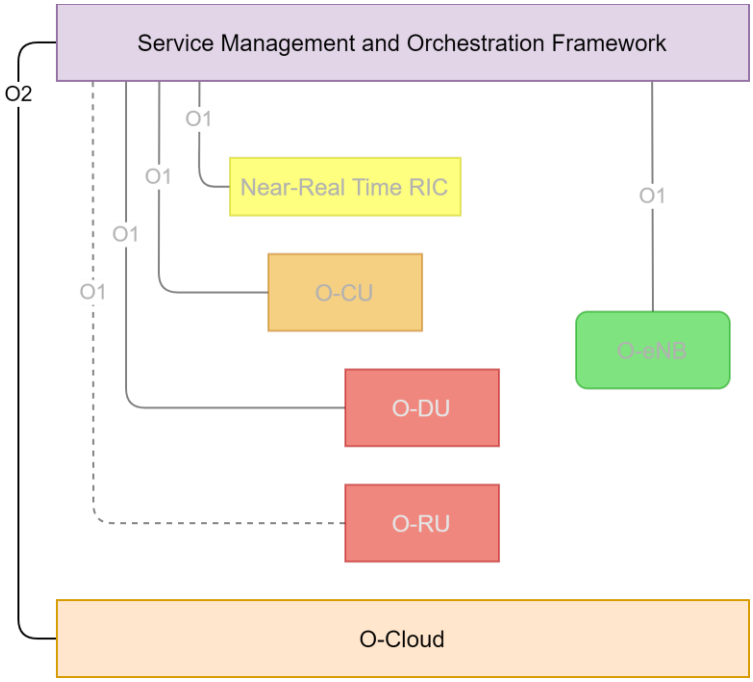

Figure 12: Management and operation of the O-Cloud platform by the SMO framework and influence of the O-Cloud on network operation: O2 interface as a connection between the SMO framework and the O-Cloud platform, taking into account the ME O-eNB for O-RAN-compatible 4G/LTE-RAN components (ref. [14])

To illustrate the O2 interface and its influence on the O-Cloud, it is highlighted in Figure 12 along with the MEs managed via the O1 interface (including the O-eNB).
Network operators can gain access to the network via the O-Cloud platform and operate and maintain the RAN using the O2 and O1 interfaces. For example, system updates/upgrades can be handled and network elements/MEs can be (re)configured.
The O-Cloud platform offers various services and functions to the SMO framework via the O2 interface. The current version of the O-RAN architecture description in [14] defines the following exemplary functions that are to be mapped via O2 and indicates that the scope of the functions does not have to be limited to those listed. In relation to the O-Cloud infrastructure, these functions are:

- Discovery and administration of O-Cloud resources
- Scale-in/scale-out of O-Cloud
- FCAPS (especially performance, configuration and fault management, as well as communication surveillance) of O-Cloud
- Software management of O-Cloud platform

In relation to software deployments on the O-Cloud infrastructure, these functions are:

- Create/delete deployments and associated allocated O-Cloud resources
- FCAPS (especially performance and fault management) of deployments and allocated O-Cloud resources
- Scale-in/scale-out of deployments and allocated O-Cloud resources
- Software management of deployments

In addition, the O-Cloud platform sends notifications via the O2 interface to the SMO framework when problems or changes are identified in the usage of O-Cloud resources. Furthermore, O2 should enable the management of hardware acceleration in the O-Cloud platform.
The lists above already show that the O2 functions can be categorised into two logical groups of services – one that addresses the O-Cloud infrastructure itself, and one that addresses software implementations (deployments) on the cloud infrastructure:

(1) Infrastructure management services (IMS)
(2) Deployment management services (DMS)

For IMS, the O2 interface provides functions that are responsible for the deployment and management of cloud infrastructures. For DMS, the O2 interface provides a set of interface functions responsible for the management of virtualised/containerised deployments on the O-Cloud infrastructure. O2 is therefore divided into two service-based interfaces (SBIs) between the SMO framework and the O-Cloud platform, each comprising its own set of functions. Figure 13 illustrates this division [26].

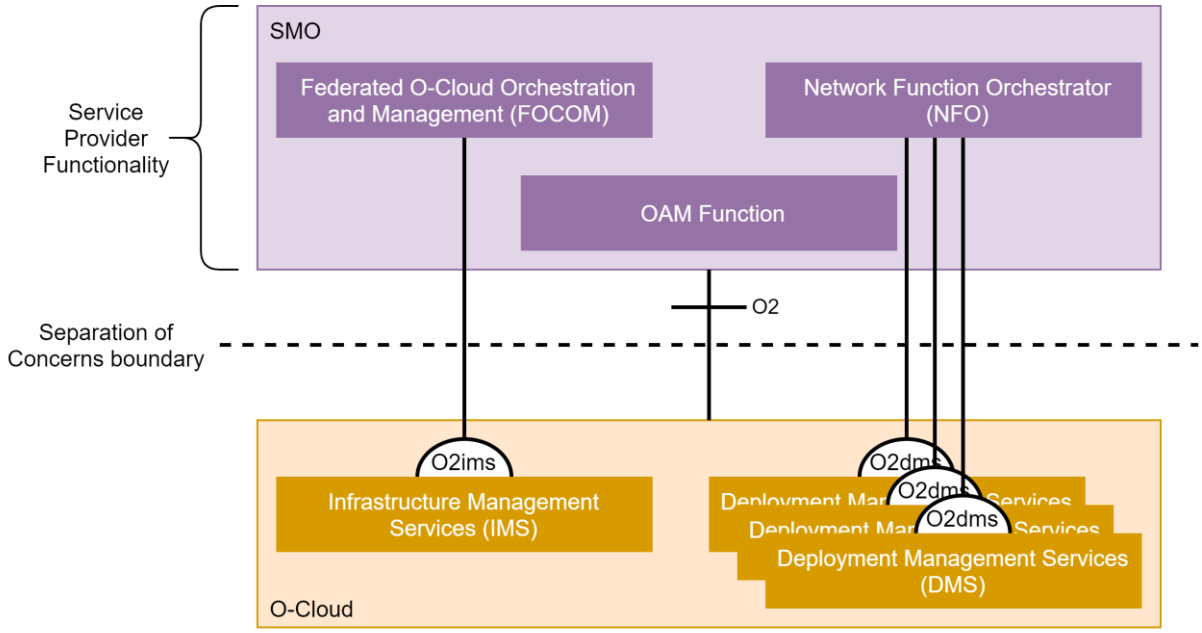

Figure 13: Division of the O2 interface into service-based interfaces (ref. [26])

Figure 14 is meant to provide a better understanding of the relationships and functions explained above in relation to the O2 interface in interaction with the relevant components of a RAN and the O-Cloud architecture. The overview shows the core components of an O-Cloud instance and the interaction between the SMO framework and the O-Cloud instances. This interaction takes place using the O2 interface and the O2 management services IMS and DMS.

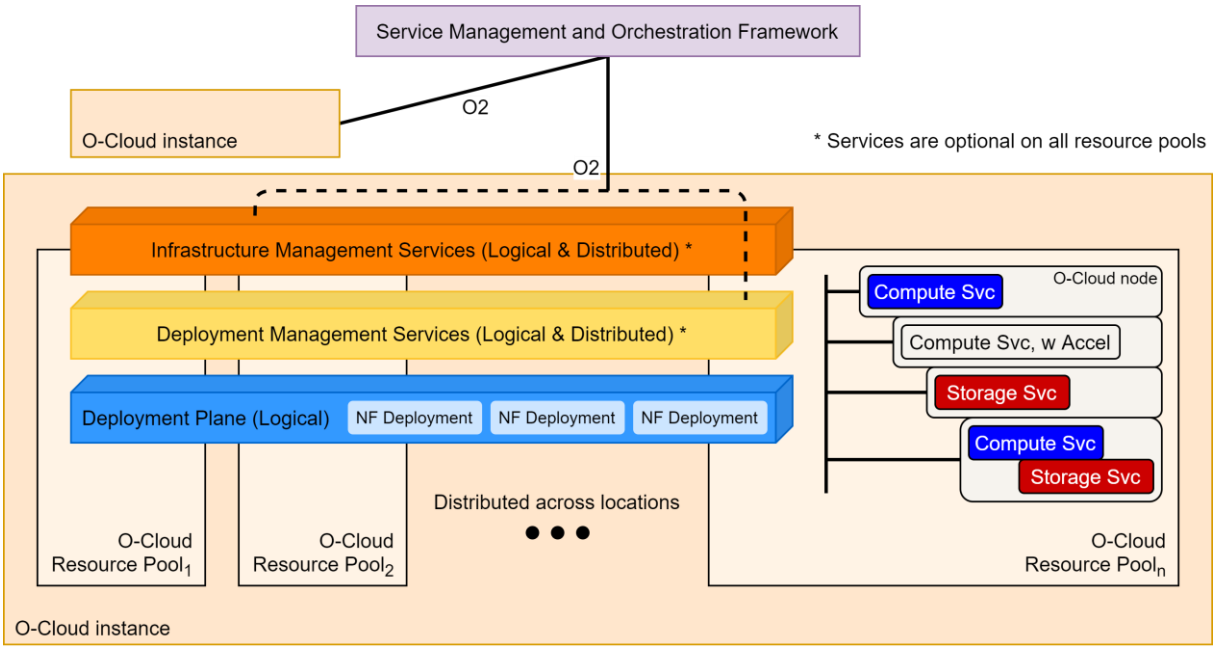

Figure 14: Interaction between SMO framework and O-Cloud instances and overview of O-Cloud core components (ref. [27])

Finally, it should be noted that the O2 services and the associated interfaces are still being specified. The writing and publication of a dedicated O2 specification is planned.

2.3.3 A1 interface

The A1 interface is being specified in the O-RAN Working Group 2 [28]. It is used for communication between the non-RT RIC and the near-RT RIC. To illustrate the A1 interface and its influence on the functions outside the RICs, Figure 15 shows the extract from the O-RAN architecture relevant for this chapter.

The non-RT RIC transmits the information captured in the SMO framework from various internal and external O-RAN sources to the near-RT RIC via the A1 interface. This information includes:

- Policy-based guidelines in declarative form (A1 policy) that contain statements on objectives and resources applicable to UEs and cells
- ML model management information (training, updating, deployment of ML models)
- A1 enrichment information from internal or external O-RAN data sources, where its availability or use is not critical for the task fulfilment of a unit, but only for its improvement

The near-RT RIC is to use this information to complete the configuration of the E2 nodes via the E2 interface. In this way, it will be possible to optimise a RAN under defined conditions (e.g. the RRM).

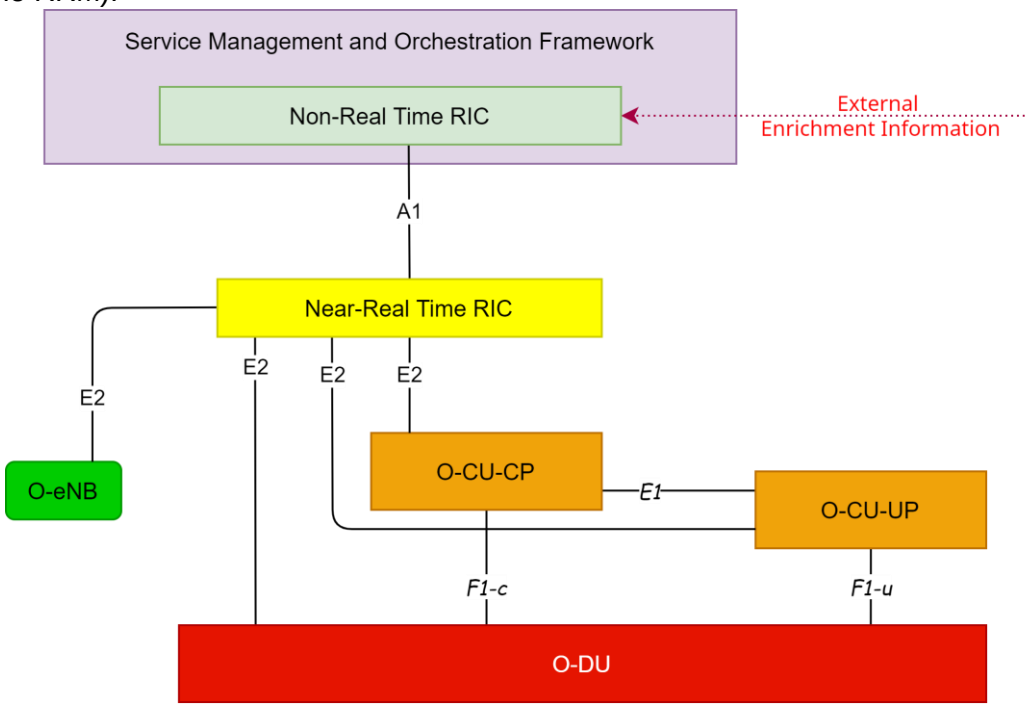

Figure 15: Extract from the O-RAN architecture: the A1 interface as a connection between the non-RT RIC and the near-RT RIC, taking into account the impact on other functions (ref. [14])

The declarative form of A1 policies means that the concrete implementation must take place in the near-RT RIC. The A1 policies are valid until modified or deleted by the non-RT RIC. The near-RT RIC has the task of informing the non-RT RIC about the status of the enforcement of an A1 policy by providing feedback via the A1 interface. It should be mentioned that the A1 policies are non-persistent, i.e. do not survive a restart of the near-RT RIC. Therefore, the task of the non-RT RIC is to check the presence of A1 policies.

If an A1 policy is related to a UE (or a group of UEs), the UE is identified by the UE Id. The UE Id is to be formed using the RAN UE Id known to the RAN and defined for the E1 interface [6] and the F1 interface [7]. The objective is to identify measurements associated with UEs in order to calculate correlations of O1-PM data with the objectives of a service (e.g. in the case of complaints about the service) and to assess the policy compliance. Neither hardware equipment nor user data is identified.

A1 enrichment information can be searched, requested and submitted via the A1 interface. When providing data from external sources, the non-RT RIC will be responsible for the source authentication and security of the connection. At this time, it has not been clarified how the connections to the external sources will be handled, nor is the search for and provision of external data defined.
For the differentiation of enrichment information, the standard defines the use of enrichment information types (EI types). However, apart from the generic term, no definitions of specific types exist to date.

2.3.4 R1 interface

The R1 interface is specified in O-RAN Working Group 2 in conjunction with the non-RT RIC [29]. Defined in the non-RT RIC, it is the connection between the inherent framework functions and the applications (rApps) running in the non-RT RIC. Figure 16 shows the R1 interface via the representation of the functions of the non-RT RIC. The functions in the non-RT RIC's framework include the management of rApps, support for rApps via access services (R1 Service Exposure Function), functions for the A1 interface, the AI/ML workflow, and potentially other functions of the non-RT RIC. Presumably due to its internal use as an API for rApps, the R1 interface is also called “Open APIs for rApps”.
rApps are designed to use the functionalities of the non-RT RIC exposed via the R1 interface to provide added value to the operation and optimisation of a RAN. These functions include:

- Providing policy-based guidelines and enrichment information via the A1 interface
- Performing data analysis, AI/ML training and data mining for RAN optimisation or for use by other rApps
- Recommending configurations that can be sent via the O1 interface

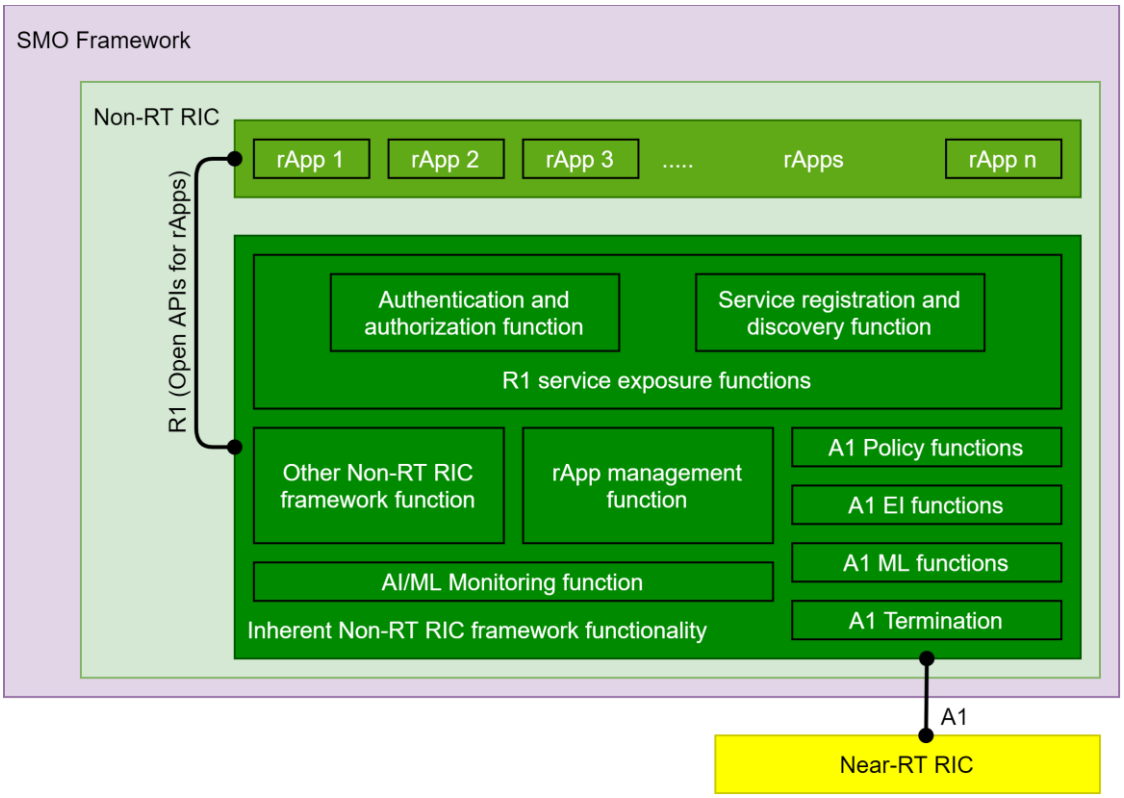

Figure 16: Illustration of the R1 interface using the functional view diagram of the non-RT RIC architecture (ref. [29])

As part of the SMO framework, the non-RT RIC has access to the SMO functions. This includes influencing the information transmitted via the O1 interface. The non-RT RIC needs this access to optimise RAN resources. However, the current O-RAN definition merely implies that the

non-RT RIC may only access the SMO framework functionality for this purpose. Accordingly, the non-RT RIC should only be allowed to influence transmissions via the O2 interface if the O-Cloud is considered a RAN resource.

2.3.5 E2 interface

The E2 interface, which is being specified in O-RAN Working Group 3 [30], connects the near-RT RIC with the E2 nodes. An E2 node is a collective term for all units to be controlled by the near-RT RIC, namely O-CU-CP, O-CU-UP and O-DU in the case of 5G NR and O-eNB in the case of 4G/LTE (see Figure 17). Accordingly, the E2 nodes are to support all the protocol layers and interfaces defined in 3GPP radio access networks, including eNB for LTE/E-UTRAN [17] and gNB/ng-eNB for NR/NG-RAN [19].

The near-RT RIC is connected to one or more E2 nodes via the E2 interface, i.e. in the case of NR to one or more O-CU-CPs, one or more O-CU-UPs and one or more O-DUs. Similarly, in the case of LTE, it connects the near-RT RIC to one or more O-eNBs. While a near-RT RIC has a one-to-many relationship with its E2 nodes, an E2 node can only have a one-to-one relationship with a near-RT RIC. Each O-CU-CP, O-CU-UP, O-DU and O-eNB can only be connected to one near-RT RIC at a time, just as a near-RT RIC can only be connected to one non-RT RIC.

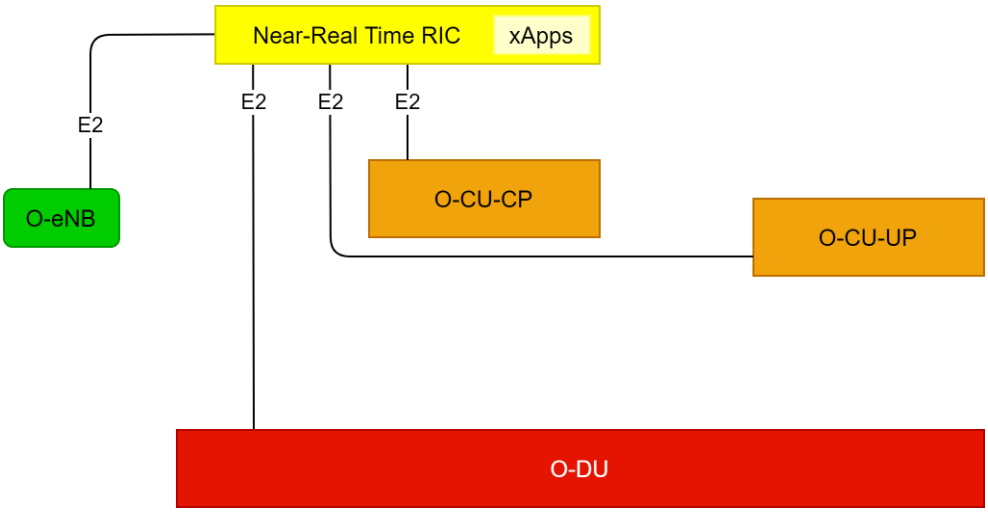

Figure 17: Extract from the O-RAN architecture: the E2 interface as the connection between the near-RT RIC and the RAN functions designated as E2 nodes (ref. [14])

As a critical interface between the radio nodes standardised by 3GPP and the functions specified by O-RAN, the first of the general principles listed in the specification states that this E2 interface must be open. This is an essential requirement for achieving interoperability. An important test of O-RAN support will be the extent to which the E2 interface is supported by RAN equipment vendors [4].

Another general principle in the O-RAN specification is that the functions of the near-RT RIC and the E2 nodes are fully separated from the transport functions. The addressing scheme used in the near-RT RIC and the E2 nodes shall not be tied to the addressing schemes of the transport functions. Furthermore, the protocols of the E2 interface are based exclusively on protocols of the control plane. Any kind of access to the user plane is not part of the specification. The goal is to control and optimise the E2 nodes and the resources they use. The xApps hosted in the near-RT RIC use the E2 interface to collect information in real time (e.g. on the basis of UEs or radio cells) to provide value-added services like the rApps do. For these purposes, the RIC services REPORT, INSERT, CONTROL and POLICY are used by the near-RT RIC to request measurement reports based on specific trigger events or to send new policies to E2 nodes.

Due to the one-to-many relationship between the near-RT RIC and the E2 nodes, the E2 interface must support failure handling and improved resilience. Even in the event of a malfunction on the E2 interface or in the near-RT RIC, an E2 node must be able to perform its function. However, failures may occur for certain value-added services that can only be provided via the near-RT RIC (e.g. when resource optimisation takes place on an O-CU with calculations performed by xApps and these calculations are based on measurement data regularly reported via E2). A cycle of this kind would not be able to continue in such a failure situation.

2.3.6 Open FH CUS interface

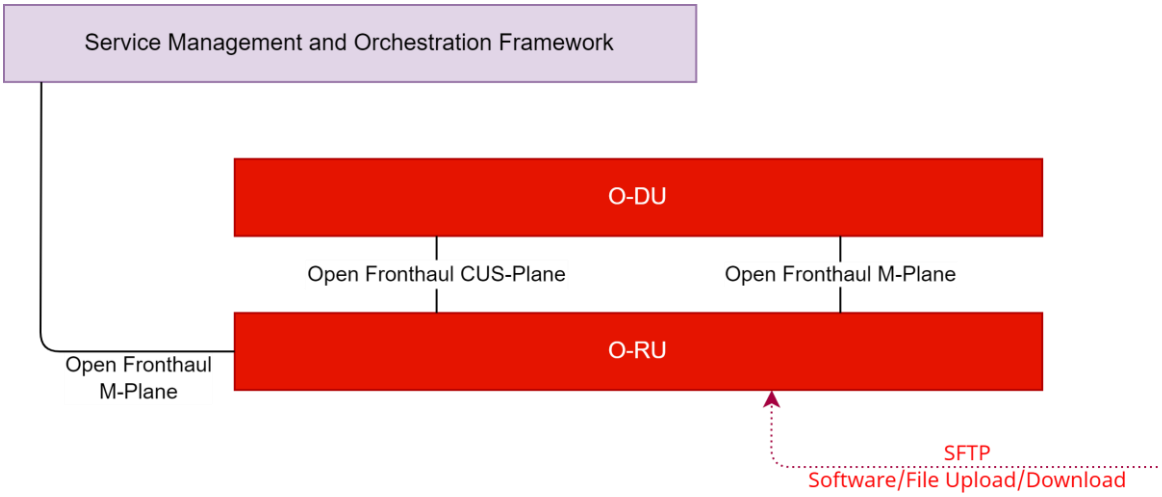

Figure 18: Excerpt from the O-RAN architecture: the Open fronthaul interface as a connection between the functions O-DU and O-RU or SMO and O-RU (ref. [14])

The Open Fronthaul Interface connects the functions O-DU and O-RU or, in the hybrid model, the SMO framework with the O-RU. The Open Fronthaul includes the CUS-Plane (Control, User and Synchronisation) and the M-Plane (Management). The M-Plane and its relation to the Open FH M-Plane interface is explained in chapter 2.3.7. Figure 18 shows the extract from the O-RAN architecture with the functions and interfaces relevant for the Open Fronthaul.
The interfaces for the C-, U- and S-Planes are described in the specification [31]. These interfaces are used to transmit the user and control plane data of the Uu interface and to carry out time synchronisation.

2.3.6.1 C-Plane (control plane)

Messages that determine the processing of user data are exchanged via the control plane. They include:

1. Information on scheduling or beamforming, if this information is not exchanged via the M-Plane (see chapter 2.3.7). These messages are transmitted separately for uplink and downlink.
2. UL- and DL-specific information on numerology, i.e. slot and subcarrier definitions
3. If pre-coding is used in the O-RU, the configuration data is transmitted from the O-DU.
4. Information for functions such as Dynamic Spectrum Sharing (DSS)

Either eCPRI or the IEEE 1914.3 protocol is used for C-Plane messages.

2.3.6.2 U-Plane (user plane)

Messages containing actual user data are transmitted via the user plane (also known as data plane). The focus is on efficient transmission, especially under the tight latency requirements in various 5G numerologies. The main functions provided are:

1. I/Q transmission of payload data, where each symbol is transmitted in a U-Plane message
2. Data compression, where different methods can be defined per physical resource block (PRB), which are specified in associated control messages
3. Downlink data precoding

It should be noted that I/Q data transmission is also possible without the C-Plane (e.g. via the Packet Random Access Channel, PRACH). In that case, the corresponding configuration must be exchanged via the M-Plane.

The supported methods for compression vary between O-RU and O-DU. It can be assumed that different O-RUs implement only one method in order to keep the complexity low. Accordingly, the O-DU must implement several methods to be interoperable with different O-RU manufacturers.

As on the C-Plane, either eCPRI or the IEEE 1914.3 protocol is used for U-Plane messages.

2.3.6.3 S-Plane (synchronisation plane):

The synchronisation requirements between O-DU and O-RU are an essential and critical part of the implementation of TDD operations, as well as for mMIMO or carrier aggregation across multiple O-RUs. Synchronisation can occur via the S-Plane using frequency, phase or common base time to align the clocks. Topologies for exchanging synchronisation information include:

- Network as master for the O-RU
- O-DU as master for the O-RU
- O-RU with local GNSS receiver as master

Protocols such as PTP and SyncE are used according to the O-RAN fronthaul specification. In case of loss of synchronisation by the O-DU, all RF connections of the connected O-RUs are terminated (“FREERUN State”).

2.3.7 Open FH M-Plane interface

According to the specification [32], the M-Plane provides the following major functionalities to the O-RU:

- “Start up” installation of procedures for commissioning
- Software management for upgrades in the operational phase
- Configuration management for initialisation and configuration of operating parameters
- Performance management and reporting via measurements and counters
- Fault management for configuration and transmission of alarms
- File management for uploads to the O-RU controller

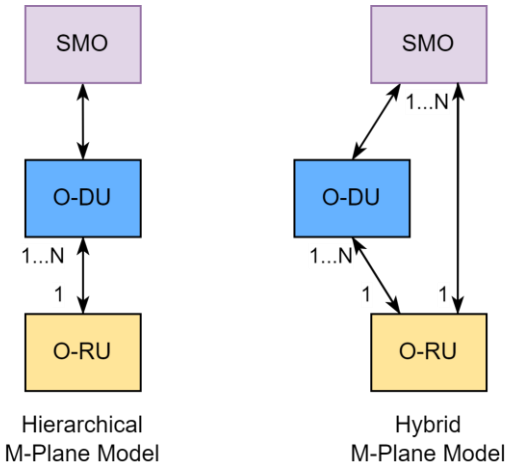

Figure 19: M-Plane architecture

The M-Plane interface can be implemented hierarchically or in a hybrid manner, as shown in Figure 19. In the hierarchical model, the O-RU is controlled by one or more DUs, (e.g. for redundancy). In the case of the hybrid model, there are simultaneous logical connections from the O-RU to the O-DU and to the SMO layer, possibly using the same physical connections. In the hybrid case, the functions for managing the O-RU can be shared between the O-RU controllers. For example, software management can be located in the SMO framework.
Typically, the configuration for the O-RU is performed initially as well as during operation, for which the following functions are used via the NETCONF protocol, for example:

- Uploading, committing and cancelling a new configuration
- Lock/unlock operations
- Rollback to a previous configuration in case of errors
- Notification of success/failure of actions

A number of market-standard functions are also available for the configuration and its triggering, as well as for the transmission of performance measurements and alarm notifications from the O-RU. These will not be discussed in detail here.

2.3.8 Cooperative Transport Interface (CTI)

CTI is an interface between the O-DUs and transport nodes of a packet-based transport network. It serves to connect the O-DUs with a large number of O-RUs [33]. CTI specifically targets transport nodes that manage a common point-to-multipoint access network. Transport nodes (routers and switches) that only manage point-to-point connections do not exchange CTI messages with the O-DUs. CTI consists of a Transport Control (TC) layer and a Transport Management (TM) layer.

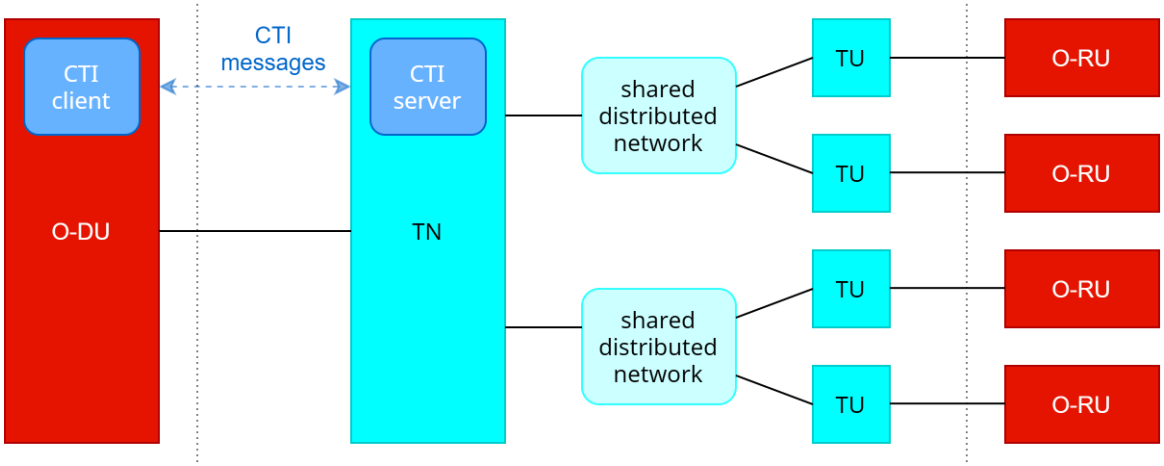

Figure 20: Relationship between CTI and the transport network used for the fronthaul (ref. [33])

CTI is intended for use in packet-based transport networks that contain transport nodes (TNs) terminating one or more common distribution networks, where each distribution network (a port on a TN) aggregates a large number of transport units (TUs; see Figure 20). This implies that the bandwidth of a distribution network is shared by multiple TUs. In the upstream direction, the TN manages the sharing by scheduling bandwidth allocations to the TUs. There are two ways to determine the allocated bandwidth: either statically or dynamically.

2.4 Optimisation aspects and machine learning

2.4.1 RIC functions for RAN optimisation

The RIC has been planned as an essential component in the O-RAN architecture. Functions embedded in the processing units of traditional base stations will be extracted in the course of the virtual split and become part of the RIC. This will enable access to the management interfaces, such as for RRM or SON functions that control radio resources and network operations.

To customise the RAN functionality, the RIC optimises the radio resources according to operator policies, thereby affecting RAN performance in three main areas [4]:

- Network intelligence: By measuring and reporting on the performance of the RAN, data is generated in standardised formats that can be analysed (e.g. using AI/ML techniques) to create new algorithms and policies.
- Resource assurance: The goal is to ensure that devices/users and services are provided with the required performance (e.g. by optimising radio link control, handover optimisation or prioritisation).
- Resource control: to ensure that the RAN system works efficiently when multiple user groups compete for adequate resources

In principle, it can be said that the intelligence envisioned by the O-RAN concept is located in the RIC. This intelligence is to be realised in the future through ML models for self-optimisation and radio network automation. The required data will be collected by the non-RT RIC from the RAN components in standardised formats via the O1 interface. First, traditional SON optimisation functions are addressed. In subsequent stages, this will enable model training by means of AI/ML so that new ML approaches for RAN optimisation can be developed and deployed. Based on this, the near-RT RIC derives policies and implements them in the CU-CPs and CU-UPs via the E2 interface. Alternatively, it applies dynamic controls in its E2 nodes. For example, certain devices (such as motor vehicles) may have high mobility requirements for which a different, superior handover management algorithm may be appropriate.

2.4.2 xApps/rApps

In addition to its function as a RAN controller, the RIC is an open platform that can host applications for RAN control. These applications are developed by specialised software providers that do not have to be part of the RIC provider itself. These “xApps” (in near-RT RIC) and “rApps” (in non-RT RIC) are intended to enable innovation in the form of RAN control algorithms and to attract software innovators to the mobile sector. xApps and rApps have the ability to process data from a RAN much faster than today's vendor-proprietary systems or centralised SON methods. As a result, a differentiated network experience can be created, with performance tailored to specific service types, user groups or locations.

Conversely, the services provided by a RIC consist of either xApps or rApps, or a combination of both. There is no definable limit on the types of xApps or rApps that can be built, and it is expected that more than one xApp or rApp will be executed in a RAN at a time.

The examples of xApps/rApps proposed so far include [4], [24]:

- Context-based dynamic handover management for vehicle-to-everything (V2X)
- Dynamic radio resource allocation for unmanned aerial vehicles
- Traffic steering
- Quality of service/quality of experience (QoS/QoE) optimisation
- Massive MIMO beamforming optimisation
- RAN sharing
- QoS-based resource optimisation
- Service assurance for RAN slices
- Multi-vendor slice performance management
- Dynamic spectrum sharing
- Optimisation of resource allocation for network slice subnet instances (NSSIs),
- Local indoor positioning in the RAN

[4] states that the first xApps will focus on “health check” functions (e.g. the operational readiness of RAN nodes). In a second phase, xApps will drill deeper into observations by capturing more granular data from RAN nodes for analysis. xApps that make changes in near-real time (i.e. time cycles of less than one second) are expected in later phases. In terms of control plane decisions, an initial approach will be to complement the current RRM functions implemented in the CU and DU, where policies from a RIC could change or override local RRM logic. A more aggressive approach of moving the RRM function entirely to the RIC is a longer-term exercise.

Initial rApps hosted in the non-RT RIC will resemble today's centralised SON applications at first. They have the potential to evolve rapidly when RAN data collection is paired with ML techniques to create algorithms that enable new forms of optimisation for the RAN.

One component of the near-RT RIC platform is the E2 Manager, which is sometimes referred to as “xApp zero”. It is used to initiate E2 connections with the RAN nodes and then store RAN configuration information that was learned during connection setup (and kept updated over time).

The first releases of the RIC specifications are available in [29] and [30]. O-RAN Working Group 3 will continue the development to add more features, such as the E2SM (E2 Service Model) specifications. The important work currently underway according to [4] includes the standardisation of E2SMs to enable traffic steering and QoS/QoE optimisation through the RIC. Traffic steering targets idle-mode mobility load balancing (MLB), inter-intra-frequency MLB, carrier aggregation and dual connectivity. QoS/QoE optimisation enables the RIC to control network functions related to QoS control, radio resource allocation, radio access control, mobility functions and connection management.

Operators such as AT&T, Deutsche Telekom, KDDI and China Mobile have publicly stated that they are testing RIC solutions [4].

2.4.3 Machine learning (ML)

The availability and high performance of AI and ML have driven the specification of architecture and processes for O-RAN in WG2. For network operations in particular, ML-assisted use cases are versatile and can therefore be deployed reasonably at various points in the O-RAN architecture. The location for the components of the ML workflow varies depending on the required response times, the availability and amount of data for ML training and ML inference (i.e. automated or computer-aided derivation of conclusions), as well as the computational complexity. The general workflow and components for ML processes are shown in Figure 21. The assignment of the ML components, the data flows and the actions regarding network functions is given as an example.

Initially, an implementation of the ML training host – primarily in the non-RT RIC – is expected, with data being transmitted via O1 and A1. In the non-RT RIC, actions are derived from the model inference and forwarded via the A1 interface (policies, configuration) to the near-RT RIC or via E2 (configuration) to O-CU, O-DU and O-RU. For processes and decisions that require shorter control loops, ML model inference is executed in the near-RT RIC to derive actions for the near-RT RIC itself or to transfer instructions/policies to the O-RAN components via E2. The intention for the future is that ML processes will also be specified for the O-DU/O-RU control loop.

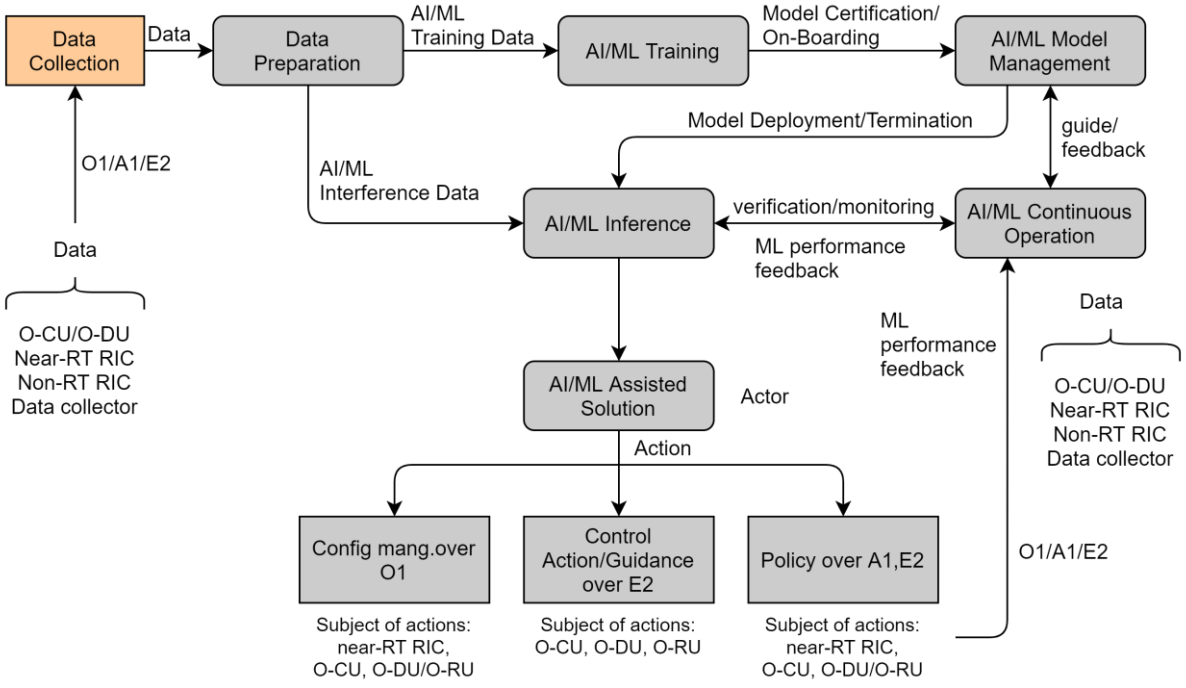

Figure 21: General procedures in AI/ML [34]

Training data is gathered by data collectors from all O-RAN components via E2, A1 and O1 and provided to the training process via aggregation and filter components. Particularly relevant is the data from the E2 interface being made available in different granularities regarding location and time. Enrichment data can be provided via all available interfaces or internally in the SMO layer. The latter is of importance in the security considerations, given that functions in the SMO layer can also have proprietary interfaces to raw data sources.

A typical application for ML-assisted processes is RAN optimisation. Due to the large amounts of data usually fetched for such a task and the resulting complex processing of these data sets, the non-RT RIC is suitable for both training and inference. In addition, enrichment data drawn from external operations support systems (OSS) can be correlated.

For secure continuous operation of the ML processes, a number of functions must also be available that monitor the components, processes and actions with regard to activity,

performance, timing, resource consumption and consistency. KPI definitions are usually utilised for this purpose, on the basis of which changes may have to be triggered. These tasks could be assisted by ML, as well.
Based on the corresponding feedback from the model inference host, decisions can be made in the SMO layer, including about the following:

- The model selection to be changed
- Additional training for the model
- The scheduling of the model

System manufacturers and third-party providers are increasingly developing ML-based algorithms that are in operation on O-CU and O-DU and generate actions in the range of

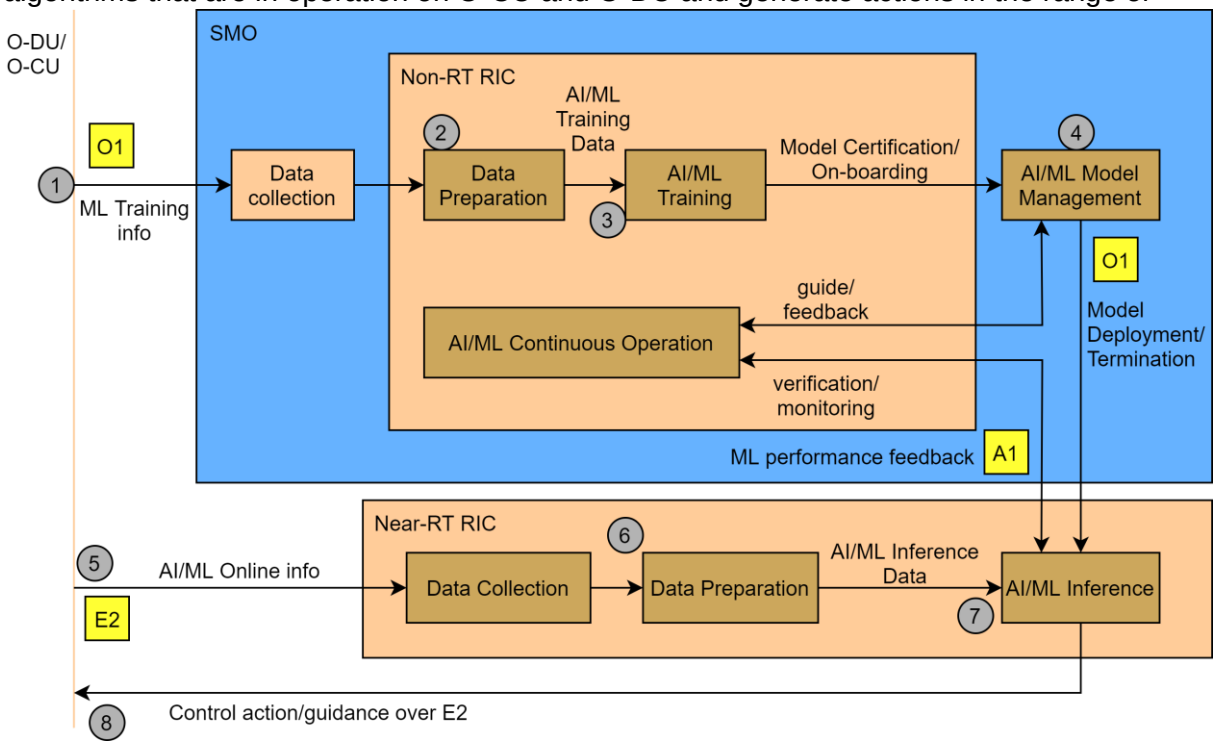

Figure 22: Deployment scenario for AI/ML [34]

milliseconds in order to react with optimisation measures in a UE-specific manner (i.e. they have a direct effect on the user experience and functional security). This makes it all the more important for O-RAN architectures to further develop not only new integration and acceptance procedures, but also operational concepts that address internal processes of the operator in a multi-vendor environment. As an example, a deployment scenario is given in Figure 22 in which the ML components and ML processes are embedded in the O-RAN functions as follows:

- Model management in the SMO
- Preparation of training data, training and the processes for continuous operation in the non-RT RIC
- The ML inference process in the near-RT RIC

2.5 O-RAN software

In addition to the specifications, an open-source reference implementation for the operation of an O-RAN is being developed by the O-RAN Alliance in cooperation with the Linux Foundation. The organisation responsible for the development is the O-RAN Software Community (OSC)[4].

[4] https://O-RAN-sc.org/

According to the white paper “O-RAN Minimum Viable Plan and Acceleration towards Commercialization” published on 29 June 2021 [35], the goal of the OSC developments is “*to achieve a solution that can be utilised for industry deployment*”. To achieve this goal, reference implementations are being developed for the various O-RAN components. These developments are coordinated by the Requirements and Software Architecture Committee (RSAC) and the Technical Oversight Committee (TOC)[5].

The TOC has 12 voting members. Currently, 10 of these positions are filled. The TOC members come from telecommunications companies (AT&T, China Mobile, Deutsche Telekom, NTT Docomo, Orange, TIM), equipment vendors (Ericsson, Nokia) and suppliers (Radisys, Wind River Systems). The meeting minutes (including recordings of the online meetings) are publicly available on the TOC website.

In addition, synergy effects with other projects are being sought. To be mentioned here, in particular, are the developments within the framework of the ONF[6] and the Open Network Automation Platform (ONAP)[7].

The plan is to publish a new release every six months. The releases are marked with letters and corresponding words. The fourth release (D, Dawn) is currently being published.

In principle, participation in the open-source software development is open to anyone interested. Anyone wishing to participate must agree to the Contributor License Agreement (CLA). Contributions are made under the Apache licence, version 2.0[8].

For source code management, `git` is used. `Gerrit` is used for the management of change proposals. The corresponding web page accessible to the public can be found at: https://gerrit.O-RAN-sc.org/

[5] https://wiki.O-RAN-sc.org/display/TOC
[6] https://opennetworking.org/
[7] https://onap.org/
[8] https://www.apache.org/licenses/LICENSE-2.0

3 Methodology and Scope

This chapter provides an overview of the methodology used for O-RAN risk analysis in this study. Existing procedure models for risk identification are described very briefly. Key principles and assumptions are also presented.

3.1 General information and scope

The risk analysis carried out in this study is limited exclusively to the 3GPP RAN in its implementation variant O-RAN. The considerations made are thus limited to the RAN interfaces to the user equipment (Uu interface) and to the 5G core. Risks for the 5G system as a whole that result from insecure user equipment (including user equipment related components such as USIM) and from an insecure 5G core do not play a role in this study.
Furthermore, this study only considers a 5G RAN (NG-RAN). In particular, this means it is assumed that the RAN is connected to a 5G core.
The information on which the risk analysis is based comes from publicly available documents. This applies in particular to the documents provided by 3GPP (standards, reports, etc.) and those provided by the O-RAN Alliance. In each case, the current, publicly accessible versions of the documents in the period in which this study was prepared (May–September 2021) were used. When referring to the standards and specifications used, the exact version of the respective document is cited. Any deviating systems actually used in practice or concrete implementations of the standards and specifications were not taken into account. Similarly, this study does not factor in the actual operation of a RAN by a network or RAN operator. It is therefore possible that a radio access network rolled out in practice is actually exposed to fewer security risks than those identified by this study because the operator may have implemented additional security measures to minimise risk.
The risk analysis essentially only considered threats and vulnerabilities that are specific to 3GPP or O-RAN. Generic IT risks that are more generally attributable to the field of information and communication technology (such as faulty implementations or configurations) are only explicitly mentioned and considered in a few instances. This is mainly because a large number of risk analyses exist for the IT sector that will not be repeated here, but which must of course be taken into account in the overall risk assessment of a specific RAN deployment.

3.2 Risk analysis methodologies

In the context of this study, a risk is considered to be an "*effect of uncertainty on objectives*" [36]. Since the "effect on objectives" is understood as damage, positive effects on objectives are not considered in the context of this study. "Uncertainty" refers to the probability of specific events occurring that would result in damage (sources/causes of risk [36]). To describe the level of a given risk in quantitative or qualitative terms, the usual formula is used:

$$\text{Risk} = \text{Likelihood} \cdot \text{Impact}$$

There are a number of procedure models regarding risk identification (ISO 27005 [37], ISO 31000 [36], IEC 31010 [38], BSI-Standard 200-3 [39], etc.). The procedure is often similar and follows the following steps:

1. Determination of the attacker to be considered
2. Determination of sensitive assets
3. Determination of the criticality of failures to meet protection goals with regard to the assets (i.e. the damage that could occur)
4. Determination of threats to protection goals and assets
5. Identification and assessment of vulnerabilities in relation to the threats determined under consideration of existing security measures
6. Determination of risk on the basis of the vulnerabilities and the potential damage

7. Assessment of risks, including the planning of measures to deal with them

The following chapters describe the protection goals and attacker model that are relevant to the study. They also explain the different scenarios considered. On this basis, the concrete procedure model chosen for this study is described.

3.3 Protection goals considered

In this study, the usual three protection goals of **confidentiality, integrity** and **availability** are considered, along with **accountability** and **privacy**. Brief explanations of the significance of these protection goals in the context of this study are provided below.

Confidentiality means that data and information can only be accessed by authorised persons.

Integrity means that data and information are complete, correct and current, or that this is detectably not the case. In the latter case, this means in particular that unauthorised manipulation of data and information can be detected. The formulation "complete, correct and current" refers to the perspective within the IT system under consideration; in other words, it is explicitly not considered whether the data and information are complete, correct and current with regard to the real world.

Accountability means that actions (e.g. the sending of data) can be verifiably attributed to a given entity (e.g. the sender) and proven to third parties. It should be noted that the risk analysis with regard to confidentiality, integrity and accountability does not take into account security measures that may be taken within a given use case at the application layer; such measures are taken outside the 5G RAN system under consideration.

Availability generally means that data, information and services are available to authorised persons where and when they are needed. In the context of this study, this protection goal also includes unauthorised impairments of quality of service (QoS) and an increase in the costs associated with providing data, information or services – for example, an increase in delay time, a reduction in throughput or an increase in energy consumption. Analyses regarding the protection goal of availability are of particular importance because generally speaking – and in contrast to the protection goals of confidentiality, integrity and accountability – the enforcement of this protection goal from the application perspective is only possible at considerable cost, or not at all if the communication system used (in the case of this study: 5G) is not reliable in terms of availability.

In this study, the protection goal "*privacy"* includes protection goals that are usually associated with data protection, such as anonymity, unlinkability or unobservability. The considerations regarding privacy are thus concerned with the extent to which metadata is generated that could be used to violate the confidentiality of the communication circumstances in question. This could include metadata that makes it possible to link several communication processes and thereby profile communication behaviour (frequency, duration, location, etc.). Furthermore, metadata in connection with location information can be used to create movement profiles or generally determine who communicates with whom (although those involved may be pseudonymised). Like availability, privacy is a protection goal that cannot be achieved (or only with considerable effort) without corresponding support from the communication network in question.

3.4 Attackers considered – attacker models

Potential attackers and the capabilities they are assumed to possess are an essential part of carrying out a risk analysis of this kind. These aspects are usually summarised in an attacker model. In the context of this study, five different attacker models are considered based on the following roles (see Figure 23):

- **Outsider**: An attacker who can only carry out attacks using the interfaces defined in the system; they initially have no control over components involved in the system. With regard to 5G RAN, this means that the individual can attack via both the wireless air interface and the interfaces specified by 3GPP or the O-RAN Alliance. It is assumed

that the attacker has full control over the transport medium used in each case; in other words, they can eavesdrop on all exchanged data and manipulate it as desired (modification, delay, deletion, generation, etc.) through both the radio connection and the (IP-based) connections used between the 5G or O-RAN components.

- **User**: An attacker who is an end user of the 5G system, meaning they have control over one or more instance of UE that can legitimately use services of the 5G system. In the context of this study, this attacker is also assumed to have the capabilities of the "outsider". The "user" attacker essentially differs from the "outsider" in that they possess and are capable of using credentials/secrets that are necessary for the legitimate use of the 5G network.
- **Cloud operator**: An attacker who has physical and logical control over the (edge) cloud infrastructure used by the 5G RAN. This affects all cloud components that are not explicitly 5G RAN components (as specified by 3GPP or the O-RAN Alliance) and includes both hardware and software components. The attacker also has all the options available as the "user" attacker.
- **Insider**: An attacker who has control over exactly one 5G RAN or O-RAN component and also has the capabilities of the "user". This attacker is particularly interesting as a means of investigating whether O-RAN's more granular separation into components can be a security benefit compared to the (at least conceptually) rather monolithic 3GPP-RAN. As detailed below, the risk analysis is first carried out individually for the relevant O-RAN interfaces and then used as a basis for a summarised overall assessment. For the risk analysis of individual interfaces, it is assumed in each case that the insider controls a component that is connected to the interface (i.e. has access to the interface). This is essentially a question of the extent to which an insider's additional knowledge or rights present new risks compared to the "user" attacker.
- **RAN operator**: An attacker who has full control over the 5G RAN. This attacker is particularly interesting as a means of assessing the risks posed by a compromised RAN. A "RAN operator" attacker possesses capabilities beyond those of the "insider" and "user" attackers.

All the attackers considered are based on the following common assumptions:

- The availability of considerable, although not unlimited, resources (computing power, storage space, money, etc.). This is intended to cover cases involving state-supported attackers or financially strong cybercriminals.
- Active, modifying attackers who are prepared to break rules outside the parts of the system they control (e.g. by manipulating transmitted data).
- Cryptography is secure, meaning it is assumed that the attackers are not able to break cryptographic algorithms and protocols that are currently considered secure.
- Cryptographic secrets are secure, meaning it is assumed that the attackers initially have no knowledge of cryptographic secrets (cryptographic keys, etc.) that they do not already know due to their role (see above).

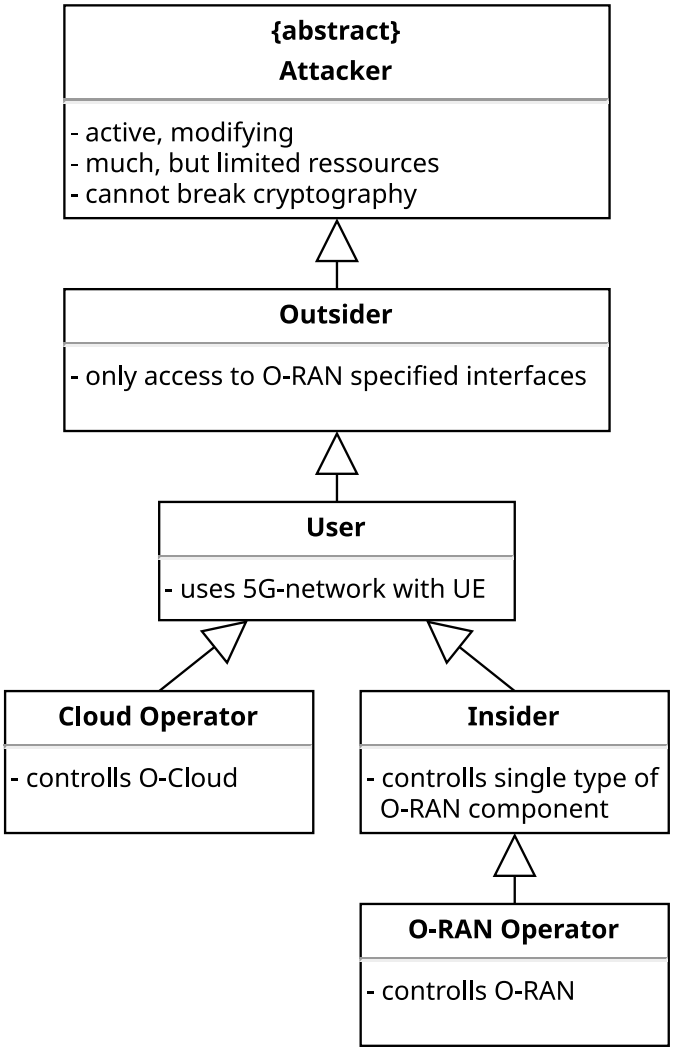

Figure 23: Hierarchy of the attacker models considered

3.5 Perspectives

The analysis of the risks regarding 5G RAN and its implementation O-RAN takes into account different perspectives that are explained in more detail below. This concerns different stakeholders and the implementation of security measures.

3.5.1 Stakeholder perspective

With the help of the stakeholder perspective, the risk analysis seeks to factor in different interest groups, each of which has specific interests and requirements with regard to the security of the 5G RAN or the 5G system as a whole. This enables a differentiated consideration of the 5G RAN risks.

3.5.1.1 End user

An end user is understood as an entity that has one or more devices (UE) connected to a 5G network and uses the services of the 5G network in question accordingly. The primary interest of the end user is assumed to be that the protection goals listed above (confidentiality, integrity, accountability, availability and privacy) are met with regard to the transmitted user data (including voice communication) and the use of the 5G network. Since no specific type of end user is assumed, the assurance of security properties concerns all types of 5G services and all use cases in which a 5G network is involved. It must therefore be assumed that any failure to meet one of the protection goals listed above is critical and can potentially cause very

significant damage. This applies in particular to the protection goals of availability and privacy, since an end user can only counter a corresponding failure with considerable effort.

3.5.1.2 Societal / governmental perspective (state)

As detailed in the introduction, mobile networks already represent important communication infrastructure. Their importance is likely to grow significantly in the coming years, with critical services or critical infrastructures (electricity, water, health care, logistics and transport, etc.) also becoming increasingly dependent on functioning 5G infrastructure. As a result, the security, trustworthiness and reliability of 5G networks are of great social relevance. This should be taken into account by governmental stakeholders.

Since this perspective is based on a large number of concrete – and also crucial – use cases, it must be assumed with regard to the societal/governmental perspective that any failure to meet the mentioned protection goals is very critical, as well. A failure of availability is considered particularly relevant here.

3.5.1.3 5G network operator / telco

A 5G RAN represents an essential component of a 5G system. Depending on the operator model, the operation of a 5G RAN may have been transferred to a third party and may therefore not be under the direct physical or logical control of the 5G network operator. In particular, 5G components (hardware, software) may be shared by different 5G network operators (RAN sharing). At the same time, a 5G RAN has the intended interfaces to the 5G core network.

It is therefore necessary for a 5G network operator[9] to know the risks posed by a potentially untrusted 5G RAN to the overall operation of the 5G network and to the associated assets. It should be noted that the 5G network operator perspective primarily considers risks related to the Control Plane; risks related to the transmitted user data are secondary, as corresponding analyses are already carried out within the framework of the end-user perspective.

3.5.2 Implementation of security safeguards

The information on which the risk analysis is based is essentially taken from the standards and specifications published by 3GPP and the O-RAN Alliance. The risk analysis is therefore not based on a concrete 5G RAN implementation with precisely specified properties. The standards and specifications rather provide a certain framework within which compliant implementations of a 5G (O-)RAN can operate. Of particular relevance to these studies are the intended security measures and mechanisms and whether their implementation is mandatory or merely optional. In order to adequately consider the optional security measures in particular in the risk analysis, two scenarios are assumed:

- **Worst case**: none of the optional security measures have been implemented.
- **Best case**: all the optional security measures have been implemented.

It should be noted here that only the security measures that are at least mentioned as optional in the standards and specifications are taken into account. It is assumed that all the conceivable additional security measures that a 5G RAN operator could implement have not been implemented.

3.5.3 Summary

The three stakeholder perspectives listed above each analyse the 5G RAN risks from a different angle. In simplified terms, the end-user perspective analyses *User Plane* risks, the 5G network operator perspective focuses on *Control Plane* risks and the government perspective combines *User Plane and Control Plane* risks. The *worst-case / best-case* considerations reflect extreme circumstances with regard to the protective measures implemented.

[9] In the context of this study, the term “telco” is also used synonymously with the term “5G network operator” for better readability.

3.6 Methodology applied for risk analysis

In accordance with the explanations given in the previous chapters, the methodology explained below was applied within the framework of the risk analysis. With regard to the risk formula, it should be noted that the risk calculation only takes into account the likelihood of occurrence, as the impact depends on the specific 5G use case and, generally speaking, can be very high (see chapter 3.5.1).

With regard to the *likelihood of occurrence*, no precise quantitative analysis is carried out because no related data is available and precise quantitative analysis also depends on the use case in question. Within the framework of the study, qualitative assessments are made on the basis of three levels:

- **High**: A high probability of occurrence exists if it is possible for a given attacker to exploit a vulnerability with little effort and thereby exploit a risk that leads to corresponding damage. This applies in particular if the exploitation of a vulnerability is within the scope of the capabilities clearly attributed to the attacker. For instance, sensitive data may be transmitted unencrypted via an interface or a transport medium to which the attacker has access according to the attacker model. It should be noted that additional efforts, such as those related to the spatial distribution of the attacker, are not taken into account here. For example, an unsecured radio interface in a cellular mobile network could mean that an attacker would have to expend a great deal of effort to conduct area-wide mass surveillance, which could be used as an argument for an overall assessment as more of a medium risk. Nevertheless, the risk of targeted attacks on individual devices or geographically limited regions remains high. One particular use case here involves attacks on campus networks. Furthermore, the prerequisite of accessing transmitting/receiving equipment distributed over a large area does not necessarily mean that this equipment must be physically set up by the attacker. Rather, the attacker can exploit insecurity in end devices (UE) to gain access and then use them for attacks. In this context, it is particularly important to consider that the number of end devices will continue to grow strongly in the future – with a high proportion of insecure end devices likely to come from the IoT sector, for example. Since a meaningful consideration of the spatial distribution aspect of an attacker is only really feasible for a concrete use case, this (possibly risk-reducing) aspect is not considered within the scope of this study.
- **Medium**: A medium likelihood of occurrence is assumed if a risk can be exploited in principle according to the attacker model, but generally only with considerable effort on the part of the attacker if no further assumptions are made. A “user” attacker could, for example, connect a large number of devices (which may be located in different geographical positions) to the 5G network in order to overload the network.[10]
- **Low**: A low probability of occurrence is assumed if the exploitation of a given risk is beyond the capabilities defined in the attacker model, or would only occur with negligible probability or extreme effort. An example of the former might involve the undetected manipulation of data transmitted securely with the help of cryptographic procedures; an example of the latter would be the guessing of secret cryptographic keys.

The risk analysis presented herein is thus about assessing how likely a breach is for each attacker (chapter 3.4), perspective (chapter 3.5) and protection goal (chapter 3.3).

[10] In some of the existing risk analyses on 5G, this attack scenario is associated with the Internet of Things. An attacker (hacker) gains access to a large number of IoT devices connected to the 5G network and then uses them to attack the network (especially its availability). This example also illustrates the difficulty of assessing how costly it actually is for an attacker to exploit risks, and that it depends on many factors. In this specific case, the attacker does not need their own devices for the attack because they can make use of third-party devices.

It should be noted that the analysis of the probability of breaches of protection goals takes into account not only known threats and vulnerabilities, but also the risk-reducing security safeguards to be applied in accordance with standards and specifications. Here, the overall risks related to 5G RAN are a combination of the threats and security safeguards implied by the 3GPP standards and the threats and security safeguards implied by the O-RAN specifications. The best-case/worst-case considerations introduced in chapter 3.5.2 are applied in the process. It also should be noted that the best-case/worst-case considerations can be applied to the 3GPP standards and the O-RAN specifications respectively. This results in four possible combinations:

- Best-case assumptions regarding both 3GPP and O-RAN (abbreviated as *bb*)
- Worst-case assumptions regarding both 3GPP and O-RAN (*ww*)
- Best-case assumptions regarding O-RAN combined with worst-case assumptions regarding 3GPP (*bw*)
- Worst-case assumptions regarding O-RAN combined with best-case assumptions regarding 3GPP (*wb*)

In particular, best-case and worst-case assumptions have been combined to determine the extent to which the O-RAN specifications have an influence on the overall risk with regard to the 5G RAN. The results can be summarised in **Fehler! Verweisquelle konnte nicht gefunden werden.**.

Attacker	**Perspective (stakeholder)**														
	End user					State					Network operator				
	Protection goals					**Protection goals**					**Protection goals**				
	C	**I**	**A**	**Z**	**P**	**C**	**I**	**A**	**Z**	**P**	**C**	**I**	**A**	**Z**	**P**
Outsider	*bb* *bw* *wb* *ww*														
User															
Cloud operator															
Insider															
RAN operator															

*Table 1: Scheme for the overview of the risk assessment. The cells in the table reflect the likelihood of occurrence of a breach of the protection goal with regard to a given attacker and a given perspective (stakeholder) in the best case (**b**) or worst case (**w**). Green means low; yellow, medium; and red, a high likelihood of occurrence, while a white field means that no statement is possible at present. The abbreviations for the protection goals represent confidentiality (C), integrity (I), availability (A), accountability (Z) and privacy (P).*

In order to prepare a (summary) risk analysis for an (overall) system, in some cases (especially with regard to O-RAN) a risk analysis is first carried out with regard to the individual components of the (overall) system. These individual assessments are then combined in the risk determination for the (overall) system in the sense of a “per security” strategy. This means that the overall risk with regard to a stakeholder, attacker and protection goal is usually at least as high as the highest individual risk with regard to this stakeholder, attacker and protection goal. If additional risks arise from the composition of the components, the overall risk can also be higher than each individual risk. In this case, appropriate justifications are given in the description of the risk analysis for the (overall) system. It should be noted that a reduction of the overall risk cannot occur through composition, since the analysis of the individual components already takes into account their role in the (overall) system (in particular the consideration of safeguards as explained above). For example, if only end-to-end secured data

is exchanged via an unsecured interface, a low risk (in terms of confidentiality and integrity) is assumed.
Finally, it should be pointed out again that the concrete assessments regarding high, medium or low risk cannot be “mathematically proven” and often contain a certain subjective evaluation. This also results from the fact that not every conceivable scenario could be analysed in detail with the time and resources available. This applies, for example, to the “insider” attacker, which represents an oversimplified generalisation in that a more detailed examination of each individual O-RAN component could determine which security risks arise if precisely the O-RAN component in question is malicious. The reader is therefore required to consider the textual comments on the individual risks. In case of doubt, these are authoritative with regard to the risk assessment.

4 Existing Studies

A number of studies exist that deal with threats, vulnerabilities and risks with regard to the 5G system as a whole and its individual components, such as the 5G RAN. A selection of these is summarised below. Particular mention should be made of the threat analysis published by the O-RAN Alliance.

4.1 ENISA Threat Landscape for 5G Networks

In December 2020, the European Union Agency for Cybersecurity (ENISA) published an update of its threat analysis [40] of 5G from November 2019 [41]. The current threat analysis presents a variety of threats to 5G systems and groups them by 5G component and 5G function. There is also a threat analysis specific to the 5G RAN, but it does not include O-RAN. The main threats identified are:

- Degradation of quality of service with respect to ultra-reliable low-latency communication (URLLC) use cases. This threat is briefly mentioned, but it is unclear from which concrete and detailed threats the assessment is derived.
- Interference with radio transmissions by means of jamming.
- "*Failure to meet General Security Assurance Requirements: a set of weaknesses will arise through the update requirements of various elements of RAN due to implementation of migration steps and the ability of early-deployed systems to comply with specification updates regarding security functions.*" It is still unclear what exactly is meant here. It is also not clear how this threat was derived from the threats listed in detail. It is assumed that reference is being made to the fact that, in practice, newer systems must also be downwardly compatible with older, insecure standards and that vulnerabilities are therefore also found in these newer systems and components.
- Attacks on the F1 interface due to security safeguards that are only optional

In addition, a variety of more or less general threats are listed that are therefore – although specifically illustrated for the 5G RAN – not specific to 5G RAN. These include, for example, implementation errors (in the various components, including the hardware/software used), configuration errors, failure to apply intended (cryptographic) security safeguards and unauthorised access to secrets used, for example, in the context of cryptographic security measures (secret cryptographic keys). The identified threats are listed in Annex E of the ENISA document, although it is noticeable that most of the stated threats are also relevant without the 5G context (for example: "*Improper authorisation and access control policy: The authorisations for accounts and applications shall be reduced to the minimum required for the tasks they have to perform.*")

4.2 EU 5G risk analysis

In October 2019, the NIS Cooperation Group published a report on 5G risk analysis [42]. The report is based on a survey of the EU member states with regard to the assessment of 5G risks. It only lists very general risks; a more detailed threat analysis will be provided by ENISA in the report described in chapter 4.1.

The risk for the Radio Access Network was assessed as "high", the second-highest level (after "critical"). Concrete information on how this assessment was reached cannot be found in the report.

With regard to the general risks dealt with in the report, the lack of diversity and the associated dependence on a few manufacturers is particularly relevant for the present study, since Open RAN in particular can be understood as an approach to countering this risk. In addition, the risk analysis refers at various points to the risk posed by state attackers who could try to manipulate manufacturers of 5G components into supplying malicious 5G components. Here, too, Open RAN could be a sensible component of the countermeasures.

In addition, general risks mentioned in the report, such as misconfigurations and insufficient access control, are also relevant for the Radio Access Network and for O-RAN in particular.

4.3 EU Toolbox

Based on the reports on risks and threats related to 5G that are mentioned in chapters 4.1 and 4.2, the NIS Cooperation Group presented a proposal on general countermeasures to mitigate the risks in January 2020 [43]. This proposal is known as the "EU Toolbox". The countermeasures are divided into three groups: strategic, technical and supportive.
The strategic countermeasures are particularly relevant in the context of this study:

- *SM05*: Diversity of manufacturers of 5G components
- *SM08*: Diversity in future network technologies; building EU expertise in this field

Increasing diversity and reducing dependency on individual manufacturers is mentioned in many different places and contexts, which underlines the role of diversity as a very important countermeasure.
As mentioned above, Open RAN can contribute to the implementation of these two strategic countermeasures. In this respect, it is important to evaluate the new risks posed by Open RAN/O-RAN in order to ultimately assess whether Open RAN/O-RAN is part of the solution or part of the problem and determine what measures must be taken so that the former is the case.
The 11 technical countermeasures recommended in the EU Toolbox essentially concern the full application of standard IT security measures. It is therefore also important to analyse the extent to which O-RAN implements these recommendations regarding technical countermeasures.
With regard to the supporting countermeasures, measure *SA03* ("Participation in 5G standardisation") should be mentioned here, as one of the goals it pursues is to increase diversity by creating well-defined interfaces.

4.4 US studies and reports

In the US, there are a number of initiatives and reports that address 5G risks and corresponding measures. They include:

- "CISA 5G Strategy" [44]
- "Potential Threat Vectors to 5G Infrastructure" [45]
- "National Strategy to Secure 5G" [46]
- "National Strategy to Secure 5G Implementation Plan" [47] (includes annexes published as a separate document [48])

In addition to the "usual" risks and threats, which can also be found in a large number of documents from other countries and organisations, two types of risks and associated measures stand out that are also related to the investigations of this study:

- Supply chain risks, especially because the US is currently dependent on international manufacturers for the 5G RAN (the "National Strategy to Secure 5G Implementation Plan" names the following manufacturers: Huawei, Ericsson, Nokia, Samsung and ZTE, whereby Huawei and ZTE are classified as untrustworthy and thus effectively excluded as suppliers)
- Risks due to the influence of standardisation

To minimise these risks, a large number of countermeasures are described. Taken together, they serve the following two goals:

- Increasing diversity in 5G component manufacturers, ideally with a high proportion of US manufacturers
- Leadership in the field of standardisation in order to influence the development of standards according to US interests

In turn, Open RAN/O-RAN is seen as a concrete way to achieve the two goals mentioned above or to promote their implementation.

4.5 "The Prague Proposals"

"The Prague Proposals"[11] is a document that resulted from the first Prague 5G Security Conference[12]. It contains some general recommendations regarding the security and roll-out of 5G networks. It should be noted that the recommendations are so general that they are not restricted to 5G networks, but apply in principle to any IT infrastructure ("*Communication networks and service should be designed with resilience and security in mind*."; "*Stakeholders should regularly conduct vulnerability assessments*."; "*Risk management framework... should be implemented*", etc.) Furthermore, supply chain risks and the need for diversity in terms of manufacturers of 5G components are emphasised. The Prague Proposals are worth mentioning mainly because they are referenced as a basis in a number of US documents and reports (see chapter 4.4).

4.6 O-RAN security threat modelling and remediation analysis

As mentioned previously, the O-RAN Alliance recently (July 2021) published its own document dealing with threats and risks related to O-RAN [49]. The analyses are based on a procedure in line with ISO 27005. While binding measures for risk minimisation are not defined in the current version of the document, this is planned for a future version.

The document does list numerous threats, some of which were also considered in the risk analysis conducted in this study. As the authors of the O-RAN Alliance study are aware, however, many of the threats listed are also of a more general nature and not (O-)RAN specific. This means that only a subset of the threats is O-RAN specific, and there is little focus – at least at present – on the Open Fronthaul interfaces (CUS, M-Plane) and other interfaces (O1, O2, etc.)

A major difference between O-RAN's own risk analysis and the one presented here is their understanding of what a risk is in the first place. In the current O-RAN document, only the "impact" of a risk is identified, rather than the likelihood of its occurrence. In this study, the opposite approach is taken: only the likelihood of occurrence is considered. However, a future version of O-RAN's own risk analysis will also consider the likelihood of occurrence

The impact classes are classified as low, medium and high, and the number of compromised components (O-DUs, O-RUs, etc.), the severity of the threat to the respective protection goals (privacy, confidentiality, integrity, availability) and the impact on possible synchronisation topologies (clock model) are taken into account.

As a result, the O-RAN Alliance identifies 32 threats as high, 16 as medium and 5 as low in terms of their impact. It remains to be seen how the O-RAN Alliance will assess the likelihood of occurrence[13] of these threats and the extent to which the other existing interfaces (O1, O2, etc.) will be examined.

4.7 GSMA mobile telecommunications security landscape

The report [50] published in March 2021 by the GSMA (Global System for Mobile Communications Association) dealt with what it considered to be important (and in some cases, new) changes in the security landscape of the mobile industry. The report covered the following subjects:

- "Signalling & Inter-connect"
- "Supply Chain"
- "Software & Virtualisation"

[11] https://www.vlada.cz/assets/media-centrum/aktualne/PRG_proposals_SP_1.pdf
[12] https://www.prague5gsecurityconference.com/
[13] These are already being discussed in internal ORAN documents that have not yet been published, which is why no reference is made to them here.

- “Cyber & Operational Security”
- “Security Skills Shortage”
- “Device & IoT”
- “Cloud Security”
- “Securing 5G”

Most of the risks derived from these thematic areas are mitigated by general measures in the report, such as the advice given to mitigate the risks from the area of “Cyber & Operational Security”: *“Good security practices can mitigate this risk through secure networks, strong authentication, least privilege practices alongside strong privileged access management (PAM).”* [50] However, these only represent general guidelines (as already mentioned) and are therefore less suitable for concrete threat prevention, which must be much more specific to the domain at hand.

With regard to 5G, the report points out that although 5G has closed many vulnerabilities through its architecture, the corresponding security safeguards have not always been fully implemented in practice. In this regard, the authors point out that most 5G architectures do not yet have a 5G core: “*At present Non Stand Alone deployments are not making full use of the standards based security, as much of this only comes when a 5G core (5GC) is deployed.*” [50]

As in the previous studies, the report also highlights the risks pertaining to supply chains, not least because of increasing national interventions: *“In 2020, we saw an increasing trend towards national responses to supply chain threats.”*

As a general guideline, the GSMA recommends examining components from different manufacturers individually with regard to the risks they pose. In particular, it recommends making concrete plans to remove a specific vendor and its component from a network should this prove necessary: *“Build business continuity plans that consider the removal of critical vendors; understand the impact if one were to be removed.”* [50] The GSMA also recommends considering open network solutions and trying them out in test environments.

In general terms, none of the risks presented in the report are specific to 5G RAN. They do not include any concrete risks presented by interfaces or components in a 5G network, for example. The guidelines mainly deal with risks that are generally found in many infrastructures, even outside of mobile networks.

5 O-RAN Risk Analysis

The following chapters consider the risk analysis in terms of O-RAN. In this context, chapter 5.1 summarises “cloud operator” and “RAN operator” attackers with regard to O-RAN as a whole. This is mainly due to the prominent position of these two entities/attackers.
For the other attackers, the risk analysis is first carried out separately with regard to the individual interfaces and basic modules specified in O-RAN. An overall risk analysis is then derived from this in chapter 5.15. In each case, the highest risk from the respective individual considerations is taken as the overall risk (see also chapter 3.2), since from the attacker's point of view it is sufficient to be able to successfully exploit a vulnerability in the system for attacks. Simply put, the weakest link in the chain is the decisive one with regard to a risk analysis.
It should also be noted that the results of the risk analysis regarding 3GPP RAN (Appendix A:) have been incorporated into the risk analysis for O-RAN, since O-RAN (as an implementation of a 3GPP RAN) “inherits” some of the security risks resulting from the 3GPP standards along with the positive effects of the 3GPP security measures.

5.1 Attackers: cloud operators and 5G RAN operators

The “cloud operator” and “5G RAN operator” attackers have prominent positions. Both ultimately have full control over the RAN: the 5G RAN operator by design, and the cloud operator because the current O-RAN specifications do not provide for security measures regarding untrusted cloud operators. It simply assumes that cloud operators can be trusted: “*Administrators, integrators, operators and orchestrators must be trustworthy, ...*” [49].
As a result, there are no differences in the O-RAN-related best-case/worst-case considerations with regard to “cloud operator” attackers. The same applies to “RAN operator” attackers, who, among other things, have full control over the O-RAN security safeguards. Therefore, only worst-case (*ww*) and best-case (*bb*) considerations are generally made below.
Worst case: If none of the 3GPP or O-RAN security safeguards specified as optional are implemented, the risk of a breach is to be assessed as high for all protection goals and for all stakeholders, since in this case the cloud operator/5G RAN operator has full access to all processed data and full control over the data processing itself. The only exception here is the end-to-end protection of the Control Plane data with regard to integrity. This improves the situation for the 3GPP Control Plane data, but the internal O-RAN configuration and management data are unprotected, which results in a medium risk overall.
Best case: Even the best-case situation is only negligibly better. For all stakeholders, failure to meet the protection goal of availability represents a high risk. From their perspective, end users must also assume a violation of the protection goals of confidentiality, integrity, privacy and accountability, since the cryptographic keys for securing the air interface (Uu) are known to the RAN. Since such an attack is easy to carry out across the board, a high risk is also seen with regard to “state” stakeholders and the violation of the protection goals of confidentiality, integrity, accountability and privacy. For the “telco” stakeholder – specifically from the Control Plane perspective – the situation is slightly better, as the corresponding Control Plane messages (NAS) between UE and the core network are end-to-end secured (between UE and the 5G core). However, manipulations of the O-RAN-specific configuration and management data are possible, resulting in the assumption of a medium risk overall.

Attacker	Perspective (stakeholder)																	
	End user					State						Network operator						
	Protection goals					Protection goals						Protection goals						
	C	I	A	Z	P	C	I	A	Z	P		C	I	A	Z	P		
Cloud operator																		
RAN operator																		

Table 2: Risk assessment for the attackers "Cloud operator" and "RAN operator" with regard to the entire O-RAN architecture

5.2 O-Cloud risk analysis

The O-Cloud is the central execution environment of the O-RAN components. As mentioned above, an untrustworthy cloud operator results in very high security risks. The same applies in the case of an O-Cloud compromised by an attacker. This is why effective security safeguards such as access control and separation are crucial. The O-RAN specifications only make a few stipulations here and even contain security-critical requirements.

For example, the “Security Requirements” specification [51] contains only a recommended requirement for access control: “*User should be authenticated and authorized*.” On the positive side, isolation mechanisms certainly are required: “*Means of isolation of control and resources among different users shall be implemented*” – although such isolation mechanisms are ineffective without mandatory user authentication.

In the O-RAN worst case, a risk situation arises with regard to “outsider” attackers that is comparable to the situation involving “RAN operator” attackers, since a compromised O-Cloud is essentially equivalent to full control over the O-RAN components.

The best-case scenario from an O-RAN perspective cannot be assessed at present, as there are practically no security requirements for the O-Cloud.

Attacker	Perspective (stakeholder)																	
	End user					State						Network operator						
	Protection goals					Protection goals						Protection goals						
	C	I	A	Z	P	C	I	A	Z	P		C	I	A	Z	P		
Outsider																		
User																		
Insider																		

Table 3: Risk assessment of the O-Cloud

5.3 O2 interface risk analysis

As described in chapter 2.3.2, the O2 interface is used to configure the O-Cloud and deploy the RAN components of O-RAN, i.e. the VNFs. The O2 interface can be considered very powerful in that it can be used to determine the entire execution environment and the software components that are executed. This is all the more true since the deployment scenarios and use cases found in the O-RAN documents provide for a great deal of flexibility, and the O2 interface must therefore be powerful enough to support this dynamic. Unauthorised access to the O2 interface has the potential to compromise the RAN, i.e. to give an attacker complete control over the entire RAN.

A detailed risk assessment is currently difficult since requirements are only described or specified in terms of the services to be made available with the help of the O2 interface. The exact details of how the interface is designed, however – that is, which concrete protocols are

used, and so on – are not specified ("*The O2 services and their associated interfaces shall be specified in the upcoming O2 specification*". [27])
The authors of the O-RAN specifications are well aware that the O2 interface must be secured, as they have defined the general requirement `REQ-O2-GEN-TLS-FUN-1` and `REQ-SEC-O2-1`. However, these are formulated in fairly "soft" terms: "*Management Service providers and consumers that use TLS shall support TLS v1.2 or higher*" [26]. The use of security protocols (specifically TLS) is therefore not prescribed in principle; only a minimum TLS version is specified if TLS is used. The associated explanatory description, "*Communications between SMO and O-Cloud are secure*", [26] can thus be described as misleading, as it represents more a wish than a fact or requirement that can be derived from the specification.
In addition to simply securing the connection (e.g. with the help of TLS), it is also important to define exactly which components or roles are allowed to use which management and deployment services. Based on this, corresponding rights management and access control must be implemented. It is important that the O2 interface has to be designed to support implementation of the least privilege principle in a natural way, i.e. that the interface offers well-defined functions with well-defined parameters that enable sufficiently granular rights management. The opposite of this would be more of a general interface – remote access via SSH, for example. Here, a clear determination of the capabilities of a principally authorised person and the implementation of corresponding restrictions on said capabilities are very complex and error-prone.
As mentioned, the exact design of the O2 interface is unclear, but the O-RAN software community does provide a reference implementation. This is based on a number of existing software components, such as OpenStack, Kubernetes and ONAP. These basic components are highly complex on their own, which makes misconfigurations likely [52] and in turn represents a vulnerability to potential attacks (e.g. privilege escalations) in the case of a less restrictive O2 interface.
With regard to the defined perspectives, attackers and protection goals, the overall assessment is as follows:
O-RAN worst case/3GPP worst case (*ww*): Since no security mechanisms are mandatory, even an outsider can access the powerful O2 interface and take complete control of the RAN. This may result in a failure to fulfil the protection goal of availability for all the stakeholders considered. With regard to the end user, a failure to meet the protection goals of confidentiality, integrity and accountability of transmitted data must also be assumed, since the attacker can gain control over the relevant encryption keys. With regard to the 5G network operator/telco, failures to fulfil the protection goals of confidentiality, accountability and privacy are to be expected. This affects all data relevant to the O-RAN, such as configuration data (including cryptographic keys), machine learning models, O-RAN software components (including rApps and xApps), log files, etc. Since a corresponding attack can easily be carried out on a large scale, a high risk with regard to all protection goals can also be assumed for governmental ("state") stakeholders. However, integrity is an exception from the network operator's point of view: even if a RAN is completely taken over, the Control Plane data between UE and AMF should be secured by the mandatory integrity assurance measures using 3GPP mechanisms.
O-RAN worst case/3GPP best case (*wb*): Here, the assessment of the risks is analogous to that of the worst-case/worst-case scenario. However, the Control Plane data is encrypted by the 3GPP specifications between UE and AMF in the best case, which means that the risk to the confidentiality protection objective can be classified as "medium" with regard to the "telco" stakeholder because the other O-RAN-specific Control Plane data is unprotected.
O-RAN best case/3GPP best case (*bb*): Since all the optional security safeguards are implemented in this case, a low probability of successful attacks by outsiders and users seeking to exploit the O2 interface can be assumed. With regard to "insider" attackers, on the other hand, a high probability can be assumed. Since there are no concrete statements on

rights management or granular access control to the O2 interface, it can currently be assumed that insider attackers can expand their rights with relative ease. This is especially true if the component compromised by the insider is the SMO framework.
By securing the O-RAN area, the risk of a failure to fulfil the availability protection goal should decrease, since incidental access to the O2 interface from the outside and by the user should no longer be possible.
O-RAN best case/3GPP worst case (*bw*): The risks of failures to meet the protection goals should be assessed here in the same way as the best-case/best-case scenario. The only difference is that no Control Plane data is encrypted between UE and AMF, which raises the risk from “medium” to “high”.

Attacker	Perspective (stakeholder)																	
	End user					State					Network operator							
	Protection goals					Protection goals					Protection goals							
	C	I	A	Z	P	C	I	A	Z	P	C	I	A	Z	P			
Outsider																		
User																		
Insider																		

Table 4: Risk assessment of the O2 interface

5.4 O1 interface risk analysis

As described in chapter 2.3.1, the O1 interface is used to manage all O-RAN components (except the management component itself). Access to this interface thus affords an attacker far-reaching possibilities that are only slightly less powerful than unauthorised access using the O2 interface.
However, several specifications exist for this interface: [17], which deals with the O1 interface in general, and [53] which examines in detail the connection between the O-DU and SMO with the help of the O1 interface. The latter is particularly relevant in the context of this study because it goes beyond the general O1 interface description in terms of security requirements and safeguards. A corresponding distinction is made in this case below.

5.4.1 Risk analysis of the general O1 interface

Optional security safeguards are provided for the O1 interface in the current specification documents. Specifically, SSH and TLS are cited here [25]. The implementation of the least privilege principle is also mentioned.
It should be noted that the Security Requirements document [51] makes more strictly formulated statements compared to the actual O1 interface specification: “*O1 interface will enforce confidentiality, integrity, authenticity through an encrypted transport, and least privilege access control using the network configuration access control model*”. However, some of the concrete requirements only use the phrase “*shall support*”, whereas the same document's notes on the security controls to be used are stricter: “*As defined in the previous chapter, the O1 will use TLS 1.2 or higher to enforce confidentiality, integrity, and authenticity; and will use NACM [10] to enforce least privileged access*”.
The O1 interface description generally refers to the Security Requirements document – but only with regard to least-privilege access control and not to transport layer protection. Overall, however, positive tendencies towards securing the O1 interface can be observed here.
Due to the partly unclear wording, it is still assumed with regard to the general O1 interface that the security safeguards are not implemented in the worst-cast scenario. In particular, this decision has been made because the considerations made in chapter 5.4.2 on the more

specific specification of the O1 interface to the O-DU will clearly show that the security safeguards must also be applied here in the worst case. Depending on one's point of view, this assessment can also be adopted as a general assessment of the O1 interface as a whole.
With regard to TLS and SSH as security options, the use of SSH must be considered more risky. This is due to the potential power of SSH, which was already mentioned in chapter 5.1. Specifically, the O1 interface uses the NETCONF protocol, which is why NETCONF-over-SSH [54] is used in the case of a combination with SSH. The SSH connection protocol [23] is explicitly used with a "session" channel type. In principle, this channel type allows the execution of any programme. From a security point of view, it is therefore necessary to implement or configure the associated SSH service in such a way that only the NETCONF subsystem can actually be started. A misconfiguration vulnerability arises here, especially when corresponding restrictions have to be implemented by means of configuration.
The explicit obligation to also support insecure cryptographic algorithms results in a further risk with regard to SSH: "*O-RAN and 3GPP interfaces that implement authentication, confidentiality and integrity using SSH shall: … Enable an O-RAN deployer to configure SSH to offer less secure ciphers using standard SSH configurations to enable backward compatibility with older SSH implementations*" [55].
Based on the SSH-based risks outlined, the recommendation is to support only NETCONF-over-TLS [56] for the O1 interface. In this respect, it is assumed with regard to the best-case analysis that the O1 interface is secured with the help of TLS 1.3 using cryptographic algorithms that are considered secure, as well as mutual authentication.
The NETCONF Access Control Model (NACM) should also be used as an option to enforce least-privilege access control. The following groups are defined in [51]:

- O1_nacm_management – allows access rights to be changed
- O1_user_management – allows the creation and deletion of users for the O1 nodes
- O1_network_management – allows reading, writing and executing on the NETCONF-<running> database, i.e. the NETCONF database that stores all configuration parameters currently in use. The same applies to the NETCON-<candidate> database (if present), i.e. the NETCONF database that contains parameters that have been configured but not yet activated.
- O1_network_monitoring – allows the reading of configurations
- O1_software_management – allows installation of new software

However, the optional nature should also be explicitly noted in the specifications: "*Management Service providers and consumers that use NETCONF SHALL support the Network Configuration Access Control Model (NACM) [...]*" [56].
O-RAN worst case/3GPP worst case (*ww*): Since no security mechanisms are mandatory and the O1 interface – at least with regard to the (O-)RAN core functionality – is similarly powerful to the O2 interface, the considerations from chapter 5.1 apply in the same way. However, the risk of an integrity violation is classified as medium from the network operator's point of view. Due to the mandatory integrity assurance of the Control Plane data (only NAS signals) on the part of 3GPP, it should not be possible for even an insider to change the Control Plane data exchanged between UE and AMF without being noticed. However, an insider still has access to the Control Plane data in the O-RAN, which is why the integrity violation is classified as medium.
O-RAN worst case/3GPP best case (*wb*): Here, the consideration is similar to the worst-case/worst-case scenario: since integrity assurance of the Control Plane data exchanged between UE and AMF is mandatory, the 3GPP best case does not offer improvement compared to the 3GPP worst-case scenario. Although the UE's User Plane data is now encrypted and its integrity secured, an attacker can access the corresponding keys via the O1 interface, so the risk remains unchanged. An important difference, however, is that the Control

Plane data between the UE and AMF is encrypted from end to end in this case. From the network operator's point of view, the risk here is therefore assessed as medium because regardless of all the other considerations, the O-RAN-specific Control Plane data is unencrypted.

O-RAN best case/3GPP best case (*bb*): As in the case of the O2 interface, the circumstances regarding the O1 interface look better in the best case due to the (optional) security safeguards provided. In relation to “user” and “outsider” attackers, the risk of a breach of the protection goals of confidentiality, integrity and accountability is assessed as low for all stakeholders. In the case of availability, the risk for all stakeholders is rated as medium. An attacker could target the availability of the O1 interface to prevent configuration changes and status messages. This in turn can lead to impairments of the quality of service of the RAN – that is, to successful availability attacks on the RAN itself. It should be noted here that, in contrast to the specification of the E2 interface, a failure of the O1 interface is currently not explicitly considered in the specification documents.

Since the O1 interface allows software management, an insider with access to this interface can in principle manipulate the executed O-RAN components and thus, as already detailed, gain access to stored cryptographic keys (for example) in the case of the O-CU. The decisive factor here is how strictly the least privileged principle is implemented and what rights a given insider actually has as a result. The risk an insider poses to the protection goals of confidentiality, integrity and accountability with regard to the “user” and “state” stakeholders is therefore assessed as medium. A similar assessment is made for the “telco” stakeholder. Although parts of Control Plane communication are secured from end to end (and thus protected against insider access), the O1 interface provides access to a great deal of network management information that is sensitive from the network operator's point of view (in terms of both confidentiality and integrity/accountability).

O-RAN best case/3GPP worst case (*bw*): Although the safeguards provided by the 3GPP protection measures are not effective here, the same risks arise as in the best-case/best-case scenario, since “outsider” and “user” attackers have no access to the O1 interface in the O-RAN best case. With regard to “insider” attackers, on the other hand, the fact that the 3GPP security safeguards no longer apply has no negative impact because they were also not effective in the 3GPP best case.

Attacker	Perspective (stakeholder)																			
	End user										State									
	Protection goals										Protection goals									
	C		I		A		Z		P		C		I		A		Z		P	
Outsider																				
User																				
Insider																				

Attacker	Perspective (stakeholder)									
	Network operator									
	Protection goals									
	C		I		A		Z		P	
Outsider										
User										
Insider										

Table 5: Risk assessment of the general O1 interface

5.4.2 Risk analysis of the O1 interface between the O-DU and SMO

In [53], the O1 interface between the O-DU and SMO is used for “start up” installation and the management of software, configurations, performance, faults and files.

For the security of the connection, the use of TLS is to be ***mandatory*** for the authentication of the O-DU, although the mandatory nature is usually rarely found in the specifications in this way. In addition to this, securing the NETCONF data exchange by means of TLS is supposed to be ***mandatory***: “*In this version of O1 Interface Specification, the security of the NETCONF protocol is realized using TLS*” [53].

Use of the NETCONF Access Control Model (NACM) is also supposed to be ***mandatory*** to enforce least-privilege access control. In general, the specification mainly covers the mandatory implementation of security mechanisms in contrast to the security mechanisms in 5.3.1.

Although the enforcement of the protection goals of confidentiality, integrity and accountability is not mentioned with the usual terminology using “shall”, the formulation: “[...] *the O1 interface will enforce confidentiality, integrity, authenticity [...] and least privilege access control [...]*”. [56] leads to the conclusion that the same meaning is intended.

O-RAN worst case/3GPP worst case (*ww*): Since the interface is secured even in the worst case by the mandatory nature of the security safeguards, there is hardly any possibility (at least for “outsiders”) to gain read, and especially write, access to this interface. However, the basic availability attacks that prevent certain status and configuration messages cannot be prevented.

As in the general assessment of the O1 interface, it can also be said that “user” attackers pose a low risk of violating the protection goals of confidentiality, integrity and accountability.

An insider, on the other hand, can breach the protection goals of confidentiality, integrity and accountability in the worst case, since the O-RAN protection measures cannot compensate for the 3GPP safeguards that do not apply here. Here, the O-DU is assumed to be an “insider”. In addition, an insider may be able to install their own software on the O-DU in the worst case (depending on their membership in rights groups) and thereby gain control of it. This results in a high risk to availability.

In spite of all this, however, an insider cannot break the integrity of the Control Plane data between UE and AMF, as it is secured by the mandatory 3GPP measures.

O-RAN worst case/3GPP best case (*wb*): Since all the measures to secure the protection goals are mandatory, the risks should be the same as in the “ww” case. From the perspective of 3GPP security, the confidentiality of the Control Plane data between UE and AMF is assured along with its integrity, which means that the risk posed by insiders can be classified as “medium”.

O-RAN best case/3GPP best case (*bb*): With regard to the best-case scenario, the assessment of “outsider” and “user” attackers is consistent with that of the worst-case scenario due to the mandatory nature of the O-RAN security safeguards. The 3GPP security safeguards can be said to prevent a violation of the protection goals of confidentiality, integrity and accountability only in cases involving “insider” attackers. With regard to availability, it can be argued that even in the best-case scenario, an insider cannot inject malware through NACM, as they do not have the necessary rights. Depending on one's point of view, however, they themselves can be O-DU, which means they can easily restrict availability.

O-RAN best case/3GPP worst case (*bw*): Here, the assessment of the risks is the same as in the previous “bb” case. In this scenario, only confidentiality is not secured by 3GPP in the Control Plane between UE and AMF, which is why the risk is assessed as “medium”.

Attacker	Perspective (stakeholder)																			
	End user					State					Network operator									
	Protection goals					Protection goals					Protection goals									
	C		I		A		Z		P		C		I		A		Z		P		C		I		A		Z		P	
Outsider																														
User																														
Insider																														

Table 6: Risk assessment of the O1 interface between O-DU and SMO

5.5 A1 interface risk analysis

As described in chapter 2.3.3, the A1 interface is primarily used to communicate guidelines to the near-real-time RIC regarding the configuration/optimisation of the RAN. These guidelines are designed as declarative policies. They express *what* is to be achieved, but not *how* it should be achieved. The concrete implementation decisions are thus part of the near-real-time RIC and the xApps.

The A1 interface is meant to provide options for enforcing confidentiality, integrity and authenticity (including protection against replay attacks) [51]. For this purpose, TLS is provided as an option for securing communications [57], [51], [49]. One threat to the A1 interface arises from the fact that, in order to implement bidirectional communication, it is intended that both end points (non-real time RIC, near-real time RIC) of the A1 interface act as servers [57]. The attack surface resulting from the open ports necessary for this should be minimised by having only one endpoint act as a server and carrying out bidirectional communication via an existing connection. To implement this, additional techniques such as reverse connections can be used.

With regard to the best-case/worst-case analysis, it should be noted that the corresponding risks here are essentially dependent on the best-case/worst-case situation pertaining to the O-RAN specifications. In contrast, due to the functionality of the A1 interface, the resulting risk is largely independent of the best-case/worst-case situation regarding the 3GPP standards.

Worst case: Since the A1 interface is unsecured in the worst case, it is possible for outsiders to gain read and write access to it. By manipulating policies accordingly, an impairment of availability (reduction of service quality) is possible at the very least. A high risk to availability is thus assumed for all stakeholders.

With regard to the protection goals of confidentiality, integrity and accountability, a low risk is assumed for the user perspective even in the worst case, as no current attacks could be identified that would allow an attacker with access to the A1 interface to read or manipulate the User Plane data.

It should be noted here that more in-depth analyses are needed to actually rule out the possibility that no User Plane data is transmitted via the A1 interface. Particularly in future versions of the specifications, it will be essential to check which services are actually offered via the A1 interface and whether User Plane data is transmitted there.

With regard to the “telco” perspective, this risk assessment changes to “medium”, since sensitive network management information (with regard to confidentiality and integrity) is transmitted via the A1 interface. This assessment was also adopted for the “state” perspective.

Best case: Due to the protection with TLS in the best case, neither an outsider nor a user can successfully attack the confidentiality or integrity of the A1 interface. Only availability attacks on the A1 interface are conceivable. Due to the functionality of the A1 interface, it can be assumed that intelligent availability attacks on this interface can impair the availability of the RAN – at least in the sense of a reduction in quality of service. This applies in particular because a failure of the A1 interface is not explicitly considered in the specifications. A medium risk is therefore assumed for all stakeholders here.

Since the security safeguards only provide protection against outsiders (connection encryption) and a rights and role concept regarding access control to the A1 interface can be inferred between the lines of the specification at best, it is currently assumed that an insider has full access to the A1 interface even in the best-case scenario. The situation pertaining to insider attackers is therefore similar to that of the worst case involving an outsider. On the other hand, no current attacks could be identified that would enable an insider to gain a significant advantage with regard to the protection goals of confidentiality, integrity and accountability based on access to the A1 interface. The comments made above regarding the need for a more in-depth analysis should be noted here.

Attacker	Perspective (stakeholder)																			
	End user						State						Network operator							
	Protection goals						Protection goals						Protection goals							
	C		I		A		Z		P		C		I		A		Z		P	
Outsider																				
User																				
Insider																				

Table 7: Risk assessment of the A1 interface

5.6 R1 interface risk analysis

As explained in chapter 2.3.4, the R1 interface is used by the rApps to access non-real-time RIC functions and other services that an rApp needs to fulfil its tasks. Here, the following is envisaged: "*The R1 interface is the sole interface between an rApp and the functionality of the Non-RT RIC and SMO. Therefore, the R1 interface should be defined to meet all functional needs of rApps, with appropriate interface extensibility capabilities as needed*." [29]. The R1 interface is thus to be regarded as very powerful. At the same time, it is currently largely unspecified. There are some general statements and some requirements, but no concrete design. The following information is relevant from a security point of view:

- "*It would be useful for the O-RAN Alliance to define an open and standard interface through which the Non-RT RIC exposes SMO Framework functionalities to 'rApps' via the R1 Services exposure functionality. We will refer to this as the 'R1' interface"* [29]
- "*Capabilities are offered for consumption and usage through the R1 services. Such services include, but are not limited to: A1-related services, O1-related services, O2-related services, …*" [29]

The way in which the interface is designed is also unclear. A certain assumption can be made that the API calls are internal and do not take place via a network interface between different computers. Among other things, this assumption arises from the fact that the R1 interface is not mentioned in various O-RAN documents regarding the new, open interfaces. Conversely, based on the general definition of the R1 interface – "*R1 Interface: Interface between rApps and Non-RT RIC framework via which R1 Services can be produced and consumed*" [58] – it cannot be ruled out that the R1 interface is implemented as a network-based interface. Within the framework of this study, it is therefore assumed with regard to the best-case scenario that the R1 interface only deals with internal APIs; in the worst-case consideration, it is assumed that the R1 interface is accessible via the network.

Security safeguards are not currently specified for the R1 interface. It is only mentioned with regard to rApps that access to R1 services should only be possible after authentication and authorisation. Details on how this is to be implemented are also not specified at present.

Due to the principle connection of O1, O2 and A1 services and the lack of current restrictions regarding the use of these interfaces, the risk analysis for the R1 interface is based on the results for the O1, O2 and A1 interfaces. The highest risk for one of these interfaces is regarded as the lower bound of the respective risk for the R1 interface. From a risk perspective, the R1 interface cannot be better than the O1, O2 or A1 interface (in the worst case).

It should also be noted that for the best-case/worst-case considerations, the situation regarding the O-RAN specifications is decisive. In the O-RAN worst case, the O1/O2 interface can be

used to take control of the O-RAN components. Therefore, any 3GPP security safeguards implemented in this case can barely take effect, or they can be circumvented. Conversely, in the O-RAN best case, misuse of the R1 interface is prevented, which is why the implementation of 3GPP protective measures is less crucial. One exception here is the telco perspective with regard to the protection goal of confidentiality: in the best case, the Control Plane is afforded end-to-end protection, which also offers protection in the O-RAN worst case.

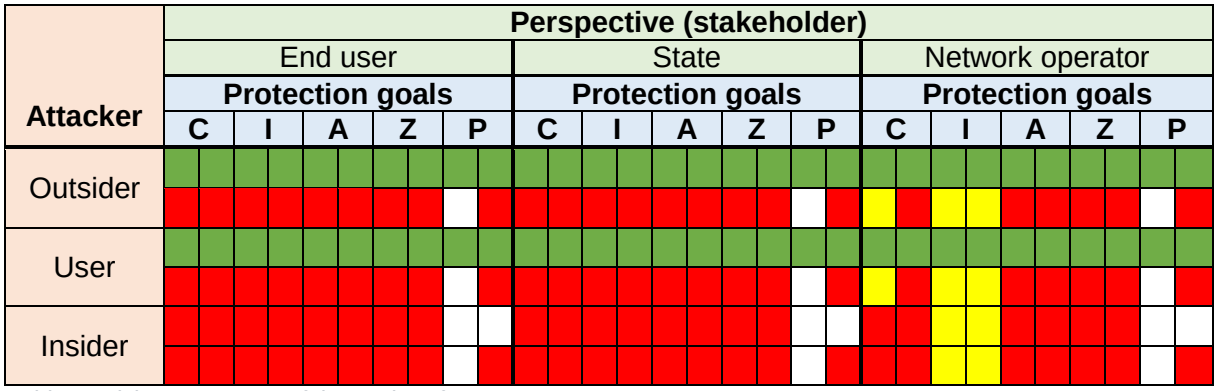

Table 8: Risk assessment of the R1 interface

5.7 E2 interface risk analysis

As detailed in chapter 2.3.5, the E2 interface is primarily used to manage the E2 nodes. In the course of this, a service ("RAN Function Network Interface (NI)" [59]) makes it possible to observe and change all data traffic on the network interfaces of the E2 nodes with the help of the E2 interface. Among other things, the following specific services are offered:

- "*Copy of Complete message with header providing network interface type, identifier and direction with optional network interface timestamp*"
- "*Injection of Complete message with header providing target network interface type, identifier and direction and optional RIC Control Message Priority*"

These services present a high risk potential from a security perspective.

Meanwhile, the security safeguards provided for in the specification are currently considered optional. One specific requirement states that the E2 interface "*shall support confidentiality, integrity and replay protection*" [30]. A similar statement is also found in [51]. IPSec is proposed as a means of protection in [49]: "*IPSEC: Should be used to protect E2 traffic*" (similarly in [51]: "*For the security protection at the IP layer on E2 interface, IPsec shall be supported*").

Based on the functional capabilities of the E2 interface and the corresponding security safeguards, the associated risks are assessed below.

O-RAN worst case/3GPP worst case (*ww*): Since no security safeguards are implemented in the worst case – apart from integrity protection of the Control Plane data – and the E2 interface allows full access to all 3GPP data traffic on the E2 nodes, a high risk to confidentiality, availability and accountability is essentially assumed for all stakeholders. With regard to "telco" stakeholders, a medium risk is assumed due to the 3GPP integrity assurance of the Control Plane data, as O-RAN-specific (configuration) data is unprotected at the same time.

O-RAN worst case/3GPP best case (*wb*): The situation changes here, as the 3GPP security safeguards offer protection related to the protection goals of confidentiality, integrity and accountability. The only question here is the extent to which "insider" attackers can gain access to the keys for securing the User Plane or to the unencrypted User Plane data. In principle, the corresponding keys are available in the CU. Further analyses are required to be able to exclude such risks with certainty, which is why the risks at hand are assessed as "medium".

O-RAN best case/3GPP best case (*bb*): In the best case, the E2 interface and the 3GPP interfaces are secured. An outsider can therefore essentially only carry out denial-of-service

attacks on the E2 interface. A low risk to confidentiality, integrity and accountability is therefore assumed for all stakeholders.
“Stupid” DoS attacks only have a limited effect here, as the O-RAN specifications stipulate that the RAN should be functional even if the E2 interface fails. However, it must be assumed that the RAN is functional with a limited quality of service – otherwise there would be no need for an E2 interface. Furthermore, the effects depend on the concrete design of the E2 node in question. By cleverly manipulating data traffic on the E2 interface, an attacker can achieve a high computing load at an E2 node due to the cryptographic operations associated with IPSec. In extreme cases (if separation is not implemented well), this could lead to insufficient computing power being available for the actual tasks of an E2 node. The extent to which “intelligent” DoS attacks can lead to a more severe impairment of quality of service also remains unclear. Attacks that deliberately delay data packets in such a way that the detection mechanisms regarding a failure of the E2 interface (timer, etc.) are not yet triggered are conceivable in this context. Furthermore, the extent to which it is possible to assign encrypted data traffic to individual E2 services or E2 functions in spite of IPSec encryption (for example, through traffic analyses regarding specific service communication traffic patterns) must also be investigated. This would enable the targeted disruption of individual E2 services or E2 functions and could lead to more severe effects compared to a complete “failure” of the E2 interface. Overall, the risk outsiders pose to availability is therefore considered medium with regard to all stakeholders.
At present, the position of “user” attackers is only slightly better. One new risk/threat stems from a user's ability to generate data traffic pertaining to the User Plane and Control Plane that can be regarded as legitimate in principle. The user can cleverly design this data traffic in such a way that it is included in the analysis and evaluation functions provided by the E2 interface (with regard to the 3GPP interfaces). Depending on how concrete these analysis and evaluation functions are, this influence can be seen as a springboard for further attacks. This is especially the case if the corresponding systems were not developed on the premise that recorded analysis and evaluation data should be regarded as potentially malicious in principle. For example, cleverly crafted User Plane or Control Plane data could be used for buffer overflow or injection attacks (SQL injection, etc.). It is also possible to influence the training of machine learning based on recorded data. Overall, such attacks are likely to primarily impact availability at first. However, since it is assumed in principle within the framework of the best-case analysis that all systems work without errors, the overall risk assessment is not higher than in the case of an “outsider” attacker.
In the context of an insider, the attacker is assumed to have access to the E2 interface. A good example of this is an xApp. There is general talk of “policies” and “authorisations” with regard to xApps here and there in the specifications, but since there is no differentiation at all and no rights management is apparent with regard to the current specification of the E2 interface, it is currently assumed that an xApp can use the E2 interface without restriction, even in the best-case scenario. If this is not the case for an xApp, it is very likely to be for the near-real-time RIC component.
Due to the assumed access to the E2 interface, availability attacks are easily possible for an insider, resulting in a high risk for all stakeholders. In terms of confidentiality, integrity and accountability, the perspective of an end user depends on which component the insider controls. If this component has both access to the E2 interface and knowledge of the keys used to secure the Uu interface, the risk is high. If the insider cannot access the cryptographic keys, the risk is low. As a general consequence, a medium risk is assumed here. This risk assessment is also adopted for the “state” stakeholder. With regard to the “telco” stakeholder, the situation is somewhat better on the one hand, since the Control Plane data traffic is secured from end to end. On the other hand, the E2 interface enables the “*exposure of selected E2 Node data (e.g. configuration information (cell configuration, supported slices, PLMNs etc.), network measurements, context information, etc.) towards the Near-RT RIC*” [30]. This data is

regarded as sensitive trade secrets and, since it forms a basis for network operation, the data's integrity is considered important. A medium risk to the protection goals of confidentiality, integrity and accountability is therefore assumed with regard to the “telco” stakeholder.
O-RAN best case/3GPP word case (*bw*): Since the E2 interface is secured in the O-RAN best case, there are no direct options for “user” and “outsider” attackers to use the E2 interface for attacks. For an insider, accessing an E2 interface with no 3GPP security safeguards makes it possible to successfully breach almost every protection goal from nearly any perspective (again, with the exception of the integrity protection of the Control Plane data, which is also present in the 3GPP worst case). This results in a correspondingly high risk.

Attacker	Perspective (stakeholder)																			
	End user					State					Network operator									
	Protection goals					Protection goals					Protection goals									
	C	I	A	Z	P	C	I	A	Z	P	C	I	A	Z	P					
Outsider																				
User																				
Insider																				

Table 9: Risk assessment of the E2 interface

5.8 Open Fronthaul M-Plane risk analysis

The Open Fronthaul M-Plane interface makes it possible to manage the O-RU components as described in chapter 2.3.7. Particularly relevant from a risk perspective is the possibility of using the Open Fronthaul M-Plane interface to perform a software update with regard to the O-RU components. In principle, this enables an attacker to take full control of the O-RU by importing appropriately manipulated software. In accordance with the O-RAN specification, software protection is not strictly defined: “*The use of compression and ciphering for the content of the software build is left to vendor implementation. The only file which shall never be ciphered is the manifest.xml file*” [32]. However, the manifest in particular contains information that should be protected from manipulation. A special risk arises from the fact that with the help of the Open Fronthaul M-Plane interface, the download URL can be freely specified with regard to the software to be installed. It would be better to provide for restrictions here so that the configured host in the O-RU (for example) is “fixed”, which would limit the negative effects of possible manipulations in the transmission of the download URL (redirection to a server under the control of the attacker).
In addition, it is assumed that access to the Open Fronthaul M-Plane interface facilitates successful attacks on availability, since the configuration data stored in the O-RU can be changed with the help of the Open Fronthaul M-Plane interface. This in turn can be used to impair quality of service or cause a complete functional failure of the O-RU.
In contrast to almost all other interfaces, the considerations for securing the Open Fronthaul M-Plane interface are comparatively extensive and could even be seen as mandatory depending on one's interpretation. Specifically, the use of SSH (for which support is mandatory) or TLS (optional) is possible here. To minimise risk, the use of SSH should be dispensed with in favour of TLS (as already detailed in connection with the O1 and O2 interfaces). With regard to security, mutual authentication is mandatory.
The room for interpretation regarding mandatory *applicable* security mechanisms is mainly due to the unclear wording in the specification documents, such as in the statement: “*The M-Plane provides end to end security as a mandatory feature*.” [32] However, a “feature” is a property or capability of a component and therefore does not have to be applied. In contrast, a

subsequent table states that the protection goals of confidentiality, integrity and accountability are implemented – however, the reason given here only refers to the basic support for SSH or TLS, but not to their mandatory use. As listed later in the table "*Mandatory and Optional Features for O-RU Authentication*", all authentication mechanisms are marked as "*optional to use*". It would be better to clarify here – if this is the intention of the specification – that although each individual mechanism is optional, at least one is mandatory.
Overall, however, the Open Fronthaul M-Plane interface contains many considerations with regard to security. It should be noted here that the description of the Open Fronthaul interfaces is also significantly more extensive overall compared to the other interfaces.
With regard to security, it should also be mentioned that a rights/roles concept is provided. There are six roles mentioned in the specification:

- sudo
- smo
- hybrid-odu
- nms
- fm-pm
- swm

With regard to the rights, a distinction is made between read, write, and execute. The rights for various functional groups (referred to as "namspace" in the specification) are then determined on the basis of the role in question. The role "sudo" has the most privileges. It is intended as an administrative role and thus grants full access to the O-RU, including the option to create users and assign roles to them.
The fact that a default user with a default password and "sudo" rights is provided for out-of-the-box is not considered advisable with regard to the current specification. In the literature, it is frequently reported that default access of this kind is a starting point for successful attacks – no matter how strongly users are advised to change their default access data during initial setup. Here, mechanisms should be provided that enable initial setup without a known default user name and password – for example, by randomly generating this data and attaching it to the O-RU as documentation.
O-RAN worst case/3GPP worst case (*ww*): Assuming that the O-RAN security safeguards are also mandatory in the worst case, "user" and "outsider" attackers pose a low risk to the protection goals of confidentiality, integrity and accountability for all stakeholders. However, due to the unclear formulations and the associated uncertainty as to whether the O-RAN security safeguards really have to be implemented, they are upgraded to a medium risk. Similarly, a medium risk is assumed with regard to availability, since availability attacks on the Open Fronthaul M-Plane interface are possible as previously described for the other interfaces. Since no explicit statements regarding the tolerance of a failure of the Open Fronthaul M-Plane interface could be found, it is also assumed with regard to the Open Fronthaul M-Plane interface that cleverly executed availability attacks can lead to an impairment of RAN service quality.
With regard to "insider" attackers, the situation is somewhat different. An insider with access to the Open Fronthaul M-Plane interface can update the software on the O-RU components, which in turn gives them access to the User Plane and Control Plane data transmitted via the Uu interface. As this data is not protected further in the worst case, it is fully exposed to manipulation – hence the assessment of the corresponding risk as "high" for all stakeholders. Only the integrity of the Control Plane data exchanged between UE and AMF is unaffected by this. Since this data has mandatory integrity protection, the related risk is low. However, it is assumed that there is also (network management) data in the O-RU that is sensitive from the telco point of view (with regard to confidentiality, integrity and accountability) and can be accessed by an insider. Therefore, a medium risk is assumed overall.

O-RAN worst case/3GPP best case (*wb*): Here, the assessment is similar to the worst-case/worst-case scenario. There are some improvements due to the User and Control Plane data, which is secured by the 3GPP measures in this case. As a result, an insider cannot break the confidentiality or integrity of the Control Plane data, which is why the risk here is also assessed as medium overall. It is also assumed that the User Plane data is protected, since insiders are assumed to be the O-RU or O-DU and neither has access to the cryptographic keys used to secure the User Plane.

O-RAN best case/3GPP best case (*bb*): The assessments regarding "outsider" and "user" attackers follow from the already very positive risk assessment in the worst case. In the case of an insider, the attacker does have access to the Uu interface, but the data (as explained previously) is protected by the 3GPP measures in the best case, which is why the risk of violating the protection goals of confidentiality, integrity and accountability is rated as low.

O-RAN best case/3GPP worst case (*bw*): With regard to "user" and "outsider" attackers, the O-RAN security safeguards prevent successful attacks on the confidentiality, integrity, accountability and privacy of User and Control Plane data. The corresponding risk is thus low. With regard to "insider" attackers (who have access to the O-RU), a high risk arises due to the non-existent protection of the User Plane data. For the network operator, the risk to confidentiality is the same, whereby integrity is again secured by the mandatory safeguards on the Control Plane data between UE and AMF. The same does not apply to the management data in the O-RU or O-DU.

Attacker	Perspective (stakeholder)															
	End user					State					Network operator					
	Protection goals					Protection goals					Protection goals					
	C	I	A	Z	P	C	I	A	Z	P	C	I	A	Z	P	
Outsider																
User																
Insider																

Table 10: Risk assessment of the Open Fronthaul M-Plane

5.9 Open Fronthaul CUS-Plane risk analysis

The User and Control Plane data of the Uu interface is transmitted via the Open Fronthaul CUS-Plane interface. In addition, time synchronisation takes place between the O-DU and O-RU. Security is explicitly *not* provided for this interface ("*security requirements: no requirements*" [31]). This is justified on the one hand by the fact that the high requirements for delay time and bandwidth do not permit security. On the other hand, it is assumed that the transmitted data is already secured by security safeguards provided for in the 3GPP standards. Since security safeguards are not provided for, the O-RAN worst-case/best-case considerations do not differ. The 3GPP security safeguards are therefore an essential determinant of risk (see also Appendix A). As a result, only the distinction between the 3GPP best and worst cases is made below.

3GPP worst case (*ww*): Access to the Open Fronthaul CUS-Plane interface enables an attacker to gain full access to the unprotected User and Control Plane data of the Uu interface. "Outsider" attackers thus pose a high risk for all stakeholders. A lower risk can only be assumed for the integrity of the Control Plane data due to the end-to-end protection between UE and AMF. However, this does not affect the integrity of the O-RAN-specific management data or, in particular, that of the very important time synchronisation data (see also chapter 2.3.6.3).

3GPP best case (*bb*): For the “user” perspective, the situation improves in the best case with regard to the protection goals of confidentiality, integrity and accountability, since the Uu interface is secured by 3GPP security safeguards. In the case of accountability, the risk is assessed as medium, since 3GPP does not provide any protective measures in this regard and a user (in contrast to an “outsider” attacker) can import User Plane messages into the system that are recognised as correct during integrity checks.
Availability attacks on the Open Fronthaul CUS Plane interface also impact the usability of the 5G system as a whole, which is why a high risk to availability is assumed even in the best-case scenario.
The 3GPP protections ensure that “insider” attackers also pose a low risk to the User Plane data transmitted on the Uu interface, since the symmetric keys used for protection are located in the O-CU.
The Control Plane data exchanged between UE and AMF is secured by the safeguards provided in the 3GPP best case, resulting in a lower risk for the network operator. However, the time synchronisation data (S-PLANE) is not secured in the 3GPP best case either, so that there is at least a medium risk here overall with regard to the protection goal of integrity.

Attacker	Perspective (stakeholder)																	
	End user					State					Network operator							
	Protection goals					Protection goals					Protection goals							
	C	I	A	Z	P	C	I	A	Z	P	C	I	A	Z	P			
Outsider																		
User																		
Insider																		

Table 11: Risk assessment of the Open Fronthaul CUS-Plane

5.10 CTI interface risk analysis

The CTI interface facilitates capacity reservation in the transport network that is used for the Open Fronthaul CUS plane. This is meant to ensure sufficient resources are always available to meet the high bandwidth and latency requirements of the Open Fronthaul CUS plane, especially in the case of a transport network shared with other services. It is currently assumed that reconfiguration of the network paths themselves (using the CTI interface) is not planned. In this respect, attacks on or access to the CTI interface do not result in any violations of the protection goals of confidentiality, integrity, accountability or privacy. It should be noted that only a digital signature is currently provided as a safeguard for the CTI interface, and its exact specification will only be defined in the future (“*The full details of the CTI Signature will be specified in a future version of this specification*.”) [33] This makes it possible to ensure integrity and accountability. Confidentiality, on the other hand, is not guaranteed. An attacker may thus be able to obtain information about network management and related strategies by eavesdropping on the CTI interface, which must be taken into account in particular with regard to the “telco” stakeholder. It should be noted that the aforementioned digital signature is optional. Furthermore, it is unclear how the mechanism is to function in concrete terms. A bit is provided in the CTI header that signals whether a digital signature is being used. Since this bit itself is not protected, it remains unclear whether the intended digital signature can actually have a meaningful protective effect.
Since the availability of the 5G RAN (especially with regard to the Uu interface) is to be ensured with the help of the CTI interface, availability attacks on the CTI interface itself enable corresponding attacks on 5G RAN availability. In the worst case, “outsider” attackers are

assumed to pose a high risk to all stakeholders. The best-case scenario, meanwhile, may indeed be somewhat better due to the digital signature. Since it is unclear whether the current design of the digital signature can actually have a meaningful protective effect, a medium risk is assumed – especially because attacks that suppress CTI messages are possible despite the digital signature and the current specifications do not explicitly take into account a failure of the CTI interface.

Attacker	Perspective (stakeholder)																			
	End user						State						Network operator							
	Protection goals						Protection goals						Protection goals							
	C		I		A		Z		P		C		I		A		Z		P	
Outsider																				
User																				
Insider																				

Table 12: Risk assessment of the Cooperative Transport Interface

5.11 Risk analysis of other interfaces

In particular, the other interfaces include those used for accessing/downloading external "enrichment information", the AI/ML interfaces and the human-machine interface intended for SMO. The exact design of these interfaces is largely unspecified at the moment.

With regard to access to enrichment information, it is envisaged that this should only take place via secure connections. In any case, this interface and the AI/ML interface represent a security risk in that untrusted providers of the associated information and data can influence management and configuration decisions of rApps/xApps. At the very least, this makes an impairment of availability possible. It is currently not assumed that a direct violation of the protection goals of confidentiality, integrity and accountability is possible with regard to the 3GPP User Plane and Control Plane data.

However, it should be noted here that the data provided by potentially untrustworthy external sources can also be cleverly generated to exploit possible security vulnerabilities (buffer overflow, etc.) in the processing components (especially in the rApps/xApps) and thereby take control of the affected components. The effects of this depend to a very large degree on the component compromised, the rights it has and the extent to which an escalation of privileges is possible. In principle, this can also result in a high risk for the other protection goals.

The human-machine interface is designed to enable a person to influence the SMO components (e.g. in terms of their configuration). The concrete security risks associated with this depend very much on the – currently essentially unspecified – possibilities offered by this user interface. If, for example, it allows software updates to be applied to the various O-RAN components, then it is possible in principle to compromise these O-RAN components, which leads to security risks comparable to those stated in chapter 5.1 in terms of "RAN operator" attackers.

5.12 Risk analysis for rApps

According to the specification, the R1 interface should be the only interface available to rApps. In this respect, the risk analysis for the R1 interface is an essential basis for the risk analysis of rApps. The only point to note here is that the specifications generally speak of access restrictions of individual rApps with regard to the R1 interface; how exactly this is to be implemented, the associated rights and role concept and, in particular, how granular a possible rights management system can be in restricting access to the R1 interface is currently unclear.

Besides the R1 interface, further threats arise from the fact that it can be assumed that a given rApp will be executed in a common hardware/execution environment together with other rApps and SMO functionalities. In the specifications at least, there is no indication that strict (physical) separation is planned with regard to rApps. Insufficient separation and isolation thus represent a possible vulnerability. This is especially true since one planned approach for implementing rApps is to use container-based separation (for example, with the help of Kubernetes). It should be noted that the original objective of separation using containers was not necessarily isolation for security reasons, but rather to avoid dependencies on the runtime environment provided by the operating system. This is why the separation mechanisms implemented and applied in current container solutions are relatively weak. As a result, there is a danger that an rApp could break out of isolation and thereby escalate the privileges granted to it.
From the “user” attacker perspective, it is conceivable to carry out parser attacks with the aim of exerting influence on an rApp through cleverly manipulated user data; in extreme cases this can lead to a compromise of the rApp. However, the data basis (in terms of message formats) for a risk assessment is not available, which is why this is only pointed out here.
Furthermore, the specification does not specify the programming languages to be used for the implementation of an rApp. In principle, this allows the use of rather “insecure” programming languages such as C or C++, where the probability of vulnerabilities that can be exploited by attacks (e.g. through buffer overflows) is significantly higher than with “secure” programming languages such as Rust. Since the communication between rApps is an essential design element of the overall architecture, a malicious rApp could exploit programming errors in other rApps to gain control of a vulnerable rApp and then carry out malicious activities using the privileges granted to that rApp.
In general, it should be noted that the risk analysis presented in the table only refers to the rApps themselves and thus does not take into account attacks on the rApps or the RAN using the rApp interfaces and other interfaces. Corresponding analyses can be found in the chapters on the individual interfaces (especially in the chapter dealing with other interfaces). This explains the low risk attributed to “user” and “outsider” attackers”, since these attackers have to attack an interface in order to compromise an rApp or to introduce a compromised rApp into the system.
In the case of an insider, on the other hand, it is assumed that the attacker has compromised an rApp. Due to the non-existent or very limited O-RAN security mechanisms, there is no need for a deeper distinction between the best and worst cases for O-RAN at this point. Incidentally, the risk assessment is derived from the risk assessment of the R1 interface with regard to “insider” attackers.

Attacker	Perspective (stakeholder)																			
	End user						State						Network operator							
	Protection goals						Protection goals						Protection goals							
	C		I		A		Z		P		C		I		A		Z		P	
Outsider																				
User																				
Insider																				

Table 13: Risk assessment for rApps

5.13 Risk analysis for xApps

For the risk analysis for xApps, the same considerations apply as for rApps. The main difference is that the xApps have access to the A1 and E2 interfaces. In this respect, the risk

analyses regarding these interfaces form the basis of the assessment. In addition, the xApps are “closer” to the 3GPP interfaces (Xn, NG, X2, E1, F1) because the xApps are part of the Central Unit (CU). Should an xApp succeed in breaking through the isolation mechanisms (assuming these are in place), there is a more immediate danger of the 3GPP User Plane or Control Plane data being accessed (in contrast to rApps[14]). Whether a higher security risk can be derived from this in general compared to rApps depends on the future, more concrete design of the respective frameworks and interfaces, since unlike xApps, rApps have the “advantage” of being able to access the O1 and O2 interfaces. In the following, it is conservatively assumed that both types of apps face the same risks with regard to the relevant security objectives.

Attacker	Perspective (stakeholder)																	
	End user					State					Network operator							
	Protection goals					Protection goals					Protection goals							
	C	I	A	Z	P	C	I	A	Z	P	C	I	A	Z	P			
Outsider																		
User																		
Insider																		

Table 14: Risk assessment for xApps

5.14 Machine learning risk analysis

Machine learning is another essential design element of the O-RAN architecture. There are different implementation variants regarding the training of the associated models. Some variants provide for training within the O-RAN components, whereby external data is to be accessed in addition to measured values that originate from the O-RAN components themselves.

In addition to the general uncertainties resulting from the use of machine learning (where decisions are based on correlations, not causalities), from a security perspective there is the particular problem that numerous attacks are known that negatively influence the trained model with the help of manipulated input data. The result is a trained model that either generally delivers incorrect results during inference (application of the model) or – and this can be the greater risk for some use cases – delivers a correct result in almost all cases and only delivers the incorrect result desired by the attacker in situations that can be determined by the attacker.

Based on the potential attacks described and the way in which machine learning is to be used in O-RAN according to the current specification – which is still quite general and unspecific with regard to machine learning – it is assumed that there is a medium risk to availability (i.e. quality of service may be negatively affected). Depending on which parameters are specifically included in the training of a model, legitimate RAN users always present a certain minimum level of risk, adding to the risk already posed by outsiders.

From the operator's (telco's) point of view, another security risk is that the confidentiality of the trained model itself may be compromised. The assumption here is that trained models represent sensitive assets, i.e. they should not become known in the sense of a trade secret. In the literature, there are corresponding “model stealing” attacks that can allow an attacker to derive the parameters of the model or at least to (strongly) restrict their respective range of values by means of clever queries or by manipulating the system either on the basis of the

[14] The risk analysis of the O-RAN Alliance describes that both types of apps are equally capable of violating the protection goals. However, it does not explain how the O-RAN Alliance arrived at this assessment, which is why it is only mentioned here.

answers or generally on the basis of the reactions of the system. Whether such attacks can actually be carried out successfully in the case of O-RAN depends to a large extent on the concrete application scenario of a given model. Due to the current lack of clarity, a medium risk for the telco is assumed here because even outsiders could carry out successful attacks, and at minimum, (supposedly) legitimate RAN users certainly could as well.
Finally, the possibility of attacks on the confidentiality of User or Control Plane data should be kept in mind if this data is used as input for the training of models. This is another area where attack strategies are known from the literature that make it possible to use a trained model to draw conclusions about the input used to train it. Whether this is possible, the consequences it has in the case of O-RAN and which concrete security risks result from this cannot be assessed meaningfully on the basis of the specifications currently available.
Overall, it should be noted that the attacks listed above must be taken into account in the progressive specification of the use of machine learning in O-RAN in order to minimise the risk of successful attacks whenever possible through appropriate design decisions regarding the O-RAN architecture.

5.15 O-RAN risk analysis summary

The summary assessment of the overall security risks associated with the O-RAN architecture is the result of the risks identified for the individual interfaces and components and their aggregation according to the “per security” approach described in chapter 3.6. With regard to insiders, it is assumed that an attacker has control over the CU, which gives them access to the cryptographic keys used to secure the User Plane in 3GPP.
In the table below, some entries are marked with a “+”. This is to express that the authors are quite sure that the assessments of the risk in question are correct, especially based on actual attack scenarios.

Attacker	Perspective (stakeholder)																			
	End user									State								Network operator									
	Protection goals									Protection goals								Protection goals									
	C		I		A		Z		P	C		I		A		Z		P	C		I		A		Z		P
Outsider		+		+	+	+		+			+		+	+	+		+			+			+	+			
	+	+	+	+	+	+	+	+		+	+	+	+	+	+	+	+			+			+	+			
User		+		+	+	+		+			+		+	+	+		+			+			+	+			
	+	+	+	+	+	+	+	+		+	+	+	+	+	+	+	+			+			+	+			
Insider	+	+	+	+	+	+	+	+		+	+	+	+	+	+	+	+			+			+	+			
	+	+	+	+	+	+	+	+		+	+	+	+	+	+	+	+			+			+	+	+	+	
Cloud operator	+	+	+	+	+	+	+	+		+	+	+	+	+	+	+	+			+			+	+			
	+	+	+	+	+	+	+	+		+	+	+	+	+	+	+	+			+			+	+	+	+	
RAN operator	+	+	+	+	+	+	+	+		+	+	+	+	+	+	+	+			+			+	+			
	+	+	+	+	+	+	+	+		+	+	+	+	+	+	+	+			+			+	+	+	+	

Table 15: Summary assessment of the overall security risks associated with O-RAN

6 Summary and Outlook

The O-RAN specifications are currently not being developed according to the paradigm "security/privacy by design/default". It is therefore not surprising that the result is a system that contains multiple security risks. The creation of the threat and risk analysis and the first attempt made to define security mechanisms that should be applied offer the first signs that the O-RAN Alliance could devote more attention to the topic of security in the future. The extent to which this will actually happen remains to be seen. In any case, experience has shown that adding security safeguards late in the process leads to either very high costs or insecure solutions – or, in no shortage of cases, to both. The developments of the 3GPP standards are a good example of this. The initial neglect of IT security has led to an insecure system. Attempts to correct these errors in subsequent versions of the standard have involved a great deal of effort and often lead to solutions that continue to be insecure, especially due to compatibility requirements that must be observed. In fact, in the real-world operations of many public mobile networks, the safety-critical legacy issues lead to uncertainties in modern 5G networks, as well. It is important to prevent a development like this with O-RAN. The O-RAN specifications should therefore be revised with a much stronger security focus before the first productive applications of O-RAN take place.

With regard to the feasibility of security solutions for risk minimisation, it is currently estimated that there are a large number of known security safeguards that can be implemented without great effort and expense to effectively counter individual threats. However, it is also estimated that in order to reduce some security risks, some – possibly even considerable – effort is required to adapt the specifications and implement the corresponding security safeguards.

The following chapter includes some suggestions for risk minimisation.

6.1 Recommendations

6.1.1 3GPP

As an implementation of a 3GPP RAN, O-RAN benefits directly from security improvements to the 3GPP standards. On the one hand, it bears mentioning that the many optional security mechanisms should be made mandatory. In addition, end-to-end security between UE and the 5G core should be introduced for the User Plane data in a manner similar to the provisions of the Control Plane data. In general, the 3GPP standards should also be developed more strongly according to the paradigm "security/privacy by design/default", in particular by leaving behind the current perimeter security model in favour of the principles of multilateral security[15] (i.e. with minimal trust assumptions regarding all stakeholders and components). This could greatly reduce the security risks currently posed by a compromised RAN or an untrusted RAN operator, or by untrusted RAN components.

In summary, this study offers the following proposals:

- Mandatory implementation of currently optional security safeguards
- End-to-end security of User Plane data between UE and the 5G core
- Stronger consideration of the paradigm "security/privacy by design/default"
- Consideration of the principles of multilateral security

6.1.2 O-RAN

One of the most important measures to reduce O-RAN security risks is a serious implementation of the paradigm "security/privacy by design/default", taking into account the principles of multilateral security. The O-RAN specification development processes should be adapted accordingly – for example, by involving security experts in the development of the

[15] This concept is also known by the misleading marketing term "zero trust".

standards and making security considerations and assessments a mandatory part of each specification, as is the case with IETF RFCs (for instance) [60], [61].
Security safeguards that are currently only optional should be made mandatory. At the same time, support for outdated security protocols or cryptographic algorithms that are considered insecure should be explicitly excluded. In particular, clear and unambiguous formulations regarding the mandatory application of the security safeguards must be included in the standards. Their mere mention is not sufficient here. The specification for the O1-O-DU interface, for example, shows that improvements are possible.
When selecting security protocols, care should be taken to ensure that these protocols themselves include as few new threats and security risks as possible. For example, instead of using SSH, it is better to use TLS for transport security. It should also be taken into account that data (including programmes) needs to be secured not only during transmission ("in transit"), but also "at rest", i.e. during persistent storage. At minimum, the protection goals of confidentiality, integrity, accountability and availability should be met.
In addition, a clear rights and roles concept should be implemented for all interfaces and rApps/xApps. The usual security principles (least privilege, need to know, etc.) should also be implemented here. The prerequisite for this is that the interfaces (O2, R1, etc.) are specified as concretely as possible and the basic security mechanisms enable access control that is as fine granular as possible. This also applies in particular to interfaces that currently only play a marginal role, such as those for enrichment information or human-machine interaction. In this context, interfaces that are currently not secured, such as the Open Fronthaul CUS interface, should also be secured. Furthermore, the interfaces should be designed to offer as little attack surface as possible by design. In particular, availability attacks on the interfaces should be taken into account, since negative effects cannot be prevented here simply by applying transport security. The protocols and the overall system should therefore be designed to limit the harmful effects of availability attacks on the interfaces. In addition to restricting functions to those that are actually necessary, this concerns design decisions that influence the number of open ports, for example. In order to have a solution that is as easy to configure and monitor as possible from a security management perspective, there should be as few service access points as possible.
Regarding xApps/rApps, a concept of strong separation and isolation should be implemented so that if one xApp/rApp is compromised, it does not result in other O-RAN components or other apps being compromised. The communication between the apps should also be secured from end to end, and there should be app specifications or implementation guidelines. These should prescribe the use of secure programming languages, for example.
Special attention should be paid to securing the O-Cloud, i.e. the underlying cloud infrastructure. In this regard, measures should be implemented that enable the highest possible level of security, even in the case of untrustworthy cloud operators. In particular, the assumption that cloud operators are trustworthy should be avoided. This also includes the consideration of potentially compromising components of cloud infrastructure. On the one hand, the O-RAN architecture should thus be (re)designed to minimise negative effects as much as possible. On the other, specifications in the sense of security requirements should be made for the components of cloud infrastructure. The extent to which current standard cloud solutions (such as containers) and their implementations (e.g. Kubernetes) meet the security requirements in principle should also be evaluated. The extent to which Trusted Execution Environments (TEEs) are suitable for enabling secure O-RAN operation even in the case of untrusted clouds should be investigated, and the use of TEEs should be made mandatory if necessary.
The risk and security analysis presented here has shown that the combination of the 3GPP standards and the O-RAN specifications results in a complexity that no longer makes a reliable assessment possible simply by "taking a closer look". Formal verification of the specifications (and ideally of related implementations, as well) should therefore be relied on as an essential

element in the development of secure systems. This should also be done in the case of O-RAN. In order to avoid an excessive amount of resulting effort for formal verification, it makes sense to first determine the amount of critical minimum functionality and use this as a basis for deriving the minimum functionality necessary for implementation, which in turn will determine the minimum trustworthy components required. These should then be subjected to formal verification. In doing so, the system as a whole should be designed to support the inclusion of untrusted components in the minimum system if required without jeopardising the secure operation of its minimum functionality. This procedure can thus be compared, for example, to the concepts of a microkernel-based execution environment (operating system, etc.). This is another case where potentially untrustworthy software components can be integrated without affecting one's ability to make determinations or guarantees with regard to security. Overall, a more formal approach can identify security risks that are typically not discovered by a “sharp eye” alone due to the complexity of the specifications.

In summary, the following recommendations can be made:

- Implement security/privacy by design/default
- Actually implement “zero trust” / multilateral security
- Require security that is currently optional at the transport layer
- Use clear wording and avoid ambiguity in clarifying the mandatory use of security mechanisms
- Replace SSH2 with TLS
- Prohibit outdated protocols and insecure cryptographic algorithms
- Secure files “at rest” (encryption, integrity protection)
- Define a clear rights/role concept regarding interfaces and services
    - Especially for the R1 and E2 interfaces (rApps, xApps)
- Clearly specify the O2 interface (cloud management)
- Clearly specify the R1 interface
- Secure the Open Fronthaul CUS interface
- Limit the impact of DoS on interfaces
- Implement firewall-friendly design
    - Minimise access points and server endpoints
- Specify security of connections to external data sources
- Clearly specify a separation concept for xApps/rApps, at least with regard to related requirements
- Secure communication between rApps
- Ensure xApps/rApps use secure programming languages (Rust, etc.)
- Provide for security safeguards (TEEs, etc.) to protect against untrustworthy cloud operators
- Prescribe security mechanisms with regard to O-Cloud, making user authentication in particular mandatory
- Prepare (and, ideally, implement) formal verifiability

7 Bibliography

[1] X. Lin and N. Lee, '5G and Beyond Fundamentals and Standards', *Springer eBook Collection*, 2021, Accessed: Sep. 15, 2021. [Online]. Available: http://wwwdb.dbod.de/login?url=https://doi.org/10.1007/978-3-030-58197-8

[2] 3GPP, 'Study on new radio access technology: Radio access architecture and interfaces', 3GPP, V14.0.0, Technical Report TR 38.801, Apr. 2017.

[3] U. Schulze, 'Endlich offen: Kurz erklärt: Open RAN', *iX*, vol. 2020, no. 9, Heise, p. 120, Aug. 26, 2020.

[4] G. Brown, 'The Role of the RAN Intelligent Controller in Open RAN Systems'. Heavy Reading White Paper produced for Sterlite Technologies Limited, Oktober 2020.

[5] 'The O-RAN Alliance and the Telecom Infra Project (TIP) Reach New Level of Collaboration for Open Radio Access Networks', Feb. 25, 2020. https://www.businesswire.com/news/home/20200225005180/en/The-O-RAN-Alliance-and-the-Telecom-Infra-Project-TIP-Reach-New-Level-of-Collaboration-for-Open-Radio-Access-Networks (accessed Sep. 17, 2021).

[6] 3GPP, 'NG-RAN; E1 Application Protocol (E1AP)', 3GPP, V16.6.0, Technical Specification TS 38.463, Jul. 2021.

[7] 3GPP, 'NG-RAN; F1 Application Protocol (F1AP)', 3GPP, V16.6.0, Technical Specification TS 38.473, Jul. 2021.

[8] 3GPP, 'Architecture enhancements for control and user plane separation of EPC nodes', 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 23.214, Dec. 2018.

[9] 3GPP, 'System architecture for the 5G System (5GS)', 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 23.501, Dec. 2020.

[10] 3GPP, 'General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)', 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 29.281, Mar. 2021.

[11] R. R. Stewart, 'Stream Control Transmission Protocol', no. 4960. RFC Editor, Sep. 2007. Accessed: Aug. 03, 2021. [Online]. Available: https://rfc-editor.org/rfc/rfc4960.txt

[12] 3GPP, 'NG-RAN; NG Application Protocol (NGAP)', 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 38.413, Jul. 2021.

[13] 3GPP, 'Non-Access-Stratum (NAS) protocol for 5G System (5GS); Stage 3', 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 24.501, Jun. 2021.

[14] O-RAN ALLIANCE e.V., 'O-RAN Architecture Description', O-RAN WG1: Use Cases and Overall Architecture Workgroup, V05.00, Technical Specification O-RAN.WG1.O-RAN-Architecture-Description-v05.00, Jul. 2021.

[15] Mavenir, Inc., 'Security in Open RAN'. White Paper, Jan. 2021.

[16] 3GPP, 'Security architecture and procedures for 5G system', 3GPP, V17.1.0, Technical Specification TS 33.501, Mar. 2021.

[17] 3GPP, 'Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Architecture description', 3GPP, V16.0.0, Technical Specification TS 36.401, Jul. 2020.

[18] O-RAN ALLIANCE e.V., 'O-RAN Operations and Maintenance Architecture', O-RAN WG1: Use Cases and Overall Architecture Workgroup, V04.00, Technical Specification O-RAN.WG1.OAM-Architecture-v04.00, 2021.

[19] 3GPP, 'NR; NR and NG-RAN Overall description; Stage-2', 3GPP, V16.6.0, Technical Specification TS 38.300, Jul. 2021.

[20] Parallel Wireless, 'Everything You Need to Know about Open RAN'. E-Book, 2020. Accessed: Jul. 29, 2021. [Online]. Available: https://www.parallelwireless.com/wp-content/uploads/Parallel-Wireless-e-Book-Everything-You-Need-to-Know-about-Open-RAN.pdf

[21] GSMA, '5G Implementation Guidelines'. E-Book, März 2019. Accessed: Jul. 29, 2021. [Online]. Available: https://www.gsma.com/futurenetworks/wp-content/uploads/2019/03/5G-Implementation-Guidelines_v1_nonconfidential-R2.pdf

[22] 3GPP, 'Network sharing; Architecture and functional description', 3GPP, V16.0.0, Technical Specification TS 23.251, Jul. 2020.

[23] Deutsche Telekom, Orange, Telefónica, TIM and Vodafone, ‘Open RAN Technical Priorities under the Open RAN MoU’. Downloadable Document. Accessed: Jul. 27, 2021. [Online]. Available: https://telecominfraproject.com/openran-mou-group/
[24] O-RAN ALLIANCE e.V., ‘O-RAN Use Cases and Deployment Scenarios; Towards Open and Smart RAN’, White Paper, Feb. 2020.
[25] O-RAN ALLIANCE e.V., ‘O-RAN Operations and Maintenance Interface Specification’, O-RAN WG1: Use Cases and Overall Architecture Workgroup, V04.00, Technical Specification O-RAN.WG1.O1-Interface.0-v04.00, Aug. 2020.
[26] O-RAN ALLIANCE e.V., ‘O-RAN O2 Interface General Aspects and Principles’, O-RAN WG6: Cloudification and Orchestration Workgroup, V01.00.04, Technical Specification O-RAN.WG6.O2-GA & P-v01.01, Jul. 2021.
[27] O-RAN ALLIANCE e.V., ‘Cloud Architecture and Deployment Scenarios for O-RAN Virtualized RAN’, O-RAN WG6: Cloudification and Orchestration Workgroup, V02.02, Technical Report O-RAN.WG6.CAD-v02.02, Jul. 2021.
[28] O-RAN ALLIANCE e.V., ‘O-RAN Working Group 2 (Non-RT RIC and A1 interface WG); A1 interface: General Aspects and Principles’, O-RAN WG2: Non-real-time RAN Intelligent Controller and A1 Interface Workgroup, V02.02, Technical Specification O-RAN.WG2.A1GAP-v02.03.01, Jun. 2021.
[29] O-RAN ALLIANCE e.V., ‘Non-RT RIC: Functional Architecture’, O-RAN WG2: Non-real-time RAN Intelligent Controller and A1 Interface Workgroup, V01.01, Technical Report O-RAN.WG2.Non-RT-RIC-ARCH-TR-v01.01, Mar. 2021.
[30] O-RAN ALLIANCE e.V., ‘Near-Real-time RAN Intelligent Controller Architecture & E2 General Aspects and Principles’, O-RAN WG3: Near-real-time RIC and E2 Interface Workgroup, V02.00, Technical Specification O-RAN.WG3.E2GAP-v02.00, Aug. 2021.
[31] O-RAN ALLIANCE e.V., ‘Control, User and Synchronization Plane Specification’, O-RAN WG4: Open Fronthaul Interfaces Workgroup, V07.00, Technical Specification O-RAN.WG4.CUS.0-v07.00, Jul. 2021.
[32] O-RAN ALLIANCE e.V., ‘Management Plane Specification’, O-RAN WG4: Open Fronthaul Interfaces Workgroup, V07.00, Technical Specification O-RAN.WG4.MP.0-v07.00, Jul. 2021.
[33] O-RAN ALLIANCE e.V., ‘Cooperative Transport Interface Transport Control Plane Specification’, O-RAN WG4: Open Fronthaul Interfaces Workgroup, V02.00, Technical Specification O-RAN.WG4.CTI-TCP.0-v02.00, Mar. 2021.
[34] O-RAN ALLIANCE e.V., ‘O-RAN Working Group 2; AI/ML workflow description and requirements’, O-RAN WG2: Non-real-time RAN Intelligent Controller and A1 Interface Workgroup, V01.02, Technical Report O-RAN.WG2.AIML-v01.03.02, Jun. 2021.
[35] O-RAN ALLIANCE e.V., ‘O-RAN Minimum Viable Plan and Acceleration towards Commercialization’, White Paper, Jun. 2021. Accessed: Jul. 28, 2021. [Online]. Available: https://www.O-RAN.org/s/O-RAN-Minimum-Viable-Plan-and-Acceleration-towards-Commercialization-White-Paper-29-June-2021.pdf
[36] DIN ISO, ‘Risikomanagement – Leitlinien (ISO 31000:2018)’, DIN, Deutsche Norm DIN ISO 31000:2018-10, Oct. 2018.
[37] ISO/IEC, ‘Information technology — Security techniques — Information security risk management’, ISO/IEC, Third Edition, International Standard ISO/IEC 27005:2018, Jul. 2018.
[38] IEC, ‘Risk management – Risk assessment techniques’, IEC, Edition 2.0, IEC 31010:2019, Jun. 2019.
[39] BSI, ‘BSI-Standard200-3 --- Risikoanalyse auf der Basis von IT-Grundschutz’, BSI, Version 1.0, BSI-Standard200-3, Oct. 2017.
[40] ENISA, ‘ENISA Threat Landscape for 5G Networks --- Updated threat assessment for the fifth generation of mobile telecommunications networks (5G)’, ENISA, Dec. 2020. [Online]. Available: https://www.enisa.europa.eu/publications/enisa-threat-landscape-report-for-5g-networks/at_download/fullReport
[41] ENISA, ‘ENISA Threat Landscape for 5G Networks --- Threat assessment for the fifth generation of mobile telecommunications networks (5G)’, ENISA, Nov. 2019. [Online].

Available: https://www.enisa.europa.eu/publications/enisa-threat-landscape-for-5g-networks/at_download/fullReport
[42] NIS Cooperation Group, 'EU coordinated risk assessment of the cybersecurity of 5G networks', Report, Oct. 2019. Accessed: Jul. 27, 2021. [Online]. Available: https://digital-strategy.ec.europa.eu/en/news/eu-wide-coordinated-risk-assessment-5g-networks-security
[43] NIS Cooperation Group, 'Cybersecurity of 5G networks EU Toolbox of risk mitigating measures', CG Publication, Jan. 2020. Accessed: Jul. 27, 2021. [Online]. Available: https://digital-strategy.ec.europa.eu/en/library/cybersecurity-5g-networks-eu-toolbox-risk-mitigating-measures
[44] CISA, 'CISA 5G STRATEGY --- Ensuring the Security and Resilience of 5G Infrastructure In Our Nation', U.S. Department of Homeland Security, Cybersecurity and Infrastructure Security Agency, 2020. Accessed: Jul. 27, 2021. [Online]. Available: https://www.cisa.gov/sites/default/files/publications/cisa_5g_strategy_508.pdf
[45] CISA, NSA, and DNI, 'POTENTIAL THREAT VECTORS TO 5G INFRASTRUCTURE', 2021. Accessed: Jul. 27, 2021. [Online]. Available: https://www.cisa.gov/sites/default/files/publications/potential-threat-vectors-5G-infrastructure_508_v2_0%20%281%29.pdf
[46] President of the United States, 'NATIONAL STRATEGY TO SECURE 5G of the United States of America', Mar. 2020. Accessed: Jul. 27, 2021. [Online]. Available: https://www.hsdl.org/?view&did=835776
[47] NTIA, 'National Strategy to Secure 5G Implementation Plan', Jan. 2021. Accessed: Jul. 27, 2021. [Online]. Available: https://www.ntia.gov/files/ntia/publications/2021-1-12_115445_national_strategy_to_secure_5g_implementation_plan_and_annexes_a_f_final.pdf
[48] NTIA, 'National Strategy to Secure 5G Implementation Plan Appendices'. Accessed: Jul. 27, 2021. [Online]. Available: https://www.ntia.gov/files/ntia/publications/5g_ip_appendices_1-5.pdf
[49] O-RAN ALLIANCE e.V., 'O-RAN Security Threat Modeling and Remediation Analysis', O-RAN SFG: Security Focus Group, V02.00.01, Technical Specification O-RAN.SFG.Threat-Model-v02.00.01, Jul. 2021.
[50] GSMA, 'Mobile Telecommunications Security Landscape', März 2021. Accessed: Sep. 29, 2021. [Online]. Available: https://www.gsma.com/security/wp-content/uploads/2021/03/id_security_landscape_02_21.pdf
[51] O-RAN ALLIANCE e.V., 'O-RAN Security Requirements Specifications', O-RAN SFG: Security Focus Group, V01.00.01, Technical Specification O-RAN.SFG.Security-Requirements-Specifications-v01.00, Jul. 2021.
[52] Red Hat, 'State of Kubernetes Security Report', E-book, Jun. 2021. Accessed: Jul. 27, 2021. [Online]. Available: https://www.redhat.com/en/resources/state-kubernetes-security-report
[53] O-RAN ALLIANCE e.V., 'ORAN O1 Interface specification for O-DU', ORAN Open F1/W1/E1/X2/Xn interface Workgroup, V02.00, Technical Specification O-RAN.WG5.MP.0-v02.00, Aug. 2021.
[54] M. Wasserman, 'Using the NETCONF Protocol over Secure Shell (SSH)', IETF, Proposed Standard RFC 6242, Jun. 2011.
[55] O-RAN ALLIANCE e.V., 'Security Protocols Specifications', O-RAN SFG: Security Focus Group, V02.00.06, Technical Specification O-RAN.SFG.Security-Protocols-Specifications-v02.00, Jul. 2021.
[56] M. Badra, A. Luchuk, and J. Schoenwaelder, 'Using the NETCONF Protocol over Transport Layer Security (TLS) with Mutual X.509 Authentication', IETF, Proposed Standard RFC 7589, Jun. 2015.
[57] O-RAN ALLIANCE e.V., 'A1 interface: Transport Protocol', O-RAN WG2: Non-real-time RAN Intelligent Controller and A1 Interface Workgroup, V01.01, Technical Specification O-RAN.WG2.A1TP-v01.01, Mar. 2021.
[58] O-RAN ALLIANCE e.V., 'Non-RT RIC & A1 Interface: Use Cases and Requirements', O-RAN WG2: Non-real-time RAN Intelligent Controller and A1 Interface Workgroup,

V04.00.03, Technical Specification O-RAN.WG2.Use-Case-Requirements-v04.00.03, Jul. 2021.
[59] O-RAN ALLIANCE e.V., 'RAN Function Network Interface (NI)', O-RAN WG3: Near-real-time RIC and E2 Interface Workgroup, V01.00.00, Technical Specification ORAN-WG3.E2SM-NI-v01.00.00, Jan. 2020.
[60] H. Flanagan and S. Ginoza, 'RFC Style Guide', IAB, Informational RFC 7997, Sep. 2014.
[61] E. Rescorla and B. Korver, 'Guidelines for Writing RFC Text on Security Considerations', Network Working Group, Best Current Practice RFC 3552 / BCP 72, Jul. 2003.
[62] 3GPP, 'Packet Data Convergence Protocol (PDCP) specification', 3GPP, V16.3.0, Technical Specification TS 38.323, Mar. 2021.
[63] 3GPP, 'Non-Access-Stratum (NAS) protocol for 5G System (5GS)', 3GPP, V17.3.1, Technical Specification TS 24.501, Jun. 2021.

8 List of Abbreviations

3GPP	3rd Generation Partnership Project
5G	5th Generation
5G SD-RAN	5G Software-Defined Radio Access Network
5GC	5G core
5GMM	5GS Mobility Management
5GS	5G system
5GSM	5G session management
AI	Artificial intelligence
AMF	Access and Mobility Management Function
API	Application programming interface
CapEx	Capital expenditure
CLA	Contributor License Agreement
CN	Core network
COTS	Commercial (or components) off-the-shelf
CP	Control Plane
CTI	Cooperative Transport Interface
CU	Centralized Unit
CU-CP	Centralized Unit Control Plane
CUPS	Control and User Plane Separation
CUS-Plane	Control, User, Synchronisation Plane
CU-UP	Centralized Unit User Plane
DMS	Deployment management services
DSS	Dynamic Spectrum Sharing
DU	Distributed Unit
E2E	End-to-end
E2SM	E2 Service Model
eCPRI	Enhanced Common Public Radio Interface
EDGE	Enhanced Data Rates for GSM Evolution
EI	Enrichment information
eNB	Evolved Node B
EPC	Evolved Packet Core
E-UTRA	Evolved UMTS Terrestrial Radio Access
FCAPS	Fault, Configuration, Accounting, Performance, Security
FH	Fronthaul
FM	Fault management
FOCOM	Federated O-Cloud Orchestration and Management
GERAN	GSM EDGE Radio Access Network
gNB	Next Generation Node B
GNSS	Global Navigation Satellite System
GPRS	General Packet Radio Service
GSM	Global System for Mobile Communications
GTP-U	GPRS Tunneling Protocol - User
HF	High-frequency
HW	Hardware
ID	Identifier
IEEE	Institute of Electrical and Electronics Engineers

IMS	Infrastructure management services
IP	Internet Protocol
ISO	International Organization for Standardization
JWI	Joint work item
KPI	Key performance indicator
L1/L2/L3	Layers 1/2/3 in the OSI model
LTE	Long-Term Evolution
MAC	Medium Access Control
MDT	Minimisation of Drive Test
ME	Managed Element
MIMO	Multiple Input/Multiple Output
ML	Machine learning
MLB	Mobility load balancing
mMIMO	Massive MIMO
MnS	Management Service
MOCN	Multi-Operator Core Network
MOI	Managed Object Instance
MORAN	Multi-Operator Radio Access Network
M-Plane	Management Plane
MSP	Managed service provider
NAS	Non-Access Stratum
NAS-MM	NAS Mobility Management
Near-RT RIC	Near-real-time RAN Intelligent Controller
NETCONF	Network Configuration Protocol
NF	Network Function
NFO	Network Function Orchestrator
NG-AP	Next Generation Application Protocol
ng-eNB	Next Generation Evolved Node B
NG-RAN	Next Generation RAN
Non-RT RIC	Non-real-time RAN Intelligent Controller
NR	New Radio
NSA	Non-standalone
NSSI	Network Slice Subnet Instance
OAM	Operation, administration and maintenance
O-Cloud	O-RAN Cloud
O-CU	O-RAN Centralized Unit
O-CU-CP	O-RAN Centralized Unit Control Plane
O-CU-UP	O-RAN Centralized Unit User Plane
O-DU	O-RAN Distributed Unit
O-eNB	O-RAN Evolved Node B
ONAP	Open Network Automation Platform
ONF	Open Network Foundation
Open FH	Open Fronthaul
OpEx	Operational expenditure
O-RU	O-RAN Radio Unit
OSA	OpenAirInterface Software Alliance
OSC	O-RAN Software Community
OSS	Operations Support System

PDCP	Packet Data Convergence Protocol
PDCP-C	Packet Data Convergence Protocol - Control
PDCP-U	Packet Data Convergence Protocol - User
PDU	Packet Data Unit
PHY	Physical Layer
PLMN	Public Land Mobile Network
PM	Performance management
PNF	Physical Network Function
PRACH	Packet Random Access Channel
PRB	Physical Resource Block
PTP	Precision Time Protocol
QoE	Quality of experience
QoS	Quality of service
RAN	Radio Access Network
RCEF	RRC Connection Establishment Failure
RIA	RAN Intelligence and Automation
RIC	RAN Intelligent Controller
RLC	Radio Link Control
RLF	Radio Link Failure
RRC	Radio Resource Control
RRM	Radio Resource Management
RRU	Remote Radio Unit
RSAC	Requirements and Software Architecture Committee
RT	Real-time
RU	Radio Unit
SA	Standalone
SBI	Service-based interface
SCTP	Stream Control Transmission Protocol
SDAP	Service Data Adaptation Protocol
SD-RAN	Software-Defined Radio Access Network
SFTP	Secure File Transfer Protocol
SI	System integrator
SLA	Service level agreement
SMF	Session Management Function
SMO	Service Management and Orchestration
SON	Self-organizing network
SSH	Secure Shell
SUCI	Subscription Concealed Identifier
SW	Software
SyncE	Synchronous Ethernet
TC	Transport Control
TDD	Time Division Duplex
TEID	Tunnel endpoint ID
TIM	Telecom Italia Mobile
TIP	Telecom Infra Project
TLS	Transport Layer Security
TM	Transport Management
TN	Transport Node

TOC	Technical Oversight Committee
T-PDU	Transport Packet Data Unit
TU	Transport Unit
UDP	User Datagram Protocol
UE	User equipment
UMTS	Universal Mobile Telecommunications System
UPF	User Plane Function
V2X	Vehicle-to-Everything
vCU	Virtual Centralized Unit
vDU	Virtual Distributed Unit
VM	Virtual machine
VNF	Virtual Network Function
vRAN	Virtual Radio Access Network
WG	Working Group (of the O-RAN Alliance)
YANG	Yet Another Next Generation

Appendix A: 3GPP 5G RAN Risk Analysis

The following summary of a risk analysis for a standardised RAN based on 3GPP is the result of an evaluation of relevant 3GPP standards and the inclusion of existing threat and risk analyses, as well as of potential attacks known from the related scientific literature. An understanding of the security risks inherent in RANs and potential countermeasures is essential to better assess and understand the risks associated with O-RAN as a concrete RAN implementation.

According to [62], [16], the User Plane data of the Uu interface is secured within the PDCP layer. Both symmetric encryption and symmetric integrity assurance (message authentication codes) are used. Due to the symmetric integrity protection, fulfilling the protection goal of accountability is not possible. The symmetric keys used here are known to the gNB. There is therefore no end-to-end security between UE and the 5G core. The transmission of User Plane data between the gNB and the 5G core via the NG-U interface is secured with the help of IPSec. Overall, it should be noted that securing both the Uu and the NG-U interface is optional: "*Confidentiality protection of user data between the UE and the gNB is optional to use. […] Integrity protection of the user data between the UE and the gNB is optional to use, and shall not use NIA0."* [16].

Whether security should be implemented is specified in the 5G core, specifically by the locally responsible SMF. In the best-case scenario, it is assumed that the security mechanisms are activated; in the worst case, it is assumed that the User Plane data is transmitted unsecured.

Similar to the User Plane, security safeguards are also provided in principle for the Control Plane [62], [16], [63] – although the security safeguards regarding confidentiality are only optional for the Control Plane: "*Confidentiality protection NAS-signalling is optional to use*." [16] The integrity of the Control Plane data, on the other hand, is mandatory: "*All NAS signalling messages except those explicitly listed in TS 24.501* [63] *as exceptions shall be integrity-protected with an algorithm different to NIA-0 except for emergency calls."* [16] The cases listed in [63] are as follows:

> *a) for an unauthenticated UE for which establishment of emergency services is allowed;*
> *b) for an W-AGF acting on behalf of an FN-RG; and*
> *c) for a W-AGF acting on behalf of an N5GC device.* [63]

This means that the edge cases in which no integrity assurance takes place are both well defined and truly marginal (for example, emergency calls). A significant difference compared to the transmission of the User Plane data is that the messages of the Control Plane are transmitted with end-to-end security (with regard to UE and the 5G core). This reduces corresponding risks posed to the protection goals of confidentiality and integrity by the 5G RAN.

"Cloud operator" attackers are not taken into account in the risk analyses, as the 3GPP specification does not make any concrete specifications for the implementation of a 5G RAN. A cloud-based solution can thus be used here, as can a cloud-free, monolithic implementation. With regard to the following analyses, it should be noted that overall, they represent the "best case" from a security perspective, as not all possible attack scenarios could be analysed in depth due to time constraints. The unanswered questions here include:

- Can a RAN influence the security mechanisms chosen for Control Plane protection and potentially result in successful attacks by a malicious RAN despite the end-to-end security provided?
- What risk arises from the lawful interception interfaces?

Worst case: Based on the statements made above, the worst case is that no security safeguards (except integrity protection with regard to the Control Plane) have been implemented. As a result, an "outsider" attacker poses a high risk of violating the protection goals of confidentiality, integrity, accountability and availability from the "user" perspective. The

same applies to the perspective of the network operator. Due to the high risk for these two stakeholders, a high risk is also assumed for the “state” stakeholder. Integrity is the only exception here, which is why the risk is assessed as low.

Best case: In the best case, the situation improves through the application of the planned security safeguards. For “outsider” attackers and all stakeholder perspectives, a low risk of a breach of the protection goals of confidentiality, integrity and accountability is assumed. With regard to “user” attackers, the only difference is that the risk of violations of the protection goal of accountability is assessed as medium, since a user can generate “valid” messages that make it impossible for third parties to determine whether they originate from the user or from the RAN. With regard to the protection goal of availability, a medium risk is assumed for both “outsider” and “user” attackers. The literature here outlines potential attacks based on a “user” attacker gaining control of a large number of (IoT) devices connected to a 5G RAN and using them for a distributed availability attack (DDoS) on the 5G RAN. With regard to “outsider” attackers, there are indications in the literature that Uu interface transmissions can be severely impaired by the intelligent broadcasting of radio signals (jamming attacks). Both the User Plane and the Control Plane are affected by availability attacks. It is therefore assumed that both a 5G user and the 5G network operator can be affected, which consequently implies a medium risk for the stakeholder “state”.

With regard to “RAN operator” attackers, a high risk is also assumed in the best-case scenario for the user perspective, since the keys for securing confidentiality and integrity are known to the RAN and these security safeguards therefore do not protect against a malicious RAN operator. The 3GPP standard should be revised so that end-to-end security between UE and the 5G core is also possible for the transmission of User Plane data. With regard to the “network operator” stakeholder, the assessment of the situation is somewhat better, since at least the Control Plane messages are secured from end to end. For all stakeholders, there is a high risk with regard to availability, as the functional and orderly RAN operation is crucial for the availability of a 5G network.

The risk posed by an insider (i.e. a compromised component) is assessed similarly to the risk regarding an untrusted RAN operator. If the attacker has access to the 5G RAN component that manages the keys for securing communications, the insider can also successfully attack the protection goals of confidentiality, integrity and accountability with regard to User Plane data. At the same time, the attacker can incidentally restrict availability by applying or making available false keys for decryption.

For the protection goal of privacy, a low risk was assumed with regard to “outsider” and “user” attackers, since a series of measures are provided by 3GPP to prevent UE tracking. It should be noted here that attacks using wireless fingerprinting of UE were not taken into account. If this type of attack is taken into account, a medium risk can be assumed. The privacy risk in the case of “insider” and “RAN operator” attackers cannot be reliably assessed. The potential attacks are mentioned in the literature, but extensive further analyses are needed to arrive at a well-founded assessment.

Attacker	Perspective (stakeholder)														
	End user					State					Network operator				
	Protection goals					Protection goals					Protection goals				
	C	I	A	Z	P	C	I	A	Z	P	C	I	A	Z	P
Outsider															
User															
Insider															
Cloud operator															
RAN operator															

Table 16: Risk assessment of the 3GPP 5G RAN