

Quantum Internet Research Group
Internet-Draft
Intended status: Informational
Expires: 23 July 2026

D. Lopez
Telefonica
V. Martin
UPM
B. Lopez
IMDEA Networks
L. M. Contreras
Telefonica
19 January 2026

A Multiplane Architecture Proposal for the Quantum Internet
draft-irtf-qirg-qi-multiplane-arch-00

Abstract

A consistent reference architecture model for the Quantum Internet is required to progress in its evolution, providing a framework for the integration of the protocols applicable to it, and enabling the advance of the applications based on it. This model has to satisfy three essential requirements: agility, so it is able to adapt to the evolution of quantum communications base technologies, sustainability, with open availability in technological and economical terms, and pliability, being able to integrate with the operations and management procedures in current networks. This document proposes such an architecture framework, with the goal of providing a conceptual common framework for the integration of technologies intended to build the Quantum Internet infrastructure and its integration with the current Internet. The framework is based on the already extensive experience in the deployment of QKD network infrastructures and on related initiatives focused on the integration of network infrastructures and services.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://dr2lopez.github.io/qi-multiplane-arch/draft-irtf-qirg-qi-multiplane-arch.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-irtf-qirg-qi-multiplane-arch/>.

Discussion of this document takes place on the Quantum Internet Research Group Research Group mailing list (<mailto:qirg@irtf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/qirg/>. Subscribe at <https://www.ietf.org/mailman/listinfo/qirg/>.

Source for this draft and an issue tracker can be found at
<https://github.com/dr2lopez/qi-multiplane-arch>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 23 July 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	3
2. Conventions and Definitions	5
3. Base Technologies: QKD Experience and Evolved SDN Concepts	5
3.1. A QKD Multi-Plane Architecture	5
3.1.1. The Service Unit Concept in QKD Networks	7
3.2. Interfacing with Classical Networks	9
3.3. Applying Network Virtualization Principles	10
3.4. CLAS and Quantum Networks	11
4. A Framework Architecture for the Quantum Internet	12
4.1. Strata for Quantum Networks	12
4.2. Identification of Interfaces and Protocols	16
4.2.1. The Role of Service Units	18
4.2.2. The Role of Synthetic Environments	20
5. Security Considerations	23
6. References	23

6.1. Normative References	23
6.2. Informative References	24
Acknowledgments	27
Authors' Addresses	27

1. Introduction

As another case of the "classical vs quantum" apparent contradictions, the nature of quantum communications [QTTI21], associated with natural physical effects that require a specific infrastructure to be used for communications, poses a significant challenge in the definition of any network reference architecture to be used for such communications. Nevertheless, the growing interest on quantum networking, its applications, and the eventual availability of a Quantum Internet, require of consensus on an architecture framework able to support the definition and evolution of different protocols and interfaces.

Several steps have been taken in this direction, including the identification of architectural principles and base technologies made in {RFC9340}}, the description of relevant use cases [RFC9583], and specific approaches to layered models for Quantum Networking, summarized and discussed in [QIPS22]. While the principles provide an extremely valuable common ground for further collaboration among quantum and network practitioners, they are not intended to provide the solid framework required for progressing in the definition of specific protocols and other interfaces for common network management tasks and interactions with user applications. On the other hand, the proposals made for a layered approach provide interesting insights on requirements and potential mechanisms to structure quantum communications, but, first, they do not include essential aspects for a network at scale and, second and most important, they do not take into account the need for direct interactions beyond the layered structure, such as those between classical and quantum networking services, between applications and the quantum network, etc.

In parallel, the operational experience with the first kind of infrastructures using quantum communication technologies to provide an actual network service, those focused on Quantum Key Distribution (QKD), has allowed practitioners to explore the solution space and identify design patterns that seem applicable to the general case of a Quantum Internet. A corpus of architectural proposals [Y3802], experimental deployments [MADQCI23] and pilot infrastructures [EUROQCI] have become available in the recent years, and can be used to derive useful conclusions, especially if combined with recent proposals in network architecture [RFC8597], intended to address the complexity of management and integration at scale beyond the basic layered constructs supporting connectivity.

This document proposes a multi-plane reference architecture for the Quantum Internet, derived from available proposals and the operational experience with QKD infrastructure. The proposal attempts to define a framework with three essential properties to guarantee a seamless evolution of the technology, and the consolidation of applications and management practices:

- * **Agility:** Provide abstractions able to incorporate new protocols and interfaces as the technology evolves, avoiding a tight coupling with specific physical technologies.
- * **Sustainability:** Considering it at all levels and in full scale, especially regarding environmental and social impacts, including open availability in technological and economical terms, and fostering infrastructure reuse.
- * **Pliability:** Facilitate the seamless integration of classical and quantum network operational procedures, applying and adapting best practices in use by the Internet community.

And trying to address three essential characteristics already identified in [PSQN22]:

- * **Universality,** so a quantum network can accommodate any application.
- * **Transparency,** so quantum networks can share physical media with classical networks.
- * **Scalability,** so quantum networking protocols can support the growth of the network.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Base Technologies: QKD Experience and Evolved SDN Concepts

The design and deployment of QKD infrastructures has followed a number of design principles, based on the best practices in network architecture and management established during the lifetime of the Internet (and even before), and focused on the separation of concerns, that have been converging on the trends around open disaggregation strategies, and the identification of separate data and control planes, connected by means of open interfaces.

In what relates to the evolution of SDN concepts, the Cooperating Layered Architecture for Software-Defined Networking (CLAS) [RFC8597] described a SDN architecture structured in two different strata, namely Service Stratum and Transport Stratum. On one hand, the Service Stratum contains the functions related to the provision of services and the capabilities offered to external applications. On the other hand, the Transport Stratum comprises the functions focused on the transfer of data between the communication endpoints, e.g., between end-user devices, between two service gateways, etc.

3.1. A QKD Multi-Plane Architecture

Applying the SDN and disaggregation principles, QKD infrastructures have been essentially structured around three different planes [QTTI21]. While we are not talking about a rigid, layered structure, where a given layer can only provide services to the immediate upper layer and consume services from the immediate lower layer, it is worth noting that interactions among elements in the different planes must use well-defined interfaces [ETSI04] [ETSI14] [ETSI15] [ETSI18], and these interactions may incorporate a layered approach.

In this approach, the Quantum Forwarding Plane (QFP) is in charge of performing the operations (quantum and classical) to ensure the forwarding of the quantum signals or enable the utilization of persistent quantum resources, like persistent, distributed entanglement. In QKD, the QFP encapsulates all the functionality required to obtain an end-to-end secret key across the network. This implies the transmission of the quantum signals and the execution of any associated protocols. Note this would require the use of classical procedures, either via a separated physical "classical

channel" [QTTI21] or the reuse of a common channel, as proposed in "packet-oriented" approaches [PSQN22]. In this sense, the forwarding of the keys at intermediate nodes in the multi-hop chains used to overcome current limitations in propagation of quantum signals or states, has to be considered part of the QFP, since it is done exclusively on behalf of the QKD functionality.

On its side, the Service Overlay Plane (SOP) supports the use of the keys derived from the QFP by applications. This includes the storage, identification, delivery, and lifecycle management of the units of consumption (keys of different length, delivered according to specific patterns) at the endpoints of the network. All network functionalities at this plane can be considered application-oriented, with a clear mapping to an overlay data plane in a classical network, though the SOP elements should be aware of the nature and specific needs of the QFP they interact with. Key management mechanisms, beyond key forwarding by intermediate nodes, fit within the SOP. This comprises methods such as hybridization and augmentation techniques, or the means for synchronizing key identifiers across API boundaries.

Finally, the Control and Management Plane (CMP) is made of the elements that create and supervise the state of the network. This decoupling between network configuration and (general) data forwarding is supported by the controller, a mediation logically centralized element between the control capabilities supported by the elements in the QFP and SOP and the management and control functions. These management and control applications rely on the controller, taking advantage of the centralization it provides, to guarantee the best performance of the network and avoid diverging local control decisions that might lead to sub-optimal configurations.

It is worth noting this management centralization does not contradict the distributed principles generally applied in current networks. Local control decisions are intended to be coordinated by centralized management. While the communication between the controller and the controlled elements relies on some kind of SDN protocol, the controller exposes a consistent abstract model of the network devices and topology, that can be structured in a hierarchy of abstractions, from lower-level, element-focused ones, up to application-oriented ones.

In summary, QKD infrastructures are converging into an extended SDN model, with two differentiated data planes, controlled in a coordinated manner through a common Control and Management Plane, that supports aggregated mechanisms for further orchestration. The QFP/SOP duality constitutes a common abstract foundation for a general approach to quantum communications networks, regardless of their final purpose.

3.1.1. The Service Unit Concept in QKD Networks

QKD infrastructures are evolving from a conglomerate of links, where keys derived from the protocol applied to a link are used to secure the communication between two entities directly associated to the endpoints of the link, into real networks, able to forward a key to be used between any two entities attached to the network. As discussed above, the entities in the SOP play a key role for this, supporting the storage, delivery and lifecycle management of the service units being consumed by the applications attached to the network. These SOP entities, are commonly referred as KME (Key Management Entity), acting as key storage for a specific element or elements in the QFP, and providing an endpoint for applications to request and consume keys for a specific secure interaction. The interfaces KMEs use to interact with the QFP elements are usually provided by specific (commonly software-based) components, acting as agents in the QFP, and therefore termed Key Management Agent (KMA).

Several of these KMEs can be logically grouped into what is called a KMS (Key Management System), supporting a set of related applications grouped into a trust domain, and therefore consistently operated by a corresponding entity in the CMP. The differentiation between KME and KMS functionalities becomes more apparent as networks expand and consolidate, with many cases of current QKD link-oriented infrastructures referring to a KMS as the entity integrating both roles.

The service units provided by a QKD network have to be uniquely identified within the network, so they can be properly managed by the SOP, including their routing across the different required KMEs, the requests of appropriate links in the QFP, and the management of the lifecycle events related to making the key available to the applications willing to use it. It is important to note we are talking about a service unit, and not a data unit associated with a particular protocol, and therefore what is relevant here are the identification of the two application endpoints (that should include a nonce mechanism to identify the specific pairing) together with relevant parameters regarding the key lifecycle, such as its length and valid time-to-live. While these are the two essential lifecycle parameters, others, as it might be the case of applicable crypto

algorithms, could be considered as well. The service unit identifier is not directional, i.e., it has no source or destination addresses, as it defines a shared state to be used by two applications. We can consider the analogy of transport flows in the current Internet, rather than packets.

The current proposal we are experimenting with advocate for using URNs [RFC8141] as endpoint identifiers, taking advantage of their nature of location-independent, persistent resource identifiers. The q-component of the endpoint URN can be used to carry the nonce part of the specific application identifier. If we consider that lifecycle parameters can be expressed using a specific URN in its q-component, we have that a service unit identifier consists of the combination of three URNs.

As an example, let's consider URNs for application endpoints use the qkd namespace id, and that lifecycle parameters use the URN qkd:lifecycle assigned name, with the parameters size and valid-until. A service unit identifier for QKD between two domains, with roots madqci and quditto, would look like:

```
urn:qkd:madqci:ccips?=nonce=177923
urn:qkd:quditto:emulator:ipsec:controller?=nid=af33017
urn:qkd:lifecycle?=size=256&valid-until=1750708945
```

The structure and delegation mechanisms provided by URNs allow for arbitrary aggregation of prefixes, enabling any kind of routing style, from the aggregation and inter-domain announcement similar (or compatible) to BGP in classical Internet to the decision on which prefixes are announced and how they are routed by means of SDN controllers, whether by means of a federation approach or in a hierarchical control structure. The approach also supports the use of non-routable identifiers that cannot be announced outside a given domain or KMS and can only establish service pairing with other applications within the same domain. These mechanisms would be applied by the corresponding elements in the CMP.

The QKD service unit identifies a shared state between two application entities, and therefore cannot be consider directional, and the concepts of source and destination do not apply here. Nevertheless, directionality is relevant in the process of establishing the QKD service unit, both in terms of its identifier and of its contents. In the case of the identifier, one of the application entities will request a service unit to the relevant KMS/KME it is attached to, identifying itself and the other peer in the service unit, together with the applicable lifecycle parameters. Relying on the available route information and the replies of the intermediate elements in the SOP, the final identifier of the QKD

service unit will be built. The associated content, i.e., the bit string defining the key to be shared between the two application endpoints, will be derived from the elements in the participating links in the QFP and the application of any additional mechanisms (key encryption, augmentation, trusted node forwarding...) required by the participating KMEs and the corresponding links.

3.2. Interfacing with Classical Networks

The interface of QKD infrastructures with classical networks (commonly identified as OTN, Optical Transport Networks) has been based on three basic principles, related to the ones we mentioned above: facilitate the reuse of physical infrastructure (sustainability and transparency), apply the abstractions commonly used in open and disaggregated networks (agility and universality), and reuse the best practices in network management being applied in current infrastructures (pliability and scalability). We can classify the interface mechanisms according to the level at which they occur.

At the application level, end-to-end key management and end-to-end key creation are obviously the main target. Since many applications of these keys are related to classical communications (direct encryption, key derivation for symmetric algorithms, peer identity...) there is a clear interface for the SOP, with classical network functions acting as consumers of the keys or, in general terms, the bit streams generated by the QFP. Further on, the application of NFV mechanisms to any network function allows for its implementation through software virtualization techniques (virtual machines, para-virtualization containers, unikernels, etc.), irrespectively of their application environments or specific plane. The lifecycle management of all network functions, of any nature, under a common MANO stack [NFV06], seems the most reasonable option.

At the control and management level, the distinct nature of network elements and the mediation nature of the controller role do not make advisable the use of common quantum/OTN controllers, but there are common abstractions able to support cross-interactions among controllers and management applications, especially regarding:

- * Quantum management applications requiring operations on topologies and physical paths in the OTN mediated by an OTN controller.
- * OTN management applications requiring operation on quantum topologies mediated by the quantum controller.
- * Topology updates exchanged between quantum and OTN controllers.

- * The coordination through an integrated controller (commonly referred as "orchestrator"), able to provide a common view to application network functions.

At the forwarding level, there is a radical difference between the network elements in quantum networks and OTN, and therefore interactions in data forwarding are not feasible, with only two exceptions: the possibility of sharing physical media, and the use of classical channels to support QKD algorithms, as it is the case of distillation channels in protocols like BB84. In this case, a proper control of the path and physical parameters has to be applied to minimize interferences of any nature and guarantee OTN connectivity for the quantum algorithms.

3.3. Applying Network Virtualization Principles

Recent proposals for QKD network management have explored the use of operational models that radically leverage the virtualization of control and key management functionalities [EVCK25]. These approaches pave the way for a tighter integration of quantum functionality with functions already established in state of the art classical networks, including support for user/function authentication and authorization, and support for user and function mobility, while adhering to established QKD network standards. Integrating these mechanisms enhance security measures and ensure that quantum networks can seamlessly interface with existing and future telecommunications infrastructure.

While SDN ensures higher degrees of flexibility and reconfigurability by allowing network functions to be easily modified and upgraded through software changes, virtualization enables the abstraction of hardware devices by creating virtual instances, which improves scalability, resource efficiency and allows the dynamic allocation of softwarized network functions in different locations. As quantum technology evolves, a virtualized layer for softwarized network functions significantly aids adaptation to these changes, ensuring pliability and responsiveness for seamless updates, and incorporating new mechanisms without extensive hardware modifications.

As for key exchange, current technology does not allow direct end-to-end quantum key exchange between distant nodes. Instead, key distribution must rely on trusted intermediary nodes to transmit keys hop-by-hop. A key management layer where the actions of all nodes are coordinated is needed to ensure secure and efficient key distribution. Virtualizing and decoupling key management from the physical QKD devices enhances flexibility and scalability, and supports the integration of hybrid cryptographic strategies, combining QKD and post-quantum algorithms to ensure security and

performance. Additionally, it allows real-time performance monitoring, data-driven control and management, and tailored access and admission mechanisms [QNSA24].

The virtualized key management layer acts as an intermediary between the clients and the cryptographic material generating devices. This layer would have as functions both those that fall within the framework of the SOP defined in previous sections, as well as key forwarding, specific to the QFP. For the latter, each functional element of this layer, identified as key managers entities in [EVCK25], has a forwarding table, which can be dynamically updated whenever necessary by the control plane. Additionally, they implement a token bucket for each application session, to control the request rate by limiting it to an agreed-upon value at the Quality of Service (QoS) level.

The virtualized control plane can have different functional elements, and, as with the key management layer, several instances of the same element can be executed as necessary for the correct operation of the network. Foundational elements include: a controller, an access control and an admission control component, a routing module, and a monitoring element. This set allows the execution of network access policies, ensuring that no unauthorized user or process enters the network, verifies the configuration parameters of new sessions opened by applications, ensuring that they are granted the appropriate QoS, and performs performance tests on the physical links and collecting statistics on the QKD modules, quickly alerting about any failure or possible attack on the QFP.

3.4. CLAS and Quantum Networks

As discussed above, SDN principles have enabled the base abstractions for the conceptualization of QKD infrastructures, including the services they provide and the required interactions in the use of classical infrastructure to support the required connectivity patterns. The original CLAS architecture, as defined by [RFC8597], addresses SDN evolution considering the forwarding (transport) and service aspects in two separated but coordinated planes. This approach matches the multi-plane approach described for QKD infrastructures, though it seems somehow limited to address the required interactions with physical connectivity, as well as to incorporate general requirements regarding automation to support convergence with operational practices.

The new extension of the CLAS architecture, as defined in [CLASEVO], intends to address the current evolution of networks and the services they support introducing new aspects, in particular the considerations of distributed computing capabilities attached to

different points in the network, and the introduction of evidence-driven techniques, such as Analytics, Artificial Intelligence (AI) and Machine Learning (ML) to improve operations by means of closed-loop automation.

The CLAS framework provides a sound foundation for incorporating the experience gained with QKD deployments in a general proposal applicable to the Quantum Internet, as it is essentially compatible with the architectural lessons learned within the QKD fields, and at the same time supports additional degrees of freedom regarding the integration of control mechanisms, and the interplay with the (shared) infrastructure and its management.

4. A Framework Architecture for the Quantum Internet

Based on the available experience on the deployment of existing QKD infrastructures and on the evolution of SDN-enabled architectures described in the previous section, this document proposes an architecture framework intended to offer a conceptual common framework for the integration of technologies intended to build the Quantum Internet infrastructure and its integration with the current Internet.

Once we presented in the previous section the lessons learned from QKD deployments, introducing a general architecture applicable to those deployments, in this section we propose the generalization of such architecture towards a Quantum Internet, augmented by the extended SDN approach proposed by the evolved CLAS in [CLASEVO]. In what follows, we will discuss how this framework architecture would support the required properties: agility, allowing for technology evolution, sustainability, fostering infrastructure reuse, and pliability, supporting operational best practices.

Furthermore, we propose here a general network architecture trying to incorporate relevant trends such as cloud nativeness, the integration of zero-touch management, or the considerations about intent. With this in mind, in what follows a CLAS-based architecture frameworks for quantum communications networks is introduced, including the proposed strata and their main characteristics.

4.1. Strata for Quantum Networks

The CLAS architecture was initially conceived from the perspective of exploiting the advantages of network programmability in operational networks, complementing and going beyond the traditional layered structured of the original SDN proposal. Following the CLAS philosophy, as proposed in its recent update [CLASEVO] of decoupling services, additional functionality, and base connectivity, the

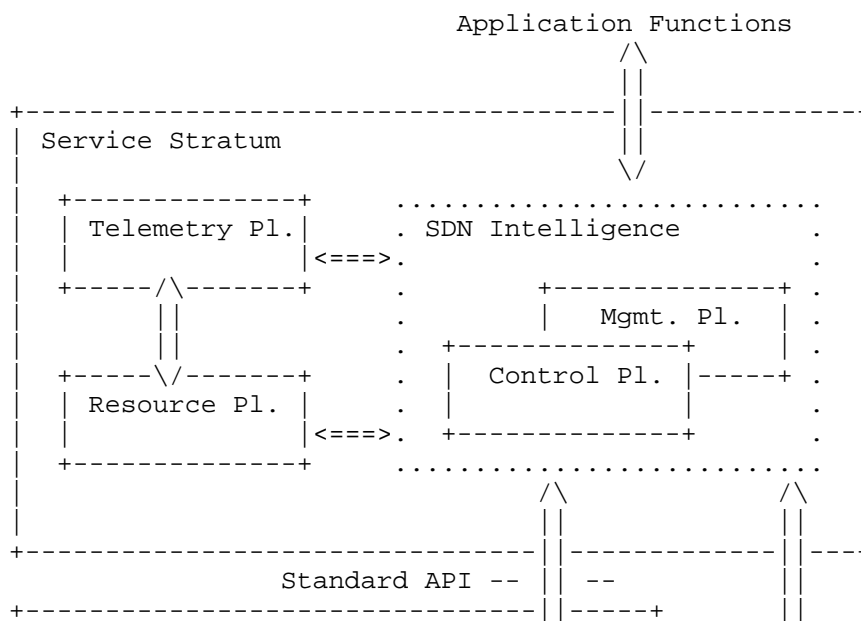
architecture of a quantum network should be composed of:

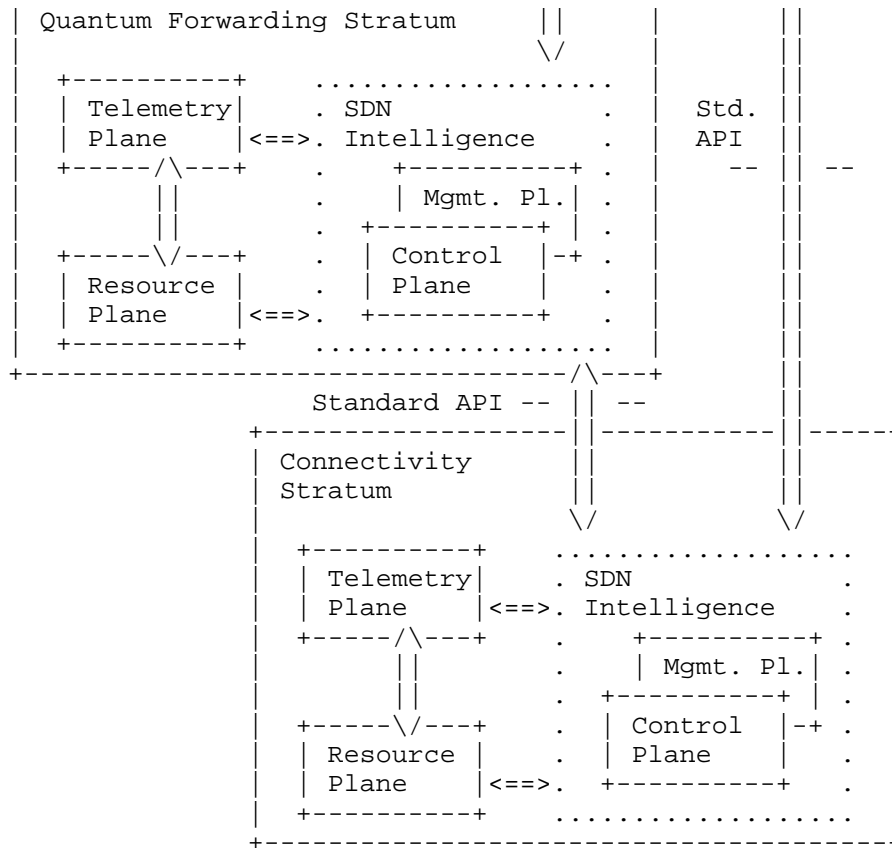
- * A Service Stratum, dealing with the functionality related to the purpose of the quantum network, and aligned with SOP described for QKD networks above. At this moment, the most general service, beyond QKD key management, is obviously entanglement distribution in a general quantum network. Others can be considered, as time synchronization, identity assurance or sensing. The service stratum would consider the relevant service units (keys, shared states, identities, timelines...), deal with their appropriate forwarding and routing, and deliver these service units as requested by the user application functions. The concept of service unit becomes essential here, as the cornerstone for fundamental network characteristics (addressing, routing, information structuring...) and for the interface to the applications using the network. As the discussion on how to identify and relate keys in a wide-area QKD network is still alive, the need to identify how to "pack" qubits in a way useful for, say, distributed computations or teleportation coding, how to route these packs, and how to request and consume services based on them is crucial to define how a global quantum network should be built and operated.
- * A Quantum Forwarding Stratum, in charge of the direct application of quantum protocols and algorithms between the two endpoints of a quantum link, even when it is a multi-hop one, very much as the QFP we described as part of QKD deployments. It is important to note that this stratum must be able to support the service units that constitute, but there is no need for a one-to-one mapping between those quantum forwarding units and the service units. As example, let us consider entanglement forwarding via swapping, which would likely occur on a pairwise basis at this stratum, but needs to be considered in a collective view to make sense to the applications interacting with the service stratum.
- * A Connectivity Stratum, taking care of providing the paths to support the quantum links used by the quantum forwarding and service strata. Typically, the connectivity stratum would be supported by OTN infrastructure, via fiber and/or open-space links, and would follow a common connectivity paradigm, specifically a circuit-based or packet-based one. While current quantum links deal with OTN infrastructure according to a circuit-based paradigm, recent proposals are addressing the idea of "quantum packets" [PSQN22] and the connectivity stratum would have to deal, in general terms, with the classical headers of such packets. Furthermore, classical links are always required for supporting quantum links procedures, and by any kind of monitoring, control, and management connections. The provisioning

of related quantum and classical links, and their consistent operation to meet service levels will be the main concern of this stratum.

This architecture, following the CLAS proposal itself, is built under the assumption that planes within and across strata communicate through well-defined, open interfaces supporting programmability, as a generalization of the common SDN architecture that defines a controller as a mediator between application and network (forwarding) devices. It includes the archetypal case of a centralized controller but is not limited to that particular realization. These broader implications of SDN principles are among the main motivation of the original CLAS proposal in [CLASEVO], and it is the main reason for using it as the base for the framework proposed by this document. The archetypal case of a centralized controller would be the most common deployment style, but the architecture is able to support more distributed approaches, in which each participating domain runs a specific instance of the different strata, providing collaboration by the exposure of tailored information to the other domains via border protocols, as proposed in [ALTOQ24]. Even configurations where a particular domain focuses on one or two of the strata, supporting the other strata being hosted in different domains is also conceivable.

Based on the images used to illustrate the strata proposed in [CLASEVO] and [RFC8597], the relationship among the strata described above would be as shown in the following diagram:





Essentially, this architecture model incorporates the findings from QKD deployments and addresses the requirements for providing a general framework for quantum networks towards the Quantum Internet. It is intended to support the evolution of network base technologies, provide the degrees of freedom necessary to encompass different deployment models, and align with relevant trends in network operation, while considering the practical aspects related to classical connectivity.

The proposed architecture will address the evolution of network base technologies by providing abstractions able to accommodate to this evolution. Considering the stages analyzed in [QIROAD18], the QKD deployment patterns described in the previous section already cover "Trusted Repeater Networks" and "Prepare and Measure Networks", and the general architecture proposed here is able to accommodate the more evolved stages, namely "Entanglement Distribution Networks", "Quantum Memory Networks", "Few Qubit Fault-Tolerant Networks", and "Quantum Computing Networks". As immediate examples we can consider

the integration of features in the Connectivity Stratum with the other two strata to support entanglement forwarding among different locations, or the incorporation of future quantum repeaters into the Quantum Forwarding Stratum to support more elaborated behaviors of the Service Stratum.

In addition, these network abstractions are intended to provide specific degrees of freedom for network design and deployment, through the incorporation of independent resource and control planes at each stratum. Given the control mechanisms identified as "SDN intelligence" on the diagram above are able to expose open interfaces, the approach for coordinating the different strata via mechanisms like those defined in [ETSI18] is totally feasible, and different aggregation patterns (multi-stratum, multi-domain...) and models (federated, hierarchical...) can be applied. These aggregation mechanisms are equally applicable in the case of telemetry data and their integration with closed-loop mechanisms for automation, in support of the required quantum network agility.

The evolved CLAS proposal in [CLASEVO] explicitly incorporates current trends in network automation, in whatever the flavor including AI and intent expressions. This architecture guarantees the future pliability of quantum networks, in alignment with the evolution of best practices in general network management.

Finally, by explicitly addressing the issues related to the connectivity of quantum links, the architecture considers the interactions with any other relevant operational aspects required for providing quantum network services. The direct integration of a stratum focused on these aspects makes the proposed architecture better aligned with the sustainability goal.

4.2. Identification of Interfaces and Protocols

The architecture proposed in this document is intended as a framework to evaluate and explore compatibility among the different proposals on protocols and interfaces for the future availability of quantum features in the global Internet, with the goal of providing a uniform reference model to choose and apply the most appropriate solutions to the Quantum Internet challenges. While the reference architecture does not intend to identify a concrete set of these protocols and interfaces, it is useful to analyze current proposals and trends, and provide some guidance on how the framework can be useful for assessing the integration of the solutions applicable to the different elements that have to converge to realize the Quantum Internet.

There is a significant corpus of standards and operational practices applicable for the Connectivity Stratum, sustained by a well established experience in the management and use of optical and, to some extent, satellite-based networks. The differentiation of the planes considered in the CLAS architecture within the Connectivity Stratum has been common practice in the deployment and operation of IP converged services over optical networks. The abstractions and topology views described in the ACTN framework defined in [RFC8453] constitute a solid foundation to describe the functionality of the planes within the Connectivity Stratum, and the interfaces to be used in the interactions with the other strata. An element like the Path Computation Element (PCE) described in [RFC8637], able to address the considerations related to quantum connectivity and the implications of entanglement-based forwarding, could constitute the core of the intelligence and telemetry planes. Specific forwarding elements, able to fulfill the conditions for quantum signals, including the potential co-propagation with classical signals, and to interface with future quantum repeaters [QREPS], would constitute the essential substrate of the resource plane. The current trends in optical disaggregation and the use of orchestrated SDN mechanisms for network path management and monitoring provide a natural path for leveraging network virtualization mechanisms within the Connectivity Stratum, facilitating their integration.

In what relates to the Quantum Forwarding Stratum, current best practices indicate that telemetry and SDN intelligence planes will follow the same directions as the other strata, with virtualized, likely cloud-native implementations for them. Even in the case of the resource plane, one can expect the availability of specific software agent elements in charge of managing devices, interacting with the Connectivity Plane and providing support to the service units relevant for the Service Stratum. A recent proposal [QUADDR], beyond the foundations described in [RFC9340], can be used to exemplify the main objective of the framework architecture described in this document. The proposal presents quantum-native mechanisms for routing procedures, and the corresponding addressing conventions supporting them, and considers network-wide mechanisms, structured in two tiers defining what could be assimilated to a local domain and an internetworking domain. This proposal can be naturally integrated in the Quantum Forwarding Stratum (QFS), and its SDN-inspired architecture would map the proposed Entanglement-Defined Controller (EDC) at the kernel of the SDN intelligence plane. The integration of an architecture like this within the framework described in this document would require to analyze the mapping between the node identifiers described in the paper and the service units discussed below. The choices for the coordination among the different strata if the QFS uses an architecture like the one proposed in the references paper would need to be also analyzed: on the one hand, the

interface between the EDC and Service Stratum should be defined, and the QFS elements should need to be extended to include its interactions with the Connectivity Stratum, or consider it oblivious to physical connectivity and leave the coordination to the Service Stratum. This is the kind of evaluations the synthetic environments discussed in Section 4.2.2 will be extremely useful.

The discussion on the foundations of the Service Stratum (SS) is made on the following section, where the concept of service units, as already introduced for the case of QKD networks, is analyzed. Furthermore, As a natural consequence of what is discussed above in the framework of cloud-native QKD, the use of network virtualization techniques would be essential for the Service Stratum, at all of their planes:

- * The SDN intelligence plane, allowing the dynamic management of service units and their association with the corresponding units in the Quantum Forwarding Stratum.
- * The telemetry plane, for dynamic monitoring and data aggregation.
- * The resource plane, in support of the different nature of the interactions at the Quantum Forwarding Stratum, like the case of entanglement persistence beyond direct physical reachability.

4.2.1. The Role of Service Units

The fact we remarked above about the QKD service unit being a shared state between two application entities supports a direct translation of the concept to apply it in a generalized quantum network. A service unit in this context will correspond to the shared quantum states to be consumed by the application entities, according to the goals of their sharing of these quantum states. This implies that a QKD service unit can be considered a specialized quantum service unit, where the shared state has been somehow pre-processed to distill the bits that define the shared key. A similar pattern could be applicable to other specialized quantum network applications, as it would be the case of distributed quantum sensing or metrology.

The identification of such service units can follow the same patterns described for the QKD service units, with three URNs, two of them identifying the pair of application entities sharing the state, and a third one defining the lifecycle parameters. Obviously, these parameters should differ from the ones postulated for QKD, and it is possible to envisage parameters such as shared state size (the number of entangled pairs), a timestamp regarding lifetime of the shared state, and others addressing aspects like fidelity. As the quantum memory technology at the foundation of these shared states evolve, a

clearer view of the parameter URN will become available. Experiments on this issue will be really useful to identify these parameters and shape the q-component of the parameter URN.

The content of a QKD service unit is a bitstring corresponding to the shared key. These bitstring is stored in the memories of the corresponding KMEs, with individual bits differentiated by their position in the string. Quantum memories must be available at the resource plane of the Service Stratum (SS), and the service unit should contain the addresses used by those quantum memories to identify the corresponding entangled pairs. The elements equivalent to the KMEs in the control plane of the SS interact with these quantum memories to identify the applicable addresses, and to require the elements in the control plane of the Quantum Forwarding Stratum (QFS) to activate the corresponding exchanges in the quantum links they operate. Each of the endpoints of these quantum links is expected to provide a functionality equivalent to the agents discussed for QKD networks, in support of the SS quantum memories.

As discussed in the case of QKD service units, directionality (the specification of an origin and a destination) is not applicable in this case either, as service units correspond to a shared state. What can be certainly considered is an originator of this shared state, corresponding to the application element requesting the establishment of the service unit. This would trigger the SS control plane element attached to the application to start its route decision procedures and to start the interactions with the relevant SS control planes to start the necessary exchanges to establish the shared quantum states. The nature of the endpoint identifiers support, as discussed for shared keys in QKD, the use of any routing mechanisms, ranging from strictly hierarchical and centralized schemas based on orchestration mechanisms to fully distributed routing algorithms.

As a result of the routing procedures and the interaction among SS control plane elements, there should be corresponding interactions with elements in the control planes of the Connectivity Strata (CS) and the QFS, to verify and require, as needed, the establishment of the individual entangled pairs and, as required, the physical links to support them. There is a consolidated corpus of interfaces (usually known as North-Bound Interfaces, NBI) for the control of classical connectivity, and specially of optical links, such as the TAPI specification [TAPI240], and different proposals to select and establish paths. It seems necessary to explore and experiment with similar interfaces and procedures for the management and control of quantum links, addressing the challenges already identified in [RFC9340] and exploring the implications of quantum-native routing proposals as made in [QUADDR]. A specially significant question is the mapping between the entangled pairs, as identified by the service unit, and the payloads exchanged within the QFS.

Finally, a word on the telemetry planes in each of the proposed strata. It should be obvious the elements in the control planes at each of the strata should start monitoring mechanisms at the involved elements in the resource planes and activate telemetry collection mechanisms. This brings the requirement of defining and experimenting with appropriate metrics and telemetry data models for both the SS and the QFS, as already being defined for QKD infrastructures [ETSI23].

4.2.2. The Role of Synthetic Environments

Due to the early stage of many, if not all, quantum technologies, experimenting with quantum devices and equipment can be seriously hindered by high costs and limited availability. This challenge is particularly evident for experimentation at the scale required to validate network protocols and inter- and intra-strata interfaces. In this context, synthetic environments, and synthetic testbeds enabled by these environments, become an essential tool. They enable the emulation of quantum network deployments in a fully controlled setting, allowing the execution of experiments and trials, protocol evaluations, and even security analyses, where potential network attacks can be tested without compromising the integrity of an already built quantum network or a significant number of physical devices.

Based on the results introduced in [QKNDT24] for QKD networks, the concept of a Quantum Network Digital Twin (QNDT) provides a foundation for such environments. QNDTs will enable a better understanding of the properties of the different network elements, interfaces, and protocols, and the applicability of the architecture proposed in this document. It is important to note that a QNDT is

not a simulation tool, even though some of its components may apply simulation functionality to adapt their behavior to that of a quantum element. Rather, a QNDT represents a distributed classical system that mirrors the operational behavior of a quantum network, responding in real time and accurately reproducing the dynamics and interactions of quantum entities.

In the case of QKD network deployments, significant progress has been achieved thanks to both practical deployments, as exemplified in [EUROQCI] and the early coordinated efforts of standardization bodies. These advances include the definition of standardized APIs that specify the communication means between quantum nodes and customer applications, like [ETSI04], and the integration of network management mechanisms widely adopted in classical communication systems, like the SDN approach defined in [ETSI15]. This coordinated efforts have translated into more flexible, programmable, and scalable control of quantum resources, facilitating seamless interoperability between quantum and classical infrastructures. Despite these advances, several aspects of QKD networking remain under active development. These include the definition of interfaces that ensure interoperability across different administrative domains, as well as the design and validation of architectures capable of supporting large-scale deployments, that is, networks comprising hundreds of interconnected nodes. In this regard, platforms such as the one described in [QUDITTO] offer a valuable opportunity, as they enable the emulation of low-level quantum network behaviors using classical computational resources. Such synthetic environments provide the means to model and analyze complex network scenarios that are currently unattainable in fully physical experimental testbeds.

When considering general-purpose quantum networks, particularly those based on entanglement distribution and management, the role of synthetic environments becomes even more significant. Unlike QKD networks, whose architectural and operational principles are relatively well understood, entanglement-based networks are still in an early stage of development. Many fundamental networking aspects, such as entanglement routing, resource scheduling, and inter-layer coordination, remain open research questions, with a crucial lack of practical validation. In this context, QNDTs offer a unique opportunity to accelerate progress: by enabling controlled emulation of quantum states, interactions, and network behaviors, they allow to test novel architectures, evaluate protocol performance, and explore scalability under realistic yet fully reproducible conditions.

However, the development of a general-purpose QNDT introduces its own set of challenges. Such a system must not only emulate the functional behavior of quantum components but also ensure that the underlying classical infrastructure responds within the same temporal

and operational constraints as its quantum counterpart, thereby enabling accurate validation of protocols and network strategies. Moreover, unlike QKD networks where standardized interfaces and APIs have already been established (or are at least emerging), no equivalent standards currently exist for general quantum networks. Consequently, a QNDT must be designed to be inherently flexible and extensible, capable of accommodating evolving definitions of interfaces, communication protocols, and architectural abstractions. In this regard, the QNDT once again becomes a key enabler for the development, integration, and testing of these foundational elements.

Building upon the above discussion, two primary challenges must be addressed as prerequisites for constructing a fully functional QNDT. First, it is necessary to develop a mechanism capable of handling the quantum-specific aspects of the system, executing simulations and distributing results across nodes, resulting in the emulation of the quantum behavior of network elements within the underlying classical infrastructure. Second, there must be a definition of a minimal set of core primitives or instructions that serves as the foundation for constructing more advanced mechanisms, such as standardized interfaces and communication methods between network elements and external systems. Together, these two pillars will establish the groundwork for a QNDT framework capable of evolving in parallel with the broader quantum networking ecosystem.

The core quantum emulation mechanism for such an environment, according to the current state of the art, would be the QNDT emulation engine, based on a centralized simulation component designed to execute the simulations needed to emulate the quantum behavior of all network elements. This engine may rely on quantum network simulators such as [NetSquid], [SeQUeNCe], or [QuNetSim]. However, these platforms alone do not fulfill the requirements of a QNDT, since, as discussed above, a QNDT is not a simulation of the network but a distributed classical system that replicates the behavior of a real quantum network. Therefore, the central simulation element must be complemented by a result distribution mechanism, for example, through a publish/subscribe (Pub/Sub) protocol. In such a setup, network elements subscribe to topics relevant to their operation and can communicate with the central simulation tool both to request simulations and receive results through asynchronous interactions.

Another essential aspect concerns the handling of temporal consistency between the “simulation time”, i.e., the time required to execute a simulation, and the “simulated time,” i.e., the time the simulation calculates the real system would take to perform the same operation. Since simulation time is generally shorter than simulated time, the QNDT must incorporate logic ensuring that results are

delivered only after the appropriate simulated delay has elapsed. This guarantees that the QNDT responds within the same temporal boundaries as its physical counterpart, thereby preserving the fidelity and realism of the emulated network behavior.

In addition, to maintain state realism within the QNDT, it is crucial to take into account the natural decoherence and noise dynamics of quantum states over time. For instance, when entangled pairs are distributed between two nodes and stored for a period before being used in subsequent operations, the QNDT must emulate the gradual evolution and degradation of these states. This entails tracking the elapsed time between state creation and use, and updating the state accordingly before executing the next instruction.

5. Security Considerations

The general considerations made in [RFC8597] apply, as well as an elaboration on the following points regarding:

- * The requirements on mutual authentication in the channels used for quantum interactions, as they should require methods rooted at physical properties.
- * Specific physical attacks related to the particular quantum mechanisms in use by the quantum forwarding stratum.
- * The interaction of these physical attacks with classical attacks to the control and monitoring activities, possibly translating into a threat surface augmentation.

Furthermore, as the identification of interfaces and protocols progresses other considerations will be required. In particular, the security considerations included in the documents referenced for the Connectivity Stratum, [RFC8453] and [RFC8637] apply to the proposed framework.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC8141] Saint-Andre, P. and J. Klensin, "Uniform Resource Names (URNs)", RFC 8141, DOI 10.17487/RFC8141, April 2017, <<https://www.rfc-editor.org/rfc/rfc8141>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.
- [RFC8597] Contreras, LM., Bernardos, CJ., Lopez, D., Boucadair, M., and P. Iovanna, "Cooperating Layered Architecture for Software-Defined Networking (CLAS)", RFC 8597, DOI 10.17487/RFC8597, May 2019, <<https://www.rfc-editor.org/rfc/rfc8597>>.

6.2. Informative References

- [ALTOQ24] Muniz, A., Canto, R., Contreras, L., Pastor, A., Lopez, D., and J. Morales, "Using Protocol to Address SD-QKD Federation in Multi-Domain Scenarios", July 2024, <<https://ieeexplore.ieee.org/document/10628176>>.
- [CLASEVO] Contreras, L. M., Boucadair, M., Lopez, D., and C. J. Bernardos, "An Evolution of Cooperating Layered Architecture for SDN (CLAS) for Compute and Data Awareness", Work in Progress, Internet-Draft, draft-contreras-coinrg-clas-evolution-03, 5 July 2024, <<https://datatracker.ietf.org/doc/html/draft-contreras-coinrg-clas-evolution-03>>.
- [ETSI04] "ETSI GS QKD 004: Quantum Key Distribution (QKD); Application Interface", August 2020, <https://www.etsi.org/deliver/etsi_gs/QKD/001_099/004/02.01.01_60/gs_QKD004v020101p.pdf>.
- [ETSI14] "ETSI GS QKD 014: Quantum Key Distribution (QKD); Protocol and data format of REST-based key delivery API", February 2019, <https://www.etsi.org/deliver/etsi_gs/QKD/001_099/014/01.01.01_60/gs_qkd014v010101p.pdf>.
- [ETSI15] "ETSI GS QKD 015: Quantum Key Distribution (QKD); Control Interface for Software Defined Networks", April 2022, <https://www.etsi.org/deliver/etsi_gs/QKD/001_099/015/02.01.01_60/gs_QKD015v020101p.pdf>.
- [ETSI18] "ETSI GS QKD 018: Quantum Key Distribution (QKD); Orchestration Interface for Software Defined Networks", April 2022, <https://www.etsi.org/deliver/etsi_gs/QKD/001_099/018/01.01.01_60/gs_QKD018v010101p.pdf>.

- [ETSI23] "ETSI Work-Item QKD 023: Quantum Key Distribution (QKD); Monitoring Interface and Data Model", n.d., <https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=69537>.
- [EUROQCI] "The European Quantum Communication Infrastructure (EuroQCI) Initiative", September 2023, <<https://digital-strategy.ec.europa.eu/en/policies/european-quantum-communication-infrastructure-euroqci>>.
- [EVCK25] Lopez, B., Vidal, I., Valera, F., and D. Lopez, "An Enhanced Virtualized Control and Key Management Model for QKD Networks", January 2025, <<https://ieeexplore.ieee.org/document/10870375>>.
- [MADQCI23] Martin, V., Brito, J. P., Ortiz, L., Brito-Mndez, R., Vicente, R., Saez-Buruaga, J., Sebastian, A. J., Aguado, D. G., Garca-Cid, M. I., Setien, J., Salas, P., Escribano, C., Dopazo, E., Rivas-Moscoso, J., Pastor-Perales, A., and D. Lopez, "The Madrid Testbed: QKD SDN Control and Key Management in a Production Network", July 2023, <<https://ieeexplore.ieee.org/document/10207295>>.
- [NetSquid] Coopmans, T., Knegjens, R., Dahlberg, A., Maier, D., Nijsten, L., Filho, J. de O., and et. al., "NetSquid, a NETwork Simulator for QUantum Information using Discrete events", July 2021, <<https://doi.org/10.1038/s42005-021-00647-8>>.
- [NFV06] "ETSI GS NFV 006: Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Architectural Framework Specification", December 2022, <https://www.etsi.org/deliver/etsi_gs/NFV/001_099/006/04.04.01_60/gs_NFV006v040401p.pdf>.
- [PSQN22] DiAdamo, S., Qi, B., Miller, G., Kompella, R., and A. Shabani, "Packet switching in quantum networks: A path to the quantum Internet", October 2022, <<https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.4.043064>>.
- [QIPS22] Illiano, J., Caleffia, M., Manzalini, A., and A. S. Cacciapuoti, "Quantum Internet Protocol Stack: a Comprehensive Survey", August 2022, <<https://www.sciencedirect.com/science/article/abs/pii/S1389128622002250>>.

- [QIROAD18] Wehner, S., Elkouss, D., and R. Hanson, "Quantum internet: A vision for the road ahead", October 2018, <<https://doi.org/10.1126/science.aam9288>>.
- [QKNDT24] Martin, R., Lopez, B., Vidal, I., Valera, F., and B. Nogales, "Service for Deploying Digital Twins of QKD Networks", January 2024, <<https://doi.org/10.3390/app14031018>>.
- [QNSA24] Lopez, B., Vidal, I., Valera, F., Lopez, D., and A. Pastor, "Unleashing Flexibility and Interoperability in QKD Networks: The Power of Softwarized Architectures", July 2024, <<https://ieeexplore.ieee.org/document/10628345>>.
- [QREPS] "Quantum repeaters: From quantum networks to the quantum internet", December 2023, <<https://doi.org/10.1103/RevModPhys.95.045006>>.
- [QTTI21] Martin, V., Brito, J. P., Escribano, C., Menchetti, M., White, C., Lord, A., Wissel, F., Gunkel, M., Gavignet, P., Genay, N., Moulton, O. L., Abellan, C., Manzalini, A., Pastor-Perales, A., Lopez, V., and D. Lopez, "Quantum Technologies in the Telecommunications Industry", July 2021, <<https://epjquantumtechnology.springeropen.com/articles/10.1140/epjqt/s40507-021-00108-9>>.
- [QUADDR] Caleffi, M. and A. S. Cacciapuoti, "Quantum Internet Architecture: unlocking Quantum-Native Routing via Quantum Addressing", July 2025, <<https://doi.org/10.48550/arXiv.2507.19655>>.
- [QUDITTO] "Quditto, a tool that allows deploying digital twins of QKD networks over classical infrastructure", April 2025, <<https://quditto.io/>>.
- [QuNetSim] Diadamo, S., Ntzel, J., Zanger, B., and M. M. Bee, "QuNetSim: A Software Framework for Quantum Networks", June 2021, <<https://doi.org/10.1109/TQE.2021.3092395>>.
- [RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <<https://www.rfc-editor.org/rfc/rfc8453>>.

- [RFC8637] Dhody, D., Lee, Y., and D. Ceccarelli, "Applicability of the Path Computation Element (PCE) to the Abstraction and Control of TE Networks (ACTN)", RFC 8637, DOI 10.17487/RFC8637, July 2019, <<https://www.rfc-editor.org/rfc/rfc8637>>.
- [RFC9340] Kozlowski, W., Wehner, S., Van Meter, R., Rijsman, B., Cacciapuoti, A. S., Caleffi, M., and S. Nagayama, "Architectural Principles for a Quantum Internet", RFC 9340, DOI 10.17487/RFC9340, March 2023, <<https://www.rfc-editor.org/rfc/rfc9340>>.
- [RFC9583] Wang, C., Rahman, A., Li, R., Aelmans, M., and K. Chakraborty, "Application Scenarios for the Quantum Internet", RFC 9583, DOI 10.17487/RFC9583, June 2024, <<https://www.rfc-editor.org/rfc/rfc9583>>.
- [SeQUeNCe] Wu, X., Kolar, A., Chung, J., Jin, D., Zhong, T., Kettimuthu, R., and M. Suchara, "SeQUeNCe: A Customizable Discrete-Event Simulator of Quantum Networks", September 2020, <<https://doi.org/10.1088/2058-9565/ac22f6>>.
- [TAPI240] "ONF Transport API SDK 2.4.0", n.d., <<https://github.com/Open-Network-Models-and-Interfaces-ONMI/TAPI/releases/tag/v2.4.0>>.
- [Y3802] "ITU-T Recommendation Y.3802: Quantum key distribution networks. Functional architecture", April 2021, <<https://www.itu.int/rec/T-REC-Y.3802>>.

Acknowledgments

This document is based on work partially funded by the EU Horizon Europe project QSNP (grant 101114043), the Spanish UNICO project OPENSEC (grant TSI-063000-2021-60), and the MadridQuantumCM project (funded by the EU, NextGenerationEU, grant PRTR-C17.I1, and by the Comunidad de Madrid, Programa de Acciones Complementarias).

Authors' Addresses

Diego Lopez
Telefonica
Email: diego.r.lopez@telefonica.com

Vicente Martin
UPM
Email: vicente.martin@upm.es

Blanca Lopez
IMDEA Networks
Email: blanca.lopez@imdea.org

Luis M. Contreras
Telefonica
Email: luismiguel.contrerasmurillo@telefonica.com